ENCYCLOPEDIA OF CONTROLLER FUNDAMENTALS AND FEATURES FIRMWARE VERSION 3.2 TO 8.0

Part Number 775041-01 Rev C.
8/16/01

NOTICE

Although reasonable efforts have been made to ensure the accuracy of the information contained herein, this publication could include technical inaccuracies or typographical errors. IBM Corporation expressly disclaims liability for any error in this information, and for damages, whether direct, indirect, special, exemplary, consequential or otherwise, that may result from such error, including but not limited to loss of profits resulting from the use or misuse of the manual or information contained therein (even if IBM Corporation has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to:

Mylex, an IBM Company
Department 6IMA - Information Development
5400 Airport Boulevard, Suite 100
Boulder, Colorado 80301

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law:

IBM CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties or the limitation or exclusion of liability for indirect, special, exemplary, incidental or consequential damages in certain transactions; therefore, this statement may not apply to you. Also, you may have other rights which vary from jurisdiction to jurisdiction.

Information in this publication is subject to change without notice and does not represent a commitment on the part of IBM Corporation. Changes may be made periodically to the information herein; these changes will be incorporated in new editions of the publication. IBM Corporation reserves the right to make improvements and/or changes at any time in the product(s) and/or program(s) described in this publication.

It is possible that this publication may contain reference to, or information about, IBM Corporation products (machines and programs), programming or services that are not announced in your country. Such references or information must not be construed to mean that IBM Corporation intends to announce, provide, or make available such Mylex Products, programming, or services in your jurisdiction.

Requests for technical information about IBM Corporation products should be made to your Mylex authorized reseller or IBM representative.

Unless otherwise noted data used in the examples in this manual are fictitious.

About This Manual

This reference manual provides general information about features incorporated in the DACSX, DACSS, DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers operating with firmware versions 3.2 through 8.0. The information is presented in an encyclopedia format, with the feature descriptions listed alphabetically. A second section provides descriptions of the currently supported controller parameters. This manual is not intended to provide step-by-step operating instructions for all procedures, however, the third section provides instructions for a few procedures not described elsewhere. The last section is a glossary of definitions for commonly used terminology. For detailed information on installing and using any of the Mylex controllers, refer to the appropriate installation guide or user manual for the controller or configuration utility you are using.

This revision of the Encyclopedia incorporates all firmware releases in an effort to provide a complete history of the controllers and the firmware. Each feature lists the firmware version in which the feature was introduced and all products that support that particular feature. Some features were introduced per customer request and are not available in subsequent firmware releases; and some features are not available on all controllers. If you are interested in a particular feature that is not available on your controller, please see your IBM representative.

Note: Firmware versions 7.0 and 7.2 require GAM version 3.1 or newer. Firmware versions 7.4 and greater require SAM/GAM version 4.0 or newer.

Summary of Changes

The following changes have been implemented since the last revision of this manual.

- The Encyclopedia of Controller Fundamentals and Features has been reorganized. The manual now has four sections: “Feature Descriptions” on page 1, “Controller Parameters” on page 69, “Frequently Used Procedures” on page 91, and “Terminology Definitions” on page 105.
- Reboot was globally replaced with restart.
- The Summary of Changes section was added to the front matter. This section will be updated with each revision to the document.
- “Controller Specifications” on page 1 and “Product Supported Features” on page 2 have been updated.
- The feature descriptions for “Alternate Path Software” on page 7, “Automatic Rebuild” on page 10, “Background Initialization” on page 11, “Battery Backup Unit (BBU)” on page 12, “Configuration Tools and Utilities” on page 21, “Drive Channel Failover and Failback” on page 28, “Drive Sizing” on page 30, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Firmware Variables” on page 45, “Replacement Controller Behavior with Forced Simplex Enabled” on page 53, and “SANmapping™” on page 56 have been modified.
- The feature descriptions for “Automatic Firmware Flash” on page 7, “Default Mode Page Settings” on page 26, “Embedded Configuration Utility” on page 35, “FFx2” on page 43, “World Wide Name (WWN) Assignments” on page 65, and “World Wide Name (WWN) Table” on page 66 have been added.
- The “Enable Background Initialization” on page 78 controller parameter has been added.
- The procedure descriptions for “Firmware Versions” on page 95 and “Replacing a Failed Controller in Existing Duplex Systems” on page 97 have been modified.
- The “Acronym List” on page 105 has been updated.
- The definitions for “Disk Loop” on page 115, “Disk Port” on page 115, “Firmware Flash” on page 119, and “Loop Redundancy Circuit (LRC)” on page 128 have been added.
- The definitions for “Gigabyte (GB)” on page 120, “Kilobyte (KB)” on page 123, “Megabyte (MB)” on page 129, “RAID Levels” on page 137, and “Terabyte (TB)” on page 150 have been modified.

Conventions

The following fonts indicate additional useful information or situations where special care is required:

Note: Text marked as a Note indicates parenthetical information that may be helpful.

ATTENTION: **Indicates information that may prevent you from losing data or damaging your equipment.**

Related Documentation

The following documents are available from IBM Corporation and provide additional information.

OEM System Reference Manual, PN 771992
OEM Firmware/Software Interface Reference Manual, PN 775067
DACSX OEM Hardware Reference Manual, PN 771974
DACSF OEM Hardware Reference Manual, PN 771991
DACFL OEM Hardware Reference Manual, PN 775019
DACSS OEM Hardware Reference Manual, PN 775023
DACFF OEM Hardware Reference Manual, PN 775026
FFx OEM Hardware Reference Manual, PN 775039
FF2 OEM Hardware Reference Manual, PN 775060
FFx2 OEM Hardware Reference Manual, PN 775094
DAC960SX Installation Guide, PN 771990
DACSF Installation Guide, PN 775016
DACFL Installation Guide, PN 775028
DACFF Installation Guide, PN 775032
DACFF Quick Installation and Configuration Guide, PN 775063
FFx Quick Installation and Configuration Guide, PN 775047
FF2 Quick Installation and Configuration Guide, PN 775061
FFx2 Quick Installation and Configuration Guide, PN 775096
DAC960SX Family User Guide, PN 771975
RAIDfx Manager User Guide, PN 771989
Global Array Manager Installation Guide and User Manual, PN 771961
Global Array Manager Server Installation Guide and User Manual, PN 08P4604
SANArray Manager Client Software for Mylex External Disk Array Controllers Installation Guide and User Manual, PN 775084
Terminal-Emulation (VT100) User Guide and LCD Front Panel Status Guide, PN 775069
Embedded Configuration Utility User’s Guide, PN 775107

CONTENTS

FEATURE DESCRIPTIONS

This section provides the theoretical and reference information describing features of the Mylex disk array controllers. The information is presented as an alphabetical listing of supported features. The information applies to all Mylex controllers, unless otherwise noted. For detailed information on installing and using any of the Mylex controllers, refer to the appropriate installation guide or user manual for the controller or configuration utility you are using.

Controller Specifications

Each Mylex controller has unique specifications. The following table describes specification details for each controller.

Controller Specifications

	Host Support	Drive Type	Host Channel Transfer Rate	Drive Channel Transfer Rate	# Host Ports	# Drive Channels	# Drives per Channel	Total # Drives Supported	FW Version Supported	Cache Memory
DACSX without Daughter Card	Single-Ended or Differential SCSI	Single-Ended SCSI	Up to 40 MB/s	Up to 40 MB/s	1	2	15 (Simplex) 14 (Duplex)	30 (Simplex) 28 (Duplex)	3.2 — 4.2	8 MB to 128 MB
DACSX with Daughter Card	Single-Ended or Differential SCSI	Single-Ended SCSI	Up to 40 MB/s	Up to 40 MB/s	1/2	5/4	15 (Simplex) 14 (Duplex)	75/60 (Simplex) 70/56 (Duplex)	3.2 — 4.2	8 MB to 128 MB
DACSS	Single-Ended or Differential SCSI	Single-Ended SCSI	Up to 40 MB/s	Up to 40 MB/s	2	4	15 (Simplex) 14 (Duplex)	60 (Simplex) 56 (Duplex)	4.4 only	32 MB to 256 MB
DACSF	Fibre	Single-Ended SCSI	100 MB/s	Up to 40 MB/s	2	6	15 (Simplex) 14 (Duplex)	90 (Simplex) 84 (Duplex)	4.0—4.2	32 MB to 256 MB
DACFL	Fibre	Low-Voltage Differential SCSI	100 MB/s	Up to 80 MB/s	2	4	15 (Simplex) 14 (Duplex) 16 (>5.0)	60 (Simplex) 56 (Duplex) 64 (>5.0)	4.2—6.3	32 MB to 256 MB
DACFF (FF)	Fibre	Fibre	100 MB/s	100 MB/s	2	4	16 (<7.0) 124 (>7.0)	64 (<7.0) 496 (>7.0)	5.0 — 7.2	32 MB to 1 GB
FFx	Fibre	Fibre	100 MB/s	100 MB/s	1	2	16 (<7.0) 124 (>7.0)	32 (<7.0) 248 (>7.0)	6.0 — 7.7	64 MB to 512 MB[a]
FF2	Fibre	Fibre	100 MB/s	100 MB/s	2	4	16 (<7.0) 124 (>7.0)	64 (<7.0) 496 (>7.0)	6.0 — 7.7	32 MB to 1 GB
FFx2	Fibre 2 Gbps	Fibre 2 Gbps	200 MB/s	200 MB/s	2	2	124	248	8.0	128 MB to 512 MB

a. Cache Memory sizes 256 MB and 512 MB are not supported with firmware version 7.0; however both sizes are supported with subsequent versions.

➢ See also "DACFF (FF)" on page 24, "DACFL" on page 24, "DACSF" on page 24, "DACSS" on page 25, "DACSX" on page 25, "Drive Channel" on page 27, "FF2" on page 42, "FFx" on page 42, "FFx2" on page 43, "Firmware Versions" on page 95, and "Host Channel" on page 121.

Product Supported Features

The following table presents a summary of all of the features, at which version of firmware the feature was implemented, and on which products the feature is supported.

Product Supported Features

Feature		Implemented	DACSX	DACSS	DACSF	DACFL	DACFF (FF)	FF2	FFx	FFx2
AEMI		pre 3.2	X							
Alarm Signal (Modified with 7.0 and 7.4)		pre 3.2	X	X	X	X	X	X	X	X
Alternate Path Software		6.0				X	X	X	X	X
Automatic Firmware Flash		8.0							X	X
Automatic Restart (Reboot) on Failure		7.4						X	X	X
Automatic Rebuild		pre 3.2	X	X	X	X	X	X	X	X
Background Initialization (Initially supported with 6.3 only, updated and re-introduced with 7.7)		6.3 and 7.7				X		X	X	X
Battery Backup Unit (BBU) and Advanced BBU		pre 3.2	X						X	X
Cache Coherency		5.4				X	X	X	X	X
Cache Line Size		Pre 3.2	X	X	X	X	X	X	X	X
Cache Mirroring		Pre 3.2	X	X	X	X	X	X	X	X
Configuration Coherency		Pre 3.2	X	X	X	X	X	X	X	X
Configuration on Disk (COD) (Modified with 5.4, 7.0, and 8.0)		pre 3.2	X	X	X	X	X	X	X	X
Configuration Tools and Utilities		pre 3.2	X	X	X	X	X	X	X	X
Conservative Cache Mode (Modified with 4.2, 6.1, and 7.4)		3.3	X	X	X	X	X	X	X	X
Consistency Check (See Parity Check)		pre 3.2	X	X	X	X	X	X	X	X
Controller-Controller Nexus (C-C Nexus) (Modified for FFx)		pre 3.2	X	X	X	X	X	X	X	X
Controller Event Logging		7.0					X	X	X	X
Controller Fault Indication		7.4						X	X	X
Controller Parameters		pre 3.2	X	X	X	X	X	X	X	X
	Automatic Failback or Autorestore	pre 3.2	X	X	X	X	X	X	X	X
	Automatic Restart (Reboot) on Failure Parameter	7.4						X	X	X
	Automatic Rebuild Management	pre 3.2	X	X	X	X	X	X	X	X
	Enable Background Initialization	7.7						X	X	X
	Block Size	pre 3.2	X	X	X	X	X	X	X	X
	Blocking Factor	7.4						X	X	X
	Cache Line Size Parameter	pre 3.2	X	X	X	X	X	X	X	X
	Command Tagging	3.2 - 4.2	X	X	X	X				
	Conservative Cache Mode Parameter	pre 3.2	X	X	X	X	X	X	X	X

Product Supported Features (Continued)

Feature		Implemented	DACSX	DACSS	DACSF	DACFL	DACFF (FF)	FF2	FFx	FFx2
	Controller Name	7.0					X	X	X	X
	Controller Present/Fault Select	pre 3.2	X							
	Data Bus Width	3.2 - 6.3	X	X	X	X				
	Device Combing (also referred to as Coalescing)	pre 3.2	X	X	X	X	X	X	X	X
	Disable BUSY Status During Failback	7.0					X	X	X	X
	Disable Check Condition for Invalid LUN	pre 3.2	X	X	X	X	X	X	X	X
	Disable Queue Full Status	pre 3.2	X	X	X	X	X	X	X	X
	Disconnect on First Command	3.2 - 4.2	X	X	X	X				
	Disk Startup Delay 1(also referred to as Spin-up Disk Delay)	pre 3.2	X	X	X	X	X	X	X	X
	Disk Startup Delay 2 (also referred to as Spin-up Sequence Delay)	pre 3.2	X	X	X	X	X	X	X	X
	Disk Startup Mode (also referred to as Spin-up Option) (Modified with 7.4)	pre 3.2	X	X	X	X	X	X	X	X
	Disk Startup Number of Devices (also referred to as Disks Per Spin)	pre 3.2	X	X	X	X	X	X	X	X
	Disk Write Through Verify (also referred to as Write Through Verify)	pre 3.2	X	X	X	X	X	X	X	X
	Duplex Fault Signals (also referred to as Controller Present/Fault Signals)	pre 3.2	X	X	X	X	X	X		
	Duplex Fault Signals on Channel 4	3.2 - 4.2	X							
	Elevator	3.2 - 4.2	X	X	X	X				
	Failover Topologies	4.0			X	X	X	X	X	X
	Force Simplex	pre 3.2	X	X	X	X	X	X	X	X
	Frame Size	4.0			X	X	X	X	X	X
	Generate Debug Dump	7.4						X	X	X
	Hard Loop IDs	4.0			X	X	X	X	X	X
	Host (SCSI) Bus Reset Delay	3.2 - 4.2	X	X						
	Host Transfer Width (also referred to as Disable Wide Operation)	3.2 - 4.2	X	X						
	Maximum IOPs	3.2 - 4.2	X	X	X	X				
	Mode Select SRA Enable	pre 3.2	X	X	X	X				
	No Pause on Controller Not Ready	pre 3.2	X	X	X	X	X	X	X	X
	Node Name Retention	4.0			X	X	X	X	X	X
	On Queue Full Give Busy (also referred to as Disable Queue Full Status)	pre 3.2	X	X	X	X	X	X	X	X
	Operational Fault Management	pre 3.2	X	X	X	X	X	X	X	X
	Override Multiport Reset	4.0			X	X	X	X	X	X
	PCI Latency Control	4.0 - 5.4			X	X	X	X		
	Queue Limit	3.2 - 4.2	X	X	X	X				
	RAID5 Algorithm	pre 3.2	X	X	X	X	X	X	X	X
	Read Ahead	pre 3.2	X	X	X	X	X	X	X	X

Product Supported Features (Continued)

Feature		Implemented	DACSX	DACSS	DACSF	DACFL	DACFF (FF)	FF2	FFx	FFx2
	Read/Write Control	pre 3.2	X	X	X	X	X	X	X	X
	Reassign Restricted to One Block	pre 3.2	X	X	X	X	X	X	X	X
	Rebuild and Check Consistency Rate	pre 3.2	X	X	X	X	X	X	X	X
	Reset Propagation	3.2 - 4.2	X	X						
	SAF-TE Data for UPS Support (also referred to as Use of UPS)	pre 3.2	X	X	X	X	X	X	X	X
	Serial Control	pre 3.2	X	X	X	X	X	X	X	X
	Serial Port Baud Rate	pre 3.2	X	X	X	X	X	X	X	X
	Serial Port Debug Type	pre 3.2	X	X	X	X	X	X	X	X
	Serial Port Usage	pre 3.2	X	X	X	X	X	X	X	X
	Simplex - No Reset	3.2 - 7.7	X	X	X	X	X	X		
	Smart Large Host Transfers	4.2	X	X	X	X	X	X	X	X
	Stripe Size Parameter	3.2 - 6.3	X	X	X	X	X	X	X	
	Super Read Ahead	pre 3.2	X	X	X	X	X	X	X	X
	Transfer Speed (also referred to as SCSI Transfer Rate or Mega-transfers/sec)	3.2 - 4.2	X	X	X	X				
	Transfer Width	3.2 - 4.2	X	X	X	X				
	True Verify	pre 3.2	X	X	X	X	X	X	X	X
	Vendor Unique Direct Command	pre 3.2	X	X	X	X	X	X	X	X
	Vendor Unique Pass-Through Command	pre 3.2	X	X	X	X	X	X	X	X
	Vendor Unique TUR	pre 3.2	X	X	X	X	X	X	X	X
Debug Dump		7.4						X	X	X
Default Mode Page Settings		7.7						X	X	X
Disk Media Error Management		pre 3.2	X	X	X	X	X	X	X	X
Drive Channel (varies with product and firmware)		pre 3.2	X	X	X	X	X	X	X	X
Drive Channel Failover and Failback (Modified with 7.75)		6.0					X	X	X	X
Drive Packs / Groups (Modified with 7.0)		pre 3.2	X	X	X	X	X	X	X	X
Drive Roaming		4.2	X	X	X	X	X	X	X	X
Drive Sizing (Modified with 7.0 and 7.4)		3.3	X	X	X	X	X	X	X	X
Drive State Management (Modified with 7.0)		pre 3.2	X	X	X	X	X	X	X	X
Drive to Enclosure Map		7.4						X	X	X
Dual-Active Controller Configuration (Not supported by all firmware releases, some simplex only)		pre 3.2	X	X	X	X	X	X	X	X
Dual-Ported Drive Support		6.0					X	X	X	X
Embedded Configuration Utility		7.7						X	X	X
Failover and Failback		pre 3.2	X	X	X	X	X	X	X	X
Failover Topology, Inactive Port		4.2			X	X	X	X		
Failover Topology, Master/Slave		5.0 - 5.4			X	X	X			
Failover Topology, Multiple Target ID		6.0						X	X	X
Failover Topology, Multiport		5.0			X	X	X	X	X	X

Product Supported Features (Continued)

Feature	Implemented	DACSX	DACSS	DACSF	DACFL	DACFF (FF)	FF2	FFx	FFx2
Fault Management (Support changed from AEMI to SAF-TE to SES)	3.2	X	X	X	X	X	X	X	X
Fibre Channel	4.0			X	X	X	X	X	X
Firmware Header Information	3.2	X	X	X	X	X	X	X	X
Firmware / Software Interface	7.0					X	X	X	X
Firmware Variables	7.4						X	X	X
Global Array Manager (GAM)	pre 3.2	X	X	X	X	X	X	X	X
Implied Reserve and Release	6.1				X	X	X	X	X
Intelligent Cache Mirroring	5.4				X	X	X	X	X
Liquid Crystal Display (LCD)	3.2 - 5.0	X	X	X	X	X			
LUN Mapping	pre 3.2	X	X	X	X	X			
Microsoft Cluster Server Configurations (MSCS)	4.2	X		X	X				
Multiple Target ID (MTID) (Modified with 6.1)	pre 3.2	X	X				X	X	X
Mylex Online RAID Expansion (MORE)	3.3	X	X	X	X	X	X	X	X
Online Spare Polling	7.4						X	X	X
Parallel Host Writes	5.4				X	X	X	X	X
Parity Check (See Consistency Check)	pre 3.2	X	X	X	X	X	X	X	X
Predictive Failure Analysis™ (PFA)	7.4						X	X	X
Programmable LUN Mapping	3.3	X	X	X	X	X			
Protected Boot ROM (PBR)	7.0					X	X	X	X
RAIDfx	3.2 - 6.3	X	X	X	X	X	X	X	
Replacement Controller Behavior with Forced Simplex Enabled	6.0				X	X	X	X	X
Reset Controllers (Modified with 4.2 and 7.2)	pre 3.2	X	X	X	X	X	X	X	X
SAF-TE	3.2	X	X	X	X				
SANArray Manager™ (SAM)	7.4	X	X	X	X	X	X	X	X
SANmapping™	6.0				X	X	X	X	X
SES (Modified with 6.1, 6.2, 7.0, and 7.4)	5.0					X	X	X	X
Self-Monitoring, Analysis and Reporting Technology (SMART)	7.4						X	X	X
Simplex	pre 3.2	X	X	X	X	X	X	X	X
Single Point of Failure (SPOF)	5.4				X				
Soft Addressing Detection	7.4						X	X	X
Spanning Drive Packs (Modified with 7.0)	pre 3.2	X	X	X	X	X	X	X	X
Standard Data Caching	3.3	X	X	X	X	X	X	X	X
Stripe Size (Modified with 7.0)	pre 3.2	X	X	X	X	X	X	X	X
System Drive Affinity	3.2 - 6.1	X	X	X	X	X	X	X	
UNIX Support and Firmware Types	pre 3.2	X	X	X	X	X	X	X	X
VT100 (Modified with 7.0)	pre 3.2	X	X	X	X	X	X	X	X
World Wide Name Assignments (Modified with 7.4 and 7.7)	7.0					X	X	X	X
World Wide Name (WWN) Table (Modified with 7.7 and 8.0)	6.0				X	X	X	X	X

➢ See also “Controller Parameters” on page 69, “DACFF (FF)” on page 24, “DACFL” on page 24, “DACSF” on page 24, “DACSS” on page 25, “DACSX” on page 25, “FF2” on page 42, “FFx” on page 42, “FFx2” on page 43, and “Firmware Versions” on page 95.

AEMI

SCOPE: This feature was implemented in firmware version 3.2 and is supported only on the DACSX controller.

Array Enclosure Management Interface (AEMI) monitors power supplies, fans, and enclosure temperature. AEMI is supported only on DACSX controllers, but DACSX controllers also support SAF-TE interfaces. If there is no SAF-TE device installed, the DACSX controller uses the AEMI interface for handling faults and error conditions external to the controller. AEMI provides the controller with information on multiple fans and power supplies, as well as reporting any out-of-limit temperature condition. AEMI also provides the controller with information on disk drives for each of the drive channels.

➢ See also “Conservative Cache Mode” on page 22, “Drive State Management” on page 31, “Environmental Device” on page 117, “Fault Management” on page 41, “SAF-TE” on page 55, and “SES” on page 58.

Alarm Signal

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers. The feature was modified with firmware version 7.0. The modifications included additional support for hardware changes and requirements. Firmware version 7.4 provided the ability to enable and disable an enclosure alarm for any operation that enables the alarm.

NOTE: The exact length of an ALARM is indeterminate. Several alarm signals continue until all conditions causing the alarm have been remedied.

An alarm signal from the controller is asserted under the following conditions.

- An alarm sounds when the primary controller status message is updated.
- An alarm sounds when the partner controller status message is updated.
- An alarm sounds when the user selects an invalid key when using the LCD Front Panel keypad user interface. (The LCD Front Panel keypad is not available on all products and is supported with firmware versions 3.2 through 6.3.)
- An alarm sounds when a physical disk drive state changes from online to offline failed (dead).
- An alarm sounds when the controller is starting. A series of two momentary alarm assertions is used to indicate progress through the early firmware start sequence. Failure to start without completing the two alarm assertions indicates that the controller is malfunctioning and must be returned to Mylex for repair.
- An alarm sounds in all enclosures that house physical disk drives configured as part of a disk pack that is associated with a critical system drive.

The following conditions were implemented in firmware version 7.0.

- An alarm sounds when a system drive is online critical and no online spares are available. The alarm continues until the condition is repaired (a rebuild completes or the drive is replaced/set back online).

- An alarm sounds in the enclosure housing a spare disk drive when it is selected for a rebuild operation.
- An alarm sounds if the SES support becomes critical.
- An alarm sounds if there is a problem in the power supply unit.
- An alarm sounds if an over temperature condition occurs.

NOTE: The alarm continues sounding in all enclosures until all of the above conditions have been remedied.

➢ See also “Critical” on page 113, “Dead” on page 113, “Drive State Management” on page 31, “Hot Spare” on page 122, “Liquid Crystal Display (LCD)” on page 47, “Logical Device States” on page 124, “Offline Failed” on page 131, “Online Critical” on page 131, “Physical Device States” on page 134, and “Reset Controllers” on page 54.

Alternate Path Software

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

A software tool that manages multiple paths between the host operating system and LUNs. The software manages the multiple paths by detecting duplicate disk objects that represent a single LUN. It then designates one disk object as the primary disk object with a primary path, while the other is designated the secondary disk object with an alternate path. If the primary path becomes inaccessible, the software redirects the data to the secondary disk object through the alternate path, preserving access to the LUN. This redirection is known as path failover. The software continuously tries to access the failed path by issuing a SCSI Test Unit Ready command. A good status returned indicates the path is repaired and restored to operational status. The software automatically redirects data back to the primary path and primary disk object. This restoration of data transfer is known as path failback.

When using alternate path software, the array must be configured to use the multiport failover topology. The multiple target ID failover topology may interfere with the operation of the alternate path driver.

When using alternate path software with the FFx2 controller, the options must be set to *Loop Preferred* for the multiport topology and *Loop Only* for the multiple target ID topology.

➢ See also “Alternate Path” on page 110, “Drive Channel Failover and Failback” on page 28, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Host / Server” on page 121, “Logical Unit Number (LUN)” on page 125, and “Primary Path” on page 137.

Automatic Firmware Flash

SCOPE: This features was implemented in firmware version 8.0 and is supported on the FFx2 controller.

This feature ensures that a replacement controller operates with the same firmware version that exists on the surviving controller. Automatic firmware flash is supported only in dual-active configurations operating with firmware version 8.0 or greater. The replacement must occur with power supplied to the system (hot swap). Only the firmware is updated on the first startup of the replacement controller. An additional restart may occur if variables stored in NVRAM do not match those used by the surviving controller. A successful automatic firmware flash procedure occurs without user intervention.

When replacing a failed controller, the automatic firmware flash feature will proceed according to the following scenarios.

Power on (hot). Replacing a controller while power is supplied to the entire system is the preferred method of recovering a failed controller. The failed controller is removed and a replacement controller is inserted into the enclosure while power is supplied to the system. For this method to proceed transparently, the *Automatic Failback or Autorestore* controller parameter must be enabled. If the *Automatic Failback or Autorestore* controller parameter is disabled, the user must issue a relinquish partner command. The replacement controller may reset multiple times before completing the startup processes. The replacement controller is updating firmware and other parameters during the resetting process.

Power off (cold). Replacing a controller while power is off to the enclosure, controllers, and host system may not provide expected results. When power is returned to the system following the removal and replacement of the failed controller, both controllers begin the startup process at the same time. One controller will reach startup complete and hold the other controller in reset. If a firmware conflict exists, the active controller will flash firmware to the controller held in reset. The active controller may or may not be the surviving controller. After the matching firmware has been flashed, both controllers will reach nexus and complete the startup process. At this time, the correct firmware version can be flashed to both controllers at the same time. This scenario requires more user intervention than the previous scenario and may require flashing firmware to both controllers.

During either of the two scenarios described above, the replacement controller becomes active, and all of the replacement controller's information is synchronized with the surviving controller's information. Once this process is complete, the surviving controller pauses I/O performed on behalf of the failed controller and the replacement controller takes over. When this failback process is complete, both controllers are fully active, in nexus, and executing I/O.

➢ See also "Controller-Controller Nexus (C-C Nexus)" on page 23, "Failed Controller" on page 118, "Failover and Failback" on page 35, "Firmware Flash" on page 119, "Hot Plugging" on page 121, "NVRAM" on page 130, "Replacement Controller" on page 145, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, and "Surviving Controller" on page 148.

Automatic Restart (Reboot) on Failure

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2 and FFx2 controllers.

This feature allows a controller to restart automatically when it encounters a fatal firmware-detected error. This feature aims to reduce, but not eliminate, the need for manual intervention when a controller failure occurs. When the automatic restart feature is enabled, a restart is triggered by one of the following firmware conditions.

- Watchdog Time Out—The watchdog timer is periodically refreshed by the firmware. Failure to refresh the timer before it expires results in a watchdog time out. When the automatic restart feature is enabled, watchdog time outs are configured to restart the controller rather than hold the controller in reset.
- Partner Reset—Partner resets occur only in dual-active controller configurations. A controller resets its partner when it has reason to believe the partner has failed. When a controller resets a partner that was enabled for automatic restart, the surviving controller releases the partner from reset, waits for it to restart, and failback if possible.

NOTE: Host-commanded failover (Kill Partner) prevents the automatic restart of the disabled controller.

- Abort Function Invoked—When the firmware encounters an unrecoverable error, the abort function is invoked to perform primitive hardware shutdown procedures, print out firmware context, and force the firmware to hang in an infinite loop. In systems with a watchdog timer, hanging the firmware causes a watchdog time out, which in turn triggers an automatic restart as described above. In dual-active controller configurations without a watchdog timer, hanging the firmware causes a partner reset, which then triggers an automatic restart as described above. In single controller configurations without a watchdog timer, the firmware attempts to restart itself when the automatic restart feature is enabled.

The automatic restart feature regulates the number of restarts attempted before manual intervention is required for recovery. Regulating the number of restarts is necessary to prevent a failed controller from continually restarting. Each controller stores in NVRAM a count of the number of times it has restarted for reasons other than host-commanded reset and host-commanded failover. This count is compared to a configurable maximum restart count to determine whether the controller will enable the automatic restart feature. The feature is enabled if and only if the controller's restart count is less than the configured maximum restart count. Possible values for the maximum automatic restart count are 0 to 15 restarts.

In dual-active controller configurations, a message is sent between controllers indicating how many restarts remain before the controller reaches its configured maximum restart count. When a controller failure occurs, the surviving controller uses this information to determine whether the failover/failback process requires manual intervention. Manual intervention is required when the surviving controller detects that the failed partner controller has no restarts remaining.

Once a failed controller has restarted, the number of restart attempts is refreshed and the controller sends a message to the partner controller indicating the new number of restarts available. The refresh process provides a troubled controller with a fresh set of restarts once it has proven that it can perform for a predetermined amount of time. Possible values for the refresh interval are from 3 minutes to 7 days, or the controller can be set to infinity, which prevents automatic refresh.

An event is generated when the controller restarts and refreshes the number of restart attempts.

Restrictions and Limitations

Single controller configurations can lose cached data when restarted if no battery backup for controller cache is installed and write-back cache is enabled.

Single controller configurations not equipped with a watchdog timer cannot restart if the firmware hangs at an arbitrary place.

Single controller configurations probably cannot restart quickly enough to prevent host time outs.

Dual-active controller configurations not equipped with a watchdog timer cannot restart if both controllers hang at approximately the same time.

Dual-active controller configurations not equipped with a watchdog timer or RSTCOM do not support automatic restart.

Restarts caused by power cycles are not differentiated from restarts caused by firmware detected errors.

➢ See also "Automatic Restart (Reboot) on Failure Parameter" on page 70, "Battery Backup Unit (BBU)" on page 12, "Controller Parameters" on page 69, "Dual-Active Controller Configuration" on page 32, "Failed Controller" on page 118, "Failover and Failback" on page 35, "NVRAM" on page 130, "Single Controller Mode" on page 146, "Surviving Controller" on page 148, and "Write-Back Cache" on page 151.

Automatic Rebuild

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers.

Mylex controllers provide automatic rebuild capabilities in the event of a physical disk drive failure. The controller performs a rebuild operation automatically when a disk drive fails and the following conditions are true:

- An online spare disk drive of identical or larger size is found attached to the same controller
  OR
  The failed disk drive is removed and replaced with a disk drive of identical or larger size;
- All system drives that are dependent on the failed disk drive are configured as a redundant array; RAID 1, RAID 3, RAID 5, or RAID 0+1;
- The *Automatic Rebuild Management* controller parameter is enabled; and
- The *Operational Fault Management* controller parameter is enabled.

NOTE: The Mylex controller always attempts to first locate a replacement (online spare) disk drive that is exactly the same size as the failed disk drive. If none is found, the controller then attempts to locate a replacement disk drive that is at least the same size as the failed disk drive. The replacement disk drive size is based on the physical size, not the configured or logical size.

During the automatic rebuild process, system activity continues as normal; however, system performance may degrade slightly.

NOTE: The priority of rebuild activity can be adjusted using the Controller Parameters to adjust the *Rebuild and Check Consistency Rate*.

Online Spare Rebuild

To use the automatic rebuild feature, it is necessary to maintain an online spare disk drive in the system. The number of online spare disk drives in a system is limited only by the maximum number of disk drives available on each drive channel.

NOTE: Mylex recommends creating an online spare disk drive as part of the original configuration, or soon after creating the original configuration. If the online spare disk drive is created after a disk drive failure has occurred, the automatic rebuild does not start until the controllers have been reset.

- A disk drive may be labeled as an online spare using the *Create Standby* or *Make Hot Spare* option of the configuration utility.
- If using the VT100 or Embedded Configuration Utility, a drive pack consisting of only one drive can be created and configured as an online spare.

Refer to the *Global Array Manager Installation Guide and User Manual,* the *SANArray Manager Client Software Installation Guide and User Manual,* the *Terminal-Emulation (VT100) User Guide and LCD Front Panel Status Guide*, or the *Embedded Configuration Utility User's Guide* for more information on creating an online spare.

Hot Swap Drive Replacement

The Mylex controllers support the ability of certain drive enclosures to perform a *hot swap* disk drive replacement while the system is online. A disk drive can be disconnected, removed, and replaced with a different disk drive without taking the system offline.

NOTE: If using SCSI disk drives, the SCSI bus termination must be arranged so that a disk drive can be removed without disrupting the termination scheme.

Restrictions and Limitations

The automatic rebuild feature is dependent upon having an online spare disk drive available or hot swapping the failed disk drive with a replacement drive. If these conditions are not met, the automatic rebuild features does not operate transparently, or without user intervention.

Automatic rebuild will not start if an online spare is configured after a disk drive has failed. A rebuild command must be issued or power must be cycled before the automatic rebuild procedure starts.

A “ghost drive” is created when a disk drive fails, power is removed from the system, the disk drive is replaced or a spare drive is added to the system, and power is returned to the system. Automatic rebuild will not occur in this situation. Additionally, the system does not recognize the replacement/spare disk drive and creates a ghost drive in the same location as the failed disk drive. If the replacement/spare disk drive was inserted into the same slot as the failed drive, the ghost drive appears in the first available empty slot, beginning with channel 0, target 0. The ghost drive represents a deleted, dead drive that still exists in the configuration and the replacement/spare disk drive has a drive state of unconfigured. In order for the rebuild to occur, the replacement/spare disk drive’s state much change from unconfigured to online spare. The rebuild procedure begins after a rebuild command has been issued or the power has been cycled to the controllers. Cycling the power also removes the “ghost drive” from the configuration.

➢ See also “Automatic Rebuild Management” on page 71, “Configuration Tools and Utilities” on page 21, “Controller Parameters” on page 69, “Controller Specifications” on page 1, “Device States” on page 114, “Drive Channel” on page 27, “Drive Sizing” on page 30, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Fault Management” on page 41, “Fault Tolerance” on page 118, “Hot Swapping” on page 122, “Online Spare” on page 132, “Operational Fault Management” on page 82, “Physical Device States” on page 134, “RAID Levels” on page 137, “Rebuild and Check Consistency Rate” on page 84, “Redundant Array” on page 145, “Replacement Drives” on page 145, “Reset Controllers” on page 54, “SCSI Termination” on page 100, “System Drives” on page 149, and “Unconfigured Offline” on page 150.

Background Initialization

SCOPE: This feature was initially available only with firmware version 6.3 and the DACFL controller. This feature was modified with firmware version 7.7 and is available on the FFx, FF2 and FFx2 controllers.

Background initialization makes uninitialized system drives consistent by setting the parity while allowing the host to have read and write access to the system drive. With the size of system drives growing dramatically, a traditional initialization (foreground initialization) takes several hours to complete. The system drive is not accessible during a foreground initialization. Background initialization occurs while the host is accessing the system drive, making the storage space instantly available.

Background initialization is invoked either automatically or via a direct SCSI command. When the controller starts up or after a configuration change, background initialization automatically initializes any uninitialized system drives one at a time. When background initialization is started by a SCSI command, only the system drive specified in the command is initialized.

Background initialization can be stopped with a direct SCSI command. The background initialization is stopped only on the system drive specified in the command. Following the next restart or reconfiguration, however, background initialization starts automatically on the uninitialized system drive. Because this command also clears the NVRAM data, the initialization starts over from the beginning.

Background initialization can be disabled using the *Background Initialization* controller parameter. When background initialization is disabled, background initialization will not start automatically, nor can it be started with a direct SCSI command. This controller parameter takes affect immediately so any initialization currently operating is stopped and the NVRAM data is reset. If the controller parameter is enabled, the firmware will check for any system drives that require initialization and begin the process.

The firmware is designed to allow the background initialization to be paused during a controller reset. The background initialization progress is saved to NVRAM. After the controllers have restarted, the background initialization resumes close to where the pause occurred.

In dual-active controller configurations, background initialization continues through a controller failover. The failed controller or replacement controller is held in reset until the background initialization completes. If a relinquish controller command is executed, the background initialization completes before relinquishing the controller (failback).

Background initialization is paused by any of the following operations: Foreground Initialization, Consistency Check and Restore, Rebuild, and MORE. If one of these operations is started while background initialization is executing, background initialization is paused until the interrupting operation is complete. At that point, background initialization continues initializing the system drive it was initializing from the point of interruption. If the interrupting operation resulted in the system drive being initialized, however, background initialization will skip that system drive and search for the next uninitialized system drive and begin initializing.

Performance is degraded during background initialization due to the fact that every write requires access to all drives in the RAID group. Once all system drives have been initialized, performance should improve.

➢ See also “Automatic Rebuild” on page 10, “Enable Background Initialization” on page 78, “Consistency Check” on page 23, “Dual-Active Controller Configuration” on page 32, “Failover and Failback” on page 35, “Initialize System Drive” on page 123, “Mylex Online RAID Expansion (MORE)” on page 49, “NVRAM” on page 130, “Reset Controllers” on page 54, and “System Drives” on page 149.

Battery Backup Unit (BBU)

SCOPE: The optional BBU was implemented in firmware versions prior to 3.2 and is supported on the DACSX, FFx, and FFx2 controllers.

The DACSX supports an optional on-board standard BBU.

The FFx supports an optional on-board standard BBU and an optional external advanced BBU.

The FFx2 supports an on-board BBU with the development kit, an optional external advanced BBU, and an optional external extended life BBU.

Each BBU and corresponding battery packs are described in greater detail in the following sections.

ATTENTION: **If the controller is operating in conservative cache mode, the BBU will not turn on. The condition forcing conservative cache must be corrected before battery backup protection is available.**

The optional Battery Backup Unit (BBU) available on the DACSX, FFx and FFx2 controllers maintains memory content in the presence of an ac power failure. The principal purpose of the BBU is to provide stable memory power during ac power glitches and short power outages, however, the BBU is capable of sustaining memory content for an extended period. The length of memory power backup is dependent upon the cache memory size and the particular battery pack used by the BBU.

The BBU is designed to work with SDRAM memory, and supports both the processor control store and ASIC cache memory SDRAM. The BBU logic detects power loss on the controller and switches the SDRAM to a self-refresh mode while transparently switching the power input from +5V to battery. Transfer to BBU power only occurs if write back cache is enabled. If the controller is operating in conservative cache mode, the BBU is disabled.

The BBU provides protection to data currently stored in the on-board cache memory during intermittent power loss to the controller. This protection is important when write-back cache is enabled, and data is waiting to be flushed to the disk drive. A warm start is required in the presence of a BBU that is supplying power to the cache memory and holding data to be flushed to the disk drive as soon as the system is operational. The data will be lost if the controller is removed and a replacement controller is inserted in the removed controller's location. The BBU guarantees memory retention only when the power loss is the result of ac loss or power supply shutdown, and for as long as the batteries remain viable. Memory retention is not guaranteed if power loss is the result of removing the controller from the system while power is supplied to the system.

Battery Reconditioning

The status of the battery pack is continuously monitored by the *Gas Gauge*. The gas gauge does not read the actual charge on a battery pack, but measures and tracks the current into and out of the battery pack. The gas gauge is synchronized to the battery pack condition by performing a full discharge cycle followed by a charge cycle. A new battery pack must first be discharged and fully charged followed by another discharge and charge cycle. This process continues until the remaining capacity value is close to the full charge capacity value. A recondition cycle is performed the first time a battery pack is connected and repeated periodically. The frequency of recondition cycles is determined by the number of discharges. Any time a battery pack is disconnected and reconnected, a recondition cycle is performed. Additionally, when a controller or the module containing the battery pack is removed, a recondition cycle is performed.

DACSX BBU Characteristics

The standard BBU utilized by the DACSX includes three 800mAh AA NiCd batteries in an on-board battery pack. The BBU contains circuitry to fully recharge the batteries from a fully discharged state in seven hours. When power is first applied to the BBU, a continuous current charge of 160 mA or 20% battery capacity is applied until a voltage of 4.2V is achieved. If this voltage is not achieved in eight hours, the BBU moves into a trickle charge state. The trickle charge state charges the battery for 0.25 seconds every eight-second time interval. Once a battery pack has achieved a fully charged state, the BBU will not execute another charge cycle until power has been removed and restored.

The biggest power drain on the BBU during the back-up state is the refresh cycle. The BBU supports three refresh cycles:

- normal memory—memory is refreshed every 15 µs
- low power memory—memory is refreshed every 31 µs
- self refresh memory—memory is refreshed at the normal 15 µs rate, but stopped for a short amount of time. The self refresh memory recognizes this state and generates the refresh signals internally. When returning from a power-down state, the self-refreshing cycle requires a complete refresh pass before any memory access occurs.

The following table describes the hold up times and conditions used to achieve the hold up times.

BBU Hold-up Times and Conditions for the DACSX

Memory Size (2 SIMMs)	SIMM Type (DRAM Configurations)	Normal Refresh Data Retention (hours)		Low Power Data Retention (hours)		Self Refresh Data Retention (hours)	
		Minimum	Typical	Minimum	Typical	Minimum	Typical
8 MB	9 — 1M x 4	15	42	48	88	70	161
8 MB	2 — 1M x 16 and 1 — 1 M x 4	34	93	71	126	82	208
16 MB	18 — 1 M x 4	8	22	30	49	48	57
16 MB	4 — 1 M x 16 and 2 — 1 M x 4	19	51	49	75	61	129
32 MB	9 — 4 M x 4	14	40	34	45	40	81
32 MB	8 — 4 M x 4 and 4 — 4 M x 1	11	31	31	42	38	78
64 MB	18 — 4 M x 4	7	20	20	24	24	45
128 MB	36 — 16 M x 1	4	12	9	9	13	16
128 MB	8 — 16 M x 4 and 4 — 16 M x 1	11	30	data not avail.	data not avail.	data not avail.	data not avail.
256 MB	18 — 16 M x 4	5	19	data not avail.	data not avail.	data not avail.	data not avail.

FFx BBU Characteristics

The FFx supports two BBUs: the standard BBU and the advanced BBU.

Standard BBU

The standard BBU utilizes three 650 mAh NiCd batteries assembled into an on-board battery pack with an over current protective device and a connector. The following table describes the hold up times and conditions used to achieve the hold up times.

Standard BBU Hold-up Times and Conditions for the FFx

Memory Size (MB)	Chip Technology (Mb)	Worst Case Hold-up Time (hours)	Typical Hold-up Time (hours)
64	64	28	62
128	64	18	40
256	256	18	40
512	256	12	24

Advanced BBU

The advanced BBU provides the same functionality as the standard BBU and may be used in place of the standard BBU. Its form factor and signaling are the same as the standard BBU, however, it offers time-based charging of an externally located battery pack. The Advanced BBU utilizes three

3800 mAh NiMH batteries assembled into a battery pack. The following table describes the hold up times and conditions used to achieve the hold up times.

Advanced BBU Hold-up Times and Conditions for the FFx

Memory Size (MB)	Chip Technology (Mb)	Worst Case Hold-up Time (hours)	Typical Hold-up Time (hours)
64	64	178	450
128	64	114	220
256	256	114	220
512	256	68	109

FFx2 BBU Characteristics

The FFx2 supports three BBUs: the Development Kit on-board BBU, the advanced BBU, and the extended life BBU. The three BBU circuits are different and are designed to properly charge one specific battery pack. Multiple battery packs are not supported.

Development Kit BBU

The Mylex development kit BBU utilizes three 880 mAh NiMH batteries assembled into an on-board battery pack with an over current protective device and a connector. This BBU and battery pack is available only with the FFx2 development kit. The following table describes the hold up times and conditions used to achieve the hold up times.

Development Kit BBU Hold-up Times and Conditions for the FFx2

Memory Size (MB)	Chip Technology (Mb)	Worst Case Hold-up Time (hours)	Typical Hold-up Time (hours)
128	64	50	68
256	256	31	34
512	256	17	21

Advanced BBU

The FFx2 advanced BBU utilizes three 3800 mAh NiMH batteries assembled into an external battery pack. The following table describes the hold up times and conditions used to achieve the hold up times.

Advanced BBU Hold-up Times and Conditions for the FFx2

Memory Size (MB)	Chip Technology (Mb)	Worst Case Hold-up Time (hours)	Typical Hold-up Time (hours)
128	64	217	295
256	256	132	145
512	256	74	90

Extended Life BBU

The extended life BBU utilizes three 3700 mAh NiMH batteries assembled into an external battery pack. The following table describes the hold up times and conditions used to achieve the hold up times.

Extended Life BBU Hold-up Times and Conditions for the FFx2

Memory Size (MB)	Chip Technology (Mb)	Worst Case Hold-up Time (hours)	Typical Hold-up Time (hours)
128	64	211	287
256	256	129	141
512	256	72	88

Modes of Operation

Each FFx and FFx2 BBU has three modes of operation: Fast Charge, Discharge, and Trickle Charge. The fast charge mode determines the rate at which a fully discharged battery pack becomes fully charged. The amount of time required for a battery pack to become fully charged is dependent upon battery age and operating temperature. The discharge mode determines the rate at which a fully charged battery pack is completely discharged. A recondition cycle requires a full discharge cycle and a full fast charge cycle. If the gas gauge cannot determine with certainty the capacity of the battery pack, an additional recondition cycle is required. The trickle charge mode determines the rate at which a battery pack is maintained at the fully charged state without over charging the battery pack. The following table lists the rates for each mode of operation for each type of BBU.

Modes of Operation for the FFx and FFx2

Type of BBU	Fast Charge	Time Required to be Fully Charged[b]	Discharge	Time Required to be Fully Discharged[b]	Length of Recondition Cycle[b]	Trickle Charge
FFx Standard BBU 650 mAh	167 mA or 26% battery capacity	5 to 6 hours	90 mA or 14% battery capacity	7 hours	13 hours	26 mA or 4% battery capacity
FFx Advanced BBU 3800 mAh	750 mA or 20% battery capacity	6 hours	300 mA or 8% battery capacity	13 hours	19 hours	35 mA or 0.9% battery capacity
FFx2 DevKit BBU 880 mAh	183 mA or 20% battery capacity	6 hours	300 mA or 34% battery capacity	3 hours	9 hours	13 mA or 1.5% battery capacity
FFx2 Advanced BBU 3800 mAh	685 mA or 18% battery capacity	6 hours	300 mA or 8% battery capacity	13 hours	19 hours	13 mA or 0.35% battery capacity
FFx2 Extended Life BBU 3700 mAh	685 mA or 18% battery capacity	6 hours	300 mA or 8% battery capacity	13 hours	19 hours	13 mA or 0.35% battery capacity

b. The times listed for each cycle are approximate. The amount of time required for each cycle is dependent upon battery age and operating temperature.

Battery Life

The NiCd battery pack has a useful life of up to five years depending on frequency of discharge/charge cycles, depth of discharge, discharge and charge rates, and operating and storage temperatures. The Mylex controller applications of the NiCd battery pack is not strenuous and should result in a battery life of close to five years. An optimal operating temperature of 25 ºC is recommended to obtain this length of battery life. If the operating temperature is allowed to be 40 ºC, and the battery is frequently discharged and charged, the life could be shortened to two years.

The life of the NiMH battery is similar to that of the NiCd. Under normal conditions and at an operating temperature of 25 ºC, the NiMH battery should have a battery life of up to five years. If the number of discharge/charge cycles exceeds 100 per year and if the temperature is greater than 25 ºC, the life will be shortened.

Determining the end of life of a battery is arbitrary and depends on the application and the user's tolerance for lower battery capacity. As a general rule, a battery is considered to have reached end of life when the capacity is reduced to 60 percent of nominal. This is probably too low for Mylex controller application. Mylex recommends a battery end of life when the capacity has reached 80 to 90 percent of nominal. The battery end of life depends on user requirements for length of backup and the size of the memory. A battery pack that provides 129 hours of backup protection when new would provide 72 hours of backup protection when degraded to 60 percent of nominal. A battery pack that provides only 72 hours of backup protection when new, however, would provide only 65 hours of backup protection when degraded to just 90 percent. Backup protection of 65 hours may be considered marginal protection.

➢ See also "Cache" on page 111, "Cache Flush" on page 112, "Conservative Cache Mode" on page 22, "Controller Specifications" on page 1, "DACSX" on page 25, "FFx" on page 42, "FFx2" on page 43, "Hot Plugging" on page 121, "Replacement Controller" on page 145, "Standard Data Caching" on page 63, "Warm Start" on page 151, "Write-Back Cache" on page 151, and "Write-Through Cache" on page 151.

Cache Coherency

SCOPE: This feature was implemented in firmware version 5.4 and is supported the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

The Cache Coherency feature implements an internal reservation mechanism that forces a controller to reserve some or all of a system drive in order to allow host access. A READ or WRITE command from a host computer implies that the controller handling the command needs to reserve the section of the system drive used by the command for the duration of the command. While the controller has the reservation, the partner controller cannot access the area reserved. Prior to allowing the reservation, the partner flushes all data in the reserved area and invalidates its read cache for the area.

The smallest area that is reserved is an entire stripe (stripe size multiplied by the number of disks in the group). The largest area that can be reserved is an entire system drive. The reservation mechanism is transparent to the host computer(s). If the entire system drive is reserved, then the controller uses a fast path mechanism that does not impact performance.

NOTE: Even with Cache Coherency, a host computer cannot safely issue a write command to both controllers for the same logical block address on the same controller. In order to guarantee the order that the data is written, a host computer must wait for command completion for a write before issuing another write for the same address on the same system drive to the other controller; or issue a reserve/release command sequence reserving the controller.

➢ See also "Cache Flush" on page 112, "Controller-Controller Nexus (C-C Nexus)" on page 23, "Dual-Active Controller Configuration" on page 32, "Host / Server" on page 121, "Implied Reserve and Release" on page 46, "Partner Controller" on page 133, "Standard Data Caching" on page 63, "Stripe Size" on page 64, "System Drives" on page 149, "Write-Back Cache" on page 151, and "Write-Through Cache" on page 151.

Cache Line Size

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers.

The Cache Line Size function is set in conjunction with stripe size and represents the size of the data “chunk” that will be read or written at one time. The cache line size should be based on the stripe size. Supported cache line sizes for Mylex external RAID controllers are 8 KB, 16 KB, 32 KB, and 64 KB. Currently, SAM/GAM automatically sets the cache line size when the stripe size is selected.

➢ See also “Controller Parameters” on page 69, and “Stripe Size” on page 64.

Cache Mirroring

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers.

The task of copying accepted host write data from the primary controller to the partner controller. This method ensures that host write data is safely mirrored to the partner controller before successful completion status is returned to the host. In the event of a controller failure, the surviving controller safely retains all mirrored data with no loss of user data.

➢ See also “Cache” on page 111, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Host / Server” on page 121, “Partner Controller” on page 133, “Primary Controller” on page 136, and “Surviving Controller” on page 148.

Configuration Coherency

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers.

Coherency maintains any configuration changes or disk drive state changes between dual-active controllers. When a controller receives and successfully executes a command resulting in a change to the current configuration, the changes are also sent to the partner controller. The partner controller makes the appropriate changes so that both controllers possess matching configuration information.

➢ See also “Configuration on Disk (COD)” on page 18, “Controller-Controller Nexus (C-C Nexus)” on page 23, “Drive State Management” on page 31, “Dual-Active Controller Configuration” on page 32, “Partner Controller” on page 133, “Physical Device States” on page 134, and “Primary Controller” on page 136.

Configuration on Disk (COD)

SCOPE: Configuration on Disk (COD) version 1.0 was implemented in firmware prior to version 3.2 and is supported on all controllers. COD 2.1 was implemented in firmware version 7.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

COD 1.0

Configuration on Disk (COD) 1.0 retains the latest version of the saved configuration at a reserved location on each disk drive configured in the system. The same information is written to a reserved location of the NVRAM of the controller(s). The configuration information includes the configuration ID number and the configuration sequence number. Both numbers reset to 0 when the configuration is changed. The configuration sequence number increments whenever there is a physical disk drive state change. This feature allows the configuration to be maintained through a controller replacement.

COD 1.0 plays a significant role during the power up sequence following a controller replacement. The replacement controller progresses through a series of checks and validation steps to determine the validity of any configuration currently present in NVRAM followed by additional checks and validation steps to determine the validity of the COD 1.0 information present on all disk drives present. The final configuration is determined by one of the following criteria:

- The majority of COD 1.0 information with the same configuration ID,
- The COD 1.0 available on the disk drive (when there is one disk drive and one controller in the system), or
- The first configuration encountered (when there is an equal number of valid configurations).

COD 2.1

COD 2.1 retains the latest version of the saved configuration at a reserved location on every physical device, not on the controller NVRAM. The COD 2.1 data contains the following information:

- Device definition. This section contains three pieces of information:
    - The logical device definition/structure for those logical devices dependent on this physical device. This information should be the same for all physical devices associated with the defined logical device.
    - Any physical device information pertaining to this physical device that is different for different physical devices even though they may be part of the same logical device definition.
    - A unique COD 2.1 signature that is used during startup to resolve the configuration. This information should be the same on all the physical devices that are part of the defined logical device. This signature may be the same on other physical devices as well.
- Data backup for data migration, MORE. This section contains required information for the MORE feature. This area also includes required information for the Background Initialization feature.
- User device name information and host software configuration parameters. This information is defined by the user and should be the same on all physical devices that are associated with the defined logical device.
- COD 2.1 locking mechanism. This area is designed to provide a locking mechanism to be used with multiple controller systems in which any one of the controllers is allowed to update the COD 2.1 area independently of the other controllers. This feature allows the controller to lock the drive’s COD 2.1 area for write access before updating the COD 2.1 to that drive.

Firmware determines COD 2.1 validity using the following criteria:

- First, use the most recent COD 2.1 independent of majority or device. The most recent COD 2.1 is updated to all configured devices. Unconfigured devices receive no COD 2.1 update, the information is all set to zeros.
- Second, use the COD 2.1 written by this controller.

ATTENTION: Downgrading firmware from 7.0 is not recommended. You need to backup all system data, clear the configuration, download the previous firmware version, recreate the configuration, and then replace the data.

Mixing controllers and/or disk drives from systems running different versions of firmware present special situations that may affect data integrity. Firmware versions prior to 7.0 write the configuration to each device using COD 1.0 format. Firmware versions 7.0 and greater write the configuration to each device using COD 2.1 format. COD 2.1 format requires a larger amount of unused space on each device than COD 1.0 format. Beginning with firmware version 5.4, all Fibre Channel disk drives reserve the larger unused space for COD information.

ATTENTION: Regardless of controller firmware version, any existing COD on a disk drive that is inserted after the controller has started (STARTUP COMPLETE) is overwritten.

Firmware 5.4 added a safeguard to ignore drives containing COD 2.1 information and mark the drives as unusable. For example, a controller using firmware version 5.4 or 6.1 is starting and encounters a disk drive with COD 2.1. The controller marks the drive as unusable to avoid overwriting COD 2.1 data with COD 1.0 data. This safeguard is intended to preserve user data in a scenario where a system is powered down and a controller using firmware version 7.0 is mistakenly replaced with a controller using firmware version 5.4 or 6.1. This safeguard does not exist in the 5.1 or earlier firmware releases. This safeguard does not affect disk drives that are installed after the controller has started. Regardless of controller firmware version, any existing COD on a disk drive that is inserted after the controller has started is overwritten.

When a controller using firmware version 7.0 is starting and finds all the disk drives have COD 1.0 data, an upgrade to COD 2.1 is performed. If a mixture of disk drives with COD 1.0 and COD 2.1 are found, the COD 1.0 is ignored and the configuration is obtained from the COD 2.1 disk drives.

If a configuration was created with firmware versions prior to 7.0 and a configuration utility other than GAM version 3.1 or newer, it is possible that the reserved space on each disk drive is not large enough to support COD 2.1. In this instance, the controller(s) aborts so that the COD information and user data is preserved. Recovery is possible by one of the following methods:

- Disconnecting the drives from the controllers and downgrading the controller firmware to the previous firmware version, or
- Backing up all data, clearing the configuration, upgrading the firmware, recreating the same configuration, and restoring data to the disk drives.

Refer to “Upgrading a Configuration Created with RAIDfx” on page 101 for a detailed description of this procedure.

➢ See also “Configuration Tools and Utilities” on page 21, “Controller Specifications” on page 1, “Drive Roaming” on page 29, “Drive State Management” on page 31, “Dual-Active Controller Configuration” on page 32, “Firmware Versions” on page 95, “Replacement Controller” on page 145, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, “Replacing a Failed Controller in Simplex Systems” on page 99, and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

Configuration Tools and Utilities

SCOPE: The configuration tools provided by Mylex have been modified due to additional feature support. Some of the configuration tools or utilities are no longer supported and others have been modified significantly.

Mylex controllers can also be accessed through customer supplied configuration tools or utilities.

RAIDfx is supported with firmware versions 3.2 through 6.3 and not supported with firmware 7.0 or greater.

The LCD front panel and keypad has limited support with firmware versions prior to 7.0. In addition, the LCD is not supported on the FFx or FFx2 controllers.

The VT100 interface has limited support with firmware versions prior to 7.0, but has been modified to support the features available with firmware versions 7.0 and greater.

Firmware versions 7.0 and 7.2 require GAM version 3.1 or greater. Firmware versions 7.4 or greater require SAM/GAM version 4.0 or greater.

Firmware version 7.7 implemented an embedded configuration utility which is accessible through a terminal-emulation interface, such as VT100.

Mylex controllers use several configuration tools or utilities. Mylex provides RAIDfx, SAM/GAM and an on-board utility accessible through the LCD front panel and keypad. Mylex controllers are also accessible through a terminal-emulation interface, such as VT100. Mylex controllers can also be accessed through customer supplied configuration tools or utilities. The configuration tools or utilities allow the user to create unique RAID configurations specific to individual needs. The tools also provide the user with the necessary procedures for performing additional management operations such as automatic rebuild, relinquish partner controller, MORE, and controller parameter modifications. Some, but not all, utilities also provide the user with monitoring abilities and error message displays.

Firmware versions prior to 7.0 supported the RAIDfx configuration tool. Beginning with firmware version 7.0, RAIDfx is no longer supported. Using firmware version 7.0 to access configurations created with RAIDfx causes the controllers to abort. Refer to “Upgrading a Configuration Created with RAIDfx” on page 101 for a detailed procedure for upgrading RAIDfx configurations.

Firmware versions prior to 7.0 also supported an on-board utility accessible through the LCD front panel and keypad. Beginning with firmware version 7.0, LCD front panel functionality is limited to a status-only display. The keypad support provided in previous firmware releases is no longer functional. The status messages are similar to those displayed on the VT100 screen.

Firmware version 7.7 implemented an embedded configuration utility which is accessible using a serial connection to a terminal-emulation interface such as VT100. The embedded configuration utility is also accessible in-band over the Fibre Channel interface.

➢ See also “Embedded Configuration Utility” on page 35, “Global Array Manager (GAM)” on page 45, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, and “VT100” on page 65.

Conservative Cache Mode

SCOPE: This feature was implemented in firmware version 3.3 and is supported on all controllers. The conservative cache mode features were modified with firmware versions 4.2, 6.1, and 7.4. Modifications implemented in firmware version 6.1 are supported on the DACFL, DACFF (FF), FFx, FF2 and FFx2 controllers. Modifications implemented in firmware version 7.4 are supported on the FFx, FF2, and FFx2 controllers.

Conservative cache mode is an operating mode in which logical devices, or system drives, configured with the write-back write caching policy are treated as though they were configured for write-through operation and the cache is flushed.

Conservative cache mode is entered to provide a higher level of data protection after a critical system component has failed. When the condition causing conservative cache is resolved, the system drives are converted back to their original settings.

Conditions that cause conservative cache execution are:

- The *Conservative Cache Mode Parameter* is enabled for a dual-active controller configuration, and a controller failure has occurred.
- A MORE (Mylex Online RAID Expansion) operation is initiated. During a MORE operation, the system automatically enters into conservative cache mode. After MORE completes, the system drives are restored to their original settings.
- A power supply has failed (not power supply removal or power supply not present).
- When an enclosure supporting redundant power (multiple power supply support) is operating on a single power supply (implemented with firmware version 4.2).
- An out-of-limit temperature condition exists (implemented with firmware version 4.2).
- The controller receives an indication of either an ac failure or a low battery from a UPS (implemented with firmware version 4.2).
- The external hardware input signal has been asserted to force the controller into conservative cache mode (implemented with firmware version 4.2).

NOTE: For dual-active controller systems, the hardware input signal must be asserted to each controller. The cache dirty output signal can be monitored to determine when the cache has been flushed. Refer to the OEM Hardware Reference Manual for your product for pinout and signal information.

- SES monitoring has failed (implemented with firmware version 6.1).
- The remaining capacity of the intelligent BBU, if supported, falls below a user-defined threshold (implemented with firmware version 7.4).

ATTENTION: **If the controller is operating in conservative cache mode, the BBU will not turn on. The condition forcing conservative cache must be corrected before battery backup protection is available.**

➢ See also “AEMI” on page 6, “Battery Backup Unit (BBU)” on page 12, “Cache Flush” on page 112, “Conservative Cache Mode Parameter” on page 72, “Controller Parameters” on page 69, “Environmental Device” on page 117, “Failed Controller” on page 118, “Fault Management” on page 41, “Mylex Online RAID Expansion (MORE)” on page 49, “SAF-TE” on page 55, “SES” on page 58, “Standard Data Caching” on page 63, “System Drives” on page 149, “Warm Start” on page 151, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Consistency Check

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers. Consistency Check is also referred to as a parity check.

The Consistency Check function is used to verify the integrity of data on a system drive. It verifies that mirror or parity information matches the stored data on redundant arrays (RAID 1, RAID 3, RAID 5, or RAID 0+1). If the parity block information is inconsistent with the data blocks, the controller has the ability to correct the inconsistencies.

A consistency check continues through a controller failure (failover). The failed controller or replacement controller is held in reset until the consistency check completes. If a relinquish controller command is executed, the consistency check completes before relinquishing the controller (failback).

➢ See also “Background Initialization” on page 11, “Failover and Failback” on page 35, “Initialize System Drive” on page 123, “Mirroring” on page 129, “Mylex Online RAID Expansion (MORE)” on page 49, “Parity” on page 133, “Parity Check” on page 51, “RAID Levels” on page 137, “Redundant Array” on page 145, and “Striping” on page 147.

Controller-Controller Nexus (C-C Nexus)

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers.

Controller-to-Controller Nexus (C-C Nexus) is a connection between the two controllers of a dual-active controller pair that enables each controller to copy write-back data to its partner controller and determine if the partner controller is operating.

➢ See also “Automatic Firmware Flash” on page 7, “Cache Coherency” on page 17, “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, “Configuration Coherency” on page 18, “Dual-Active Controller Configuration” on page 32, “Partner Controller” on page 133, “Secondary Controller” on page 146, “Standard Data Caching” on page 63, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Controller Event Logging

SCOPE: This feature was implemented in firmware version 7.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controller.

Mylex controllers contain a circular buffer of 512 events that are transferred to the configurator.

➢ See also “Configuration Tools and Utilities” on page 21, “Embedded Configuration Utility” on page 35, “Global Array Manager (GAM)” on page 45, “SANArray Manager™ (SAM)” on page 56, and “VT100” on page 65.

Controller Fault Indication

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This features adds the capability of the controller to store why it has faulted in persistent memory. This value must be retrievable by software applications after the controller is restored. A more complete implementation of this feature adds the capability of informing a partner of the fault value during the shutdown of a faulted controller. The method of extracting the fault condition value remains the same, only now the survivor provides the information. This capability is provided if the debug dump data is written to disk, but not otherwise.

Currently the only fault supported is Memory Fault Indication which indicates that a multi-bit ECC error occurred in the data cache memory. This value can be used to determine if the memory DIMM for data cache has experienced an error. The user can use this value as an indication to replace the memory DIMM.

➢ See also “Debug Dump” on page 25, “Dual-Active Controller Configuration” on page 32, “Embedded Configuration Utility” on page 35, “Failed Controller” on page 118, and “Surviving Controller” on page 148.

DACFF (FF)

A Mylex disk array controller with Fibre Channel host interface and Fibre Channel disk drive interface. The DACFF (FF) supports two host channels and four drive channels. The DACFF (FF) utilizes firmware versions 5.0 through 7.2.

ATTENTION: **Do not install an FF2 controller in a dual-active configuration with a DACFF (FF) controller. The current configuration and data may be lost.**

The FF2 controller appears similar to the DACFF (FF) controller, however, the two controllers are not interchangeable. A dual-active controller configuration must consist of two identical controllers. A DACFF (FF) system can be upgraded to an FF2 system following the procedure described under “Upgrading a DACFF (FF) System to an FF2 System” on page 102.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “FF2” on page 42, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, and “Host Channel” on page 121.

DACFL

A Mylex disk array controller with Fibre Channel host interface and low-voltage differential SCSI disk drive interface. The DACFL supports two host channels and four drive channels. The DACFL utilizes firmware versions 4.2 through 6.3.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, “Host Channel” on page 121, “SCSI” on page 145, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

DACSF

A Mylex disk array controller with Fibre Channel host interface and single-ended SCSI disk drive interface. The DACSF supports two host channels and six drive channels. The DACSF utilizes firmware versions 4.0 through 4.2.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, “Host Channel” on page 121, “SCSI” on page 145, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

DACSS

A Mylex disk array controller with single-ended or differential SCSI host interface and single-ended SCSI disk drive interface. The DACSS supports two host channels and four drive channels. The DACSS utilizes firmware version 4.4 only.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Firmware Versions” on page 95, “Host Channel” on page 121, “SCSI” on page 145, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

DACSX

A Mylex disk array controller with single-ended or differential SCSI host interface and single-ended SCSI disk drive interface. The DACSX supports one host channel and two drive channels without a daughter card attached. The DACSX supports one or two host channels and four or five drive channels with a daughter card attached. The DACSX utilizes firmware versions 3.2 through 4.2.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Firmware Versions” on page 95, “Host Channel” on page 121, “SCSI” on page 145, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

Debug Dump

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This features records controller state information when an abort occurs. After the abort has completed the abort information can be retrieved and analyzed to help determine why the abort occurred. The information is generated while a controller abort is in progress and recorded to NVRAM and a disk drive, if enabled. The aborting controller blocks requests from the partner controller so that the dump can be generated. The Debug Dump is written to a Reserved Disk Area (RDA) on one selected disk drive. The contents of the dump are not user configurable.

This feature has two levels of support.

1. The first level records the reason why the abort occurred and is always enabled.
2. The second level of support records more information about the abort and is configurable. The second level of support is enabled or disabled using the VT100 user interface. The default for this feature is on.

When the dump is complete, status information is written to the Debug Dump header. A “scan for dumps” starts on the surviving controller and an event is generated if a dump is found. The host uses the Retrieve Dump direct command to collect the dump data from the controller.

The NVRAM Debug Dump data structure holds ten data entries with the most recently generated entry being returned to the host when requested. The Debug Dump data entry consists of an abort code, an error code, and an event flag. The abort code specifies where in the firmware image the abort occurred. The abort code is written to NVRAM when the controller is aborting and generates an event. The abort code takes up 2 bytes of data.

Retrieve Debug Dump

The Retrieve Debug Dump command returns status information and data previously saved to disk by an aborting controller. The status information returned by the command is the abort code, possibly an error code, and an event flag. This information is generated regardless of the Debug Dump feature and is overwritten by each successive abort. The data saved to disk by the aborting controller includes the data described in the abort header and the data previously written to the trace buffer. This data is generated only when the Debug Dump feature is enabled.

Debug Dump data can be retrieved two different ways.

1. Scan for the first unread dump.
2. Specify the disk from which to read the dump. This method can be used to retrieve secondary dump data after a primary dump has been read.

➢ See also “Controller Fault Indication” on page 23, “Failed Controller” on page 118, “NVRAM” on page 130, “Partner Controller” on page 133, “Surviving Controller” on page 148, and “VT100” on page 65.

Default Mode Page Settings

SCOPE: This feature was implemented in firmware version 7.7 and is supported on the FFx, FF2, and FFx2 controllers.

This feature defines a set of drive mode page settings and the preferred values required for all disk drives in the system. This provides more control over the disk drives by providing behavior consistency and predictability, and eliminating data integrity issues associated with some mode page settings. When the primary controller is started and the disk drives have completed the spin-up process, the current mode page settings for each disk drive are compared to the preferred values. If the disk drive settings are inconsistent, a Mode Select command is issued to set the preferred values for the disk drive. When a new disk drive is detected during a device scan, the mode page settings are checked and updated if necessary.

The following preferred mode page settings are defined by this feature.

- Automatic Write Reallocation Enable (AWRE) and Automatic Read Reallocation Enable (ARRE) are enabled from the Read-Write Error Recovery mode page (0x01).
- Write Cache Enable (WCE) is cleared from the Cache mode page (0x08).
- Require Hard Address (RHA)/Disable Soft Addressing (DSA) and Prevent Loop Port Bypass (PLPB) bits are set from the Fibre Channel Port Control mode page (0x19) and all other bits on this page are cleared.

➢ See also “Fibre Channel” on page 43, “Hard Loop ID” on page 120, “Mode Page Parameters” on page 96, “Physical Disk Drive” on page 135, and “Port” on page 135.

Disk Media Error Management

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers.

Mylex controllers transparently manage disk media errors. Disks are programmed to report errors, even ECC-recoverable errors.

When a disk reports a media error during a read, the controller reads the data from the mirror (RAID 1 or RAID 0+1) or computes the data from the parity blocks (RAID 3 or RAID 5) and writes the data back to the disk that encountered the error. If the write fails (media error on write), the controller issues a “reassign” command to the disk, and then writes the data to a new location. Since the problem has been resolved, no error is reported to the system.

➢ See also “Controller Parameters” on page 69, “Reassign Restricted to One Block” on page 83, “Mirroring” on page 129, “Mode Page Parameters” on page 96, “RAID Levels” on page 137, and “Parity” on page 133.

Drive Channel

SCOPE: The number of drive channels available on a controller varies with each controller. The number of drives supported on each drive channel also varies with the controller type and firmware version.

The drive channel serves as a path for the transfer of data and control information between the physical disk drives and the controller. Mylex disk array controllers have up to six drive channels. Depending on the firmware version, each drive channel can support from 14 up to 124 physical disk drives. The following table lists each controller and the number of disk drives supported with various levels of firmware.

Maximum Number of Disk Drives per Drive Channel

Controller	Firmware Version	# Drive Channels	# Disks per Channel	Max # Disks Supported
DACSX without Daughter Card	3.2 through 4.2	2	15 (simplex) 14 (duplex)	30 (simplex) 28 (duplex)
DACSX with Daughter Card	3.2 through 4.2	5/4	15 (simplex) 14 (duplex)	75/60 (simplex) 70/56 (duplex)
DACSS	4.4	4	15 (simplex) 14 (duplex)	60 (simplex) 56 (duplex)
DACSF	4.0 through 4.2	6	15 (simplex) 14 (duplex	90 (simplex) 84 (duplex)
DACFL	4.2	4	15 (simplex) 14 (duplex)	60 (simplex) 56 (duplex)
	5.0 through 6.3	4	16	64
DACFF (FF)	5.0 through 6.2	4	16	64
	7.0 and greater	4	124	496
FF2	6.0 through 6.2	4	16	64
	7.0 and greater	4	124	496
FFx	6.0 through 6.2	2	16	32
	7.0 and greater	2	124	248
FFx2	8.0	2	124	248

➢ See also “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “DACFF (FF)” on page 24, “DACFL” on page 24, “DACSF” on page 24, “DACSS” on page 25, “DACSX” on page 25, “Dual-Active Controller Configuration” on page 32, “Duplex” on

page 116, “Drive Distribution” on page 116, “FF2” on page 42, “FFx” on page 42, “FFx2” on page 43, “Soft Addressing Detection” on page 61, “Host Bus Adapter (HBA)” on page 120, “Physical Disk Drive” on page 135, and “Simplex” on page 60.

Drive Channel Failover and Failback

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers and requires Fibre Channel, dual-ported disk drives.

This feature was modified with firmware version 7.7 and is supported on the FFx, FF2, and FFx2 controllers and requires Fibre Channel, dual-ported disk drives.

All Fibre Channel disk drives are manufactured with two ports to provide redundant access to the disk drive should one port fail or become inaccessible. The two disk ports should not be connected to the same disk loop. In most configurations, both disk ports are connected to separate disk loops, although some configurations may only connect one disk port and leave the other disk port disconnected.

At startup time, the controller obtains a list of all ports (disks, controllers, and other devices) that are connected to each disk loop. Information about each port is obtained to determine if any two ports reference the same Fibre Channel node. For example, a Fibre Channel disk drive represents one Fibre Channel node with two ports. Each port of a disk drive has a unique port name but the same node name. When the controller determines that a dual-ported device is present, an algorithm is used to select the preferred port and alternate port for the device. The preferred port is used for normal access whereas the alternate port is used only when access to the preferred port is not possible.

This feature does not allow the use of disk drive load balancing. At any point in time, the controller accesses Fibre Channel disk drives through a single port only. Some Fibre Channel drives that support dual ports do not currently allow commands to both ports simultaneously.

Enhancements implemented in firmware version 7.7 do not automatically perform failover to the alternate loop or port. In the modified implementation, if a controller determines that a disk loop is down or a disk port is no longer accessible, an algorithm is executed before failing over to the alternate loop or port. The new algorithm maintains an agreement between the controllers as to which disk loops and disk ports are active. In the previous implementation, if a disk drive was being accessed through its alternate port and the preferred port became accessible, the controller would failback to the preferred port automatically. In the modified implementation the failback does not occur until the new algorithm has determined that failing back is warranted.

For example, disk drives that formerly appeared on two loops but now appear on only one loop may fail over to their alternate port. Disk drives that formerly appeared on only one loop but now appear on two loops may failback to their preferred port. Disk drives that formerly did not appear on any loop are assigned a preferred port and possibly an alternate port.

Some system configurations require that dual-active controllers reside in separate enclosures to provide for the expansion of additional enclosures. This requirement yields a system configuration in which pulling any disk loop cable results in each controller seeing different behavior of the loop or a different subset of disk ports. For example, removal or failure of a disk cable downstream from both controller nodes results in those disk ports being inaccessible to both controllers. If the failure occurs between the controller nodes, each controller will see a different subset of disk ports on the loop and cannot see the other controller. This situation requires that both controllers stop using any disk loop where the other controller cannot be seen.

Drive channel failover/failback requires that disk drives be assigned hard IDs. Soft addressing is not permitted.

➢ See also “Alternate Channel” on page 110, “Channel” on page 112, “Controller Parameters” on page 69, “Disk Loop” on page 115, “Disk Port” on page 115, “Dual-Ported Drive Support” on page 34, “Fibre Channel” on page 43, “Soft Addressing Detection” on page 61, “Hard Loop ID” on page 120, “Port” on page 135, “Preferred Channel” on page 136, and “Single Point of Failure (SPOF)” on page 61.

Drive Packs / Groups

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers. The feature was modified with firmware version 7.0 to allow 16 disk drives per pack.

A drive pack is a group of disk drives created during the array configuration process. Drive packs or groups have the following properties:

- From one to 8 or 16 (firmware version 7.0 and greater) disk drives are included in an individual drive pack or group. This enables configuration of drives into drive packs of 7 or 15 data drives plus 1 parity drive, or 7 + 1 or 15 + 1. Configurations of 2 + 1 through 15 + 1 are supported.
- The drive pack can include physical disk drives located on different drive channels.
- The number of disk drives in a drive pack determines the possible RAID level.

The capacity of a drive pack is determined by multiplying the size of the smallest disk drive in the drive pack times the number of disk drives in the drive pack. A drive pack can be used in many ways. For example it can be divided into 8 or 32 (firmware version 7.0 and greater) system drives or logical devices; it can be a single system drive; or it can be the last part of one spanned system drive and the first part of the next spanned system drive.

➢ See also “Array” on page 111, “CAP Strategy for Selecting a RAID Level” on page 92, “Configuration Strategies” on page 94, Embedded Configuration Utility, “Firmware Versions” on page 95, “Global Array Manager (GAM)” on page 45, “Mylex Online RAID Expansion (MORE)” on page 49, “Parity” on page 133, “RAID Levels” on page 137, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, “Spanning Drive Packs” on page 62, “System Drives” on page 149, and “VT100” on page 65.

Drive Roaming

SCOPE: This feature was implemented in firmware version 4.2 and is supported by all controllers.

Drive roaming allows disk drives to be moved to other channel/target ID locations while the system is powered down. Drive roaming allows for easier disassembly and assembly of systems, and potential performance enhancement by optimizing channel usage.

Drive roaming uses the Configuration on Disk (COD) information stored on the physical disk drive. When the system restarts, a table is generated for each disk drive’s current location and where each disk drive was last located before the system was powered down. This information is used to remap the physical disk drives into the proper location in the system drive.

This feature is designed for use within one system environment, for example, a single system or a cluster of systems sharing a simplex or dual-active controller configuration. Foreign disk drives containing valid COD information from other systems must not be introduced into a system. If the

COD information on a replacement disk drive is questionable or invalid, the disk drive will be labeled unconfigured offline or dead.

➢ See also “Channel” on page 112, “Configuration on Disk (COD)” on page 18, “Dead” on page 113, “Drive State Management” on page 31, “Dual-Active Controller Configuration” on page 32, “Physical Device States” on page 134, “Physical Disk Drive” on page 135, “Replacement Drives” on page 145, “Simplex” on page 60, “System Drives” on page 149, “Target ID (TID)” on page 149, and “Unconfigured Offline” on page 150.

Drive Sizing

SCOPE: This feature was implemented in firmware version 3.3 and is supported on all controllers. Drive Size Truncation was implemented in firmware version 7.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

Drive sizing takes various similarly sized physical disk drives and assigns a virtual capacity to each so that they all appear to be the same size from the perspective of the controller. As of firmware version 7.0, drive sizing is accomplished via a table driven drive size truncation algorithm. Prior to firmware version 7.0, a drive size file was used to control drive sizes.

The controller firmware automatically truncates a physical disk drive’s capacity by a factor of no more than 5%. This function groups similarly sized disk drives into one truncated size. For example, this allows the user to mix similarly sized drives from different vendors and obtain the same virtual capacity for all. The following example demonstrates drive sizing.

Given two vendors and a block size of 512 bytes.

Vendor A’s 18 GB drive has an actual capacity of 35,843,670 blocks and is marketed as 18.35 GB.

Vendor B’s 18 GB drive has an actual capacity of 35,885,167 blocks and is marketed as 18.37 GB.

After applying the drive sizing truncation algorithm, the virtual drive capacity is 17,166 MB for both vendors’ 18 GB drives. The 17,166 MB value is based on 1 MB = 1024^2, where disk drive manufactuerers report drive sizes based on 1 MB = 10^6. Using the second method to report drive sizes can overstate a drive’s actual apacity and result in large discrepancies between the marketing value and the engineering value.

Disk drives previously configured using a drive size file, available only with firmware versions prior to 7.0, maintain that configured size until the configuration is cleared. Once the configuration is cleared, all devices are considered unconfigured offline and the drive size is set according to the 7.0 drive size truncation function. The drive size may be different using drive size truncation rather than the configured drive size file.

When an online spare disk drive is selected for a rebuild procedure, the drive size truncation is ignored. An online spare disk drive is selected based on the physical size of the disk drive.

➢ See also “Automatic Rebuild” on page 10, “Configuration on Disk (COD)” on page 18, “Failed Disk Drive” on page 118, “Firmware Versions” on page 95, “Gigabyte (GB)” on page 120, “Global Array Manager (GAM)” on page 45, “Hot Spare” on page 122, “Kilobyte (KB)” on page 123, “Megabyte (MB)” on page 129, “Online Spare” on page 132, “Physical Disk Drive” on page 135, “RAIDfx” on page 53, “Rebuild” on page 144, “Replacement Drives” on page 145, “Terabyte (TB)” on page 150, and “Upgrading a Configuration Created with RAIDfx” on page 101.

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers. The terms used by the firmware and configuration utilities have evolved through the various firmware versions and releases. The terms used by firmware versions 7.0 and greater are listed and the equivalent term for previous firmware versions are given in parenthesis.

The controller provides information to the configuration tool or utility to assist in maintaining operational information for each physical disk drive. The controller keeps a record of each physical disk drive's operational state and information on any available target ID addresses. The controller first determines the presence or absence of the physical disk drive or available target ID and then determines the operational state of disk drives present. If the controller determines the disk drive to be present, the location of the disk drive is considered configured and the operational state of the disk drive is then determined. If the controller determines the disk drive at the available target ID location to be absent, the location of the disk drive is considered unconfigured and the operational state is determined as unconfigured offline or dead.

If using firmware versions earlier than 7.0, the controller considers any available target ID locations to be configured and any disk drives later inserted into the available locations will be determined to be standby. Empty locations should be manually set to dead before the configuration is saved to the controller.

If a configured disk drive is removed or fails, and is then replaced with a new disk drive, that new disk drive is set to online spare. This allows the automatic rebuild feature to work for replaced drives. When a disk drive is inserted into the system, the AEMI, SAF-TE, or SES device recognizes that the drive has been replaced. If a configured disk drive fails and then the controller power is cycled, when the controller power returns, the disk drive remains offline failed or dead.

Unconfigured disk drives can be removed and the device state will remain unconfigured. New disk drives added to the system are considered unconfigured until used in a new configuration. Unconfigured disk drive fault lights (LEDs) are disabled and any insertion, removal, or errors related to these unconfigured devices do not result in fault light activity or error message generation.

If using firmware version 7.0 or greater, COD information is written, or updated, to all configured devices. Unconfigured devices receive no COD update, the information is all set to zeros.

Descriptions of possible physical device states are described in the following table. This information applies only to physical devices, not to logical devices, or system drives.

Device State Descriptions

Device State	Description of Device State
Online Optimal (Online)	The disk drive is powered on, has been defined as a member of a drive pack, and is operating properly.
Online Spare (Standby or Hot Spare)	The disk drive is powered on, is able to operate properly, and has been defined as a standby or hot spare.
Offline Failed or Unconfigured Offline (Dead or Unconfigured)	The disk drive is not present, is present but not powered on, is a newly inserted replacement disk drive, or fails to operate properly and is labeled offline by the controller (whether or not it has been defined as a member of a drive pack).

Device State Descriptions

Device State	Description of Device State
Online Rebuild (Rebuild)	The disk drive is in the process of being rebuilt, that is, data is being copied from the mirrored disk drive to the replacement disk drive in a RAID 1 or 0+1 array, or data is being regenerated via the exclusive OR (XOR) redundancy algorithm and written to the replacement disk drive in a RAID 3 or 5 array.
Unconfigured	This location is unconfigured.
Environmental	An environmental device is present at this address.

➢ See also “AEMI” on page 6, “Alarm Signal” on page 6, “Automatic Rebuild” on page 10, “Configuration on Disk (COD)” on page 18, “Configuration Tools and Utilities” on page 21, “Critical” on page 113, “Dead” on page 113, “Device States” on page 114, “Environmental Device” on page 117, “Logical Device States” on page 124, “Offline” on page 130, “Offline Failed” on page 131, “Online” on page 131, “Online Optimal” on page 132, “Online Rebuild” on page 132, “Online Spare” on page 132, “Physical Device” on page 133, “Physical Device States” on page 134, “Physical Disk Drive” on page 135, “Rebuild” on page 144, “Replacement Drives” on page 145, “Standby” on page 147, “Target ID (TID)” on page 149, and “Unconfigured Offline” on page 150.

Drive to Enclosure Map

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

The firmware provides enough drive information so that each disk drive can be mapped to the enclosure in which it resides.

Dual-Active Controller Configuration

SCOPE: This feature was implemented in firmware prior to 3.2 and is supported on all controllers. Some firmware versions were released that did not support dual-active controller configurations and include 4.0, 4.1, 5.0, and 6.0. In addition, the firmware must be type 4, 5, 6, or 7.

Dual-active controller configurations interconnect two *identical* controllers that share a common set of disk drives. In addition to increasing overall performance, this method allows a surviving controller to take over resources of a failed controller. This *failover* process is transparent to the host.

Dual-active controller support provides the system with the mechanisms for initializing two controllers to process host requests; establishing a controller-controller nexus (C-C nexus) for controller redundancy; detecting failure of a controller which is part of a C-C nexus and initiating a failover; and detecting the replacement of a failed controller, initiating a failback and re-establishment of a C-C nexus. The two active controllers can be connected to a single host. The two controllers communicate with each other to verify that both are functioning properly through a ping/acknowledgement sequence. Failure to acknowledge the ping triggers failover.

The drive channels also carry signals necessary to the controllers’ functions. For firmware versions earlier than 6.0, this requires all drive channels to be connected and terminated even if they do not contain any disk drives. External termination provides maximum fault tolerance by preserving the ability to hot swap failed disk drives and hot plug failed controllers. For firmware versions 6.0 and greater, this requires one drive channel to be interconnected between the dual-active controller pair.

The dual-active controller configuration also supports dual- or multiple-host communication offering the advantages of being able to sustain data access in the event of a host failure. If configured in a cluster or high availability environment, it is also able to sustain data access in the event of the failure of a server or an HBA. This configuration requires implementation of alternate path software on the server.

ATTENTION: **If two systems independently access the same volume of data, and the operating system does not support file locking, data corruption may occur. To avoid this, create two or more volumes (or LUNs) and configure each volume to be accessed by one system only.**

Dual-Active Requirements

When configuring a dual-active controller system, replacing a failed controller in an existing dual-active controller system, or upgrading from a simplex system to a dual-active controller system, both controller cards must be identical. Additional hardware and cabling requirements must also be satisfied in order to provide optimal operation. When replacing a failed controller or adding a second controller, the configuration is updated by the surviving controller when the replacement/added controller starts, or from the COD stored on the disk drives.

NOTE: Mylex recommends a complete data backup before upgrading hardware or firmware.

Firmware Requirements

NOTE: The Automatic Firmware Flash feature implemented in firmware version 8.0 ensures that dual-active controllers operate with the same firmware version, firmware type, and controller parameter settings.

The firmware requirements for dual-active controller systems are as follows:

- Both controllers must have the same firmware version.
- Both controllers must have the same dual-active firmware type (4, 5, 6, or 7).
- The following controller parameters must be set accordingly:
    - Disable *Force Simplex*,
    - Enable *Controller Present/Fault Signals* or *Duplex Fault Signals* (not required for FFx or FFx2 controllers),
    - *Controller Present/Fault Select*, set to A (DACSX only), and
    - Disable *Simplex—No Reset*.

Hardware Requirements

The hardware requirements for dual-active controller systems are as follows:

- Both controllers must have the same amount of memory.
- Both controllers must have the same number of host and drive channels.
- For firmware versions earlier than 6.0, all drive channels must be connected one to one. For example, Channel 0 on the primary controller (C0) must be connected to Channel 0 on the partner controller (C1), and channel 1 on C0 must be connected to Channel 1 on C1. This is true even when a drive channel does not have any other devices connected to it.
- For firmware versions 6.0 and greater, one drive channel must be connected one to one between the dual-active controller pair.
- Hot plugging a replacement controller in a dual-active controller configuration requires that the controllers' dual-active signals be isolated. Refer to the appropriate installation manual for details.

ATTENTION: **Do not install an FF2 controller in a dual-active configuration with a DACFF (FF) controller. The current configuration and data may be lost.**

NOTE: The FF2 controller is very similar to the DACFF (FF) controller, however, the two controllers are not interchangeable. A dual-active controller configuration must consist of two identical controllers. A DACFF (FF) system can be upgraded to an FF2 system following the procedure described under "Upgrading a DACFF (FF) System to an FF2 System" on page 102.

Cabling Requirements

The cabling requirements for dual-active controller systems are as follows:

- Disk drives must be connected in series to the same drive port on both distribution boards; for example, disk drives to Channel 0 on Controller 0 and Channel 0 on Controller 1.
- Host ports on both controllers must be interconnected for transparent failover to take place; that is, Host Port 0 must be connected to Host Port 0, and Host Port 1 must be connected to Host Port 1.
- The dual-active communications ports on both controllers must be connected using a dual-active (heartbeat) cable. The dual-active cable is not needed for dual-active FFx and FFx2 controller systems.
- All SCSI drive and host channels must be terminated externally—not on the devices—to ensure access to data in the event of a disk drive or controller failure.
- Fibre Channel drive and host channels must employ a hub LRC (Loop Resiliency Circuit) device to recover the loop in the event of a controller failure.

➢ See also "Alternate Path Software" on page 7, "Automatic Firmware Flash" on page 7, "Channel" on page 112, "Configuration on Disk (COD)" on page 18, "Controller-Controller Nexus (C-C Nexus)" on page 23, "Controller Parameters" on page 69, "Controller Specifications" on page 1, "Drive Channel" on page 27, "Duplex Fault Signals" on page 77, "Failed Controller" on page 118, "Failover and Failback" on page 35, "Failover Topology, Inactive Port" on page 37, "Failover Topology, Master/Slave" on page 38, "Failover Topology, Multiple Target ID" on page 39, "Failover Topology, Multiport" on page 40, "Fault Tolerance" on page 118, "Soft Addressing Detection" on page 61, "Firmware Versions" on page 95, "Force Simplex" on page 78, "Host Bus Adapter (HBA)" on page 120, "Host Channel" on page 121, "Host / Server" on page 121, "Hot Plugging" on page 121, "Hot Swapping" on page 122, "Hub" on page 122, "Loop Redundancy Circuit (LRC)" on page 128, "Microsoft Cluster Server Configurations (MSCS)" on page 47, "Partner Controller" on page 133, "Primary Controller" on page 136, "Redundant Array" on page 145, "Replacement Controller" on page 145, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, "SCSI Cabling" on page 99, "SCSI Termination" on page 100, "Secondary Controller" on page 146, "Simplex" on page 60, "Simplex - No Reset" on page 86, "Surviving Controller" on page 148, and "Upgrading From Simplex to Dual-Active Configurations" on page 103.

Dual-Ported Drive Support

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

Failover and failback operations of up to 124 multi-ported drive IDs are supported. The disk drives must be assigned hard IDs for this feature to operate successfully. All failover/failback and addressing activity is transparent to the host.

Each dual-ported drive can be connected to two drive channels. When the controller firmware looks for drives, one of the drive ports is assigned a *preferred* (channel-target ID) address; the other is assigned an *alternate* address. When a drive channel fails, if that channel is associated with the preferred address, the drive is accessed through its alternate address. When the failed channel is restored, the drive is again accessed through its preferred address.

If a controller starts with a failed drive channel, then all drives are accessed through the available channel. This channel becomes the preferred address by default. When the failed channel is restored, it becomes the alternate channel or alternate address.

In dual-active dual-ported drive configurations, both controllers will start, provided at least one accessible drive channel is active.

If a single drive on a loop fails, the drive can be rebuilt through normal rebuild operations. This is true even in the case of a single port failure on a drive.

NOTE: It is possible to configure drive channels such that the alternate channel may be different for individual drives on the same preferred channel.

➢ See also “Alternate Address” on page 110, “Alternate Channel” on page 110, “Channel” on page 112, “Controller Parameters” on page 69, “Drive Channel Failover and Failback” on page 28, “Dual-Active Controller Configuration” on page 32, “Failed Disk Drive” on page 118, “Soft Addressing Detection” on page 61, “Hard Loop ID” on page 120, “Host / Server” on page 121, “Loop Address (Loop ID)” on page 125, “Port” on page 135, “Preferred Address” on page 135, “Preferred Channel” on page 136, “Rebuild” on page 144, and “Target ID (TID)” on page 149.

Embedded Configuration Utility

SCOPE: This feature was implemented in firmware version 7.7 and is supported on the FFx, FF2, and FFx2 controllers.

The embedded configuration utility is an on-board configuration utility accessible through a serial port to a terminal-emulation utility such as VT100. The embedded configuration utility is also accessible in-band over the Fibre Channel interface. The embedded configuration utility allows the user to create unique RAID configurations specific to individual needs. Refer to the *Embedded Configuration Utility User’s Guide* for more information on using the embedded configuration utility.

➢ See also “Configuration Tools and Utilities” on page 21, “Fibre Channel” on page 43, and “VT100” on page 65

Failover and Failback

SCOPE: This feature was implemented in firmware prior to 3.2 and is supported on all controllers.

In dual-active controller configurations, maintaining continuous access to data requires that a failed controller be replaced in a manner that is transparent to the host. Two different methods are used to implement the failover/failback operation. Both SCSI and Fibre Channel controllers use the processor’s ability to respond to multiple target IDs to implement the failover/failback process. The Fibre Channel controllers can also configure ports on the controllers to be *inactive* ports until such time as those ports are needed to respond to requests of a failed controller. The FFx, FF2 and FFx2 controllers support multiple target IDs per port.

In the event of a controller failure in a dual-active controller system, the failed controller's operations are assumed by the surviving controller. The failed controller can then be removed and replaced while the system is online. When the surviving controller detects the presence of the new controller, the new controller resumes processing array operations. During failover and failback, write cache coherency is maintained with the disk drives.

When a controller is participating in a C-C nexus and detects a communication error with its partner controller, it initiates the failover process. The following steps outline the failover process executed by the surviving controller:

1. On detection of the controller failure, the surviving controller holds the failed controller in reset.
2. The surviving controller activates the failover port or enables its secondary ID for the partner controller.
3. Cached data is flushed to the disk drives.
4. Conservative cache is enabled (if the *Conservative Cache Mode* controller parameter has been enabled).
5. The surviving controller begins handling I/O requests for the failed controller.

When a failed controller is replaced, the system either automatically detects the replacement (if configured for *Automatic Failback*) or is informed of the replacement by the Relinquish Controller/Partner direct command. The following steps outline the failback process executed by the surviving controller:

1. A replacement controller is detected.
2. The surviving controller releases its partner from reset.
3. Once the replacement controller completes initialization and is ready to resume I/O requests, the surviving controller quiesces both ports by responding with BUSY status to new I/O requests.
4. The surviving controller disables the failover port or secondary ID.
5. The surviving controller clears the BUSY condition.
6. The replacement controller enables its primary ports or primary ID.
7. Both controllers disable conservative cache (if enabled) for write-back system drives and resume normal dual-active controller operation.

➢ See also "Active Controller" on page 109, "Active Fibre Port" on page 109, "Automatic Failback or Autorestore" on page 70, "Background Initialization" on page 11, "Cache Flush" on page 112, "Configuration on Disk (COD)" on page 18, "Conservative Cache Mode" on page 22, "Conservative Cache Mode Parameter" on page 72, "Consistency Check" on page 23, "Controller-Controller Nexus (C-C Nexus)" on page 23, "Controller Parameters" on page 69, "Dual-Active Controller Configuration" on page 32, "Failed Controller" on page 118, "Failover Port" on page 118, "Fault Tolerance" on page 118, "Host / Server" on page 121, "Hot Plugging" on page 121, "Inactive Fibre Port" on page 123, "Multiple Target ID (MTID)" on page 48, "Partner Controller" on page 133, "Port" on page 135, "Primary ID" on page 136, "Programmable LUN Mapping" on page 52, "Replacement Controller" on page 145, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, "Replacing a Failed Controller in Simplex Systems" on page 99, "SANmapping™" on page 56, "SCSI Cabling" on page 99, "SCSI Termination" on page 100, "Secondary ID" on page 146, "Surviving Controller" on page 148, "System Drive Affinity" on page 64, and "Write-Back Cache" on page 151.

SCOPE: This feature was implemented in firmware version 4.2. This failover topology is only recommended for the DACSF, DACFL, DACFF (FF), and FF2 controllers.

In a Fibre Channel inactive port failover topology, one port is active and the other is inactive on each controller. This topology is shown in the following illustration.

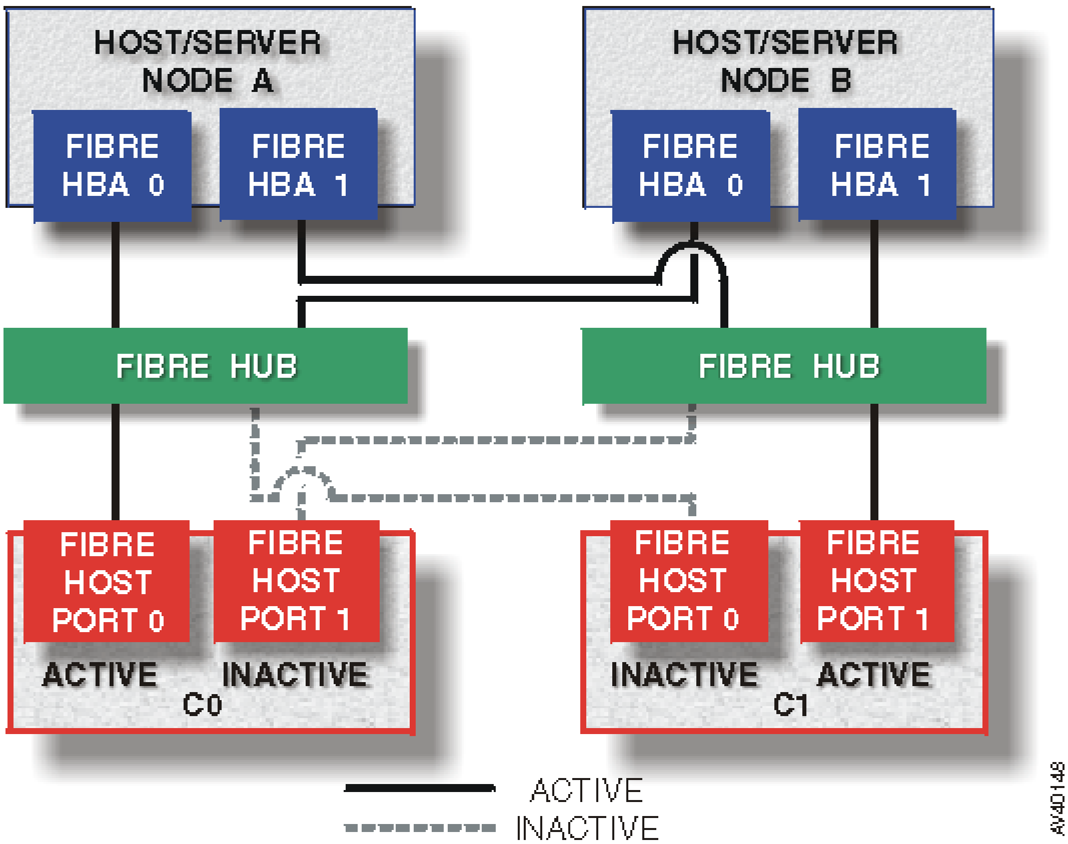

Dual-Active Controller Inactive Port Topology

A different port is used on each controller to allow each controller to process I/O requests independent of the other. This active/inactive relationship allows for transparent failover of I/O requests to the surviving controller in the event of a controller failure. The dual-active controllers communicate across a serial data link. The inactive port on each controller serves as the failover port for the partner controller. When failure occurs, the surviving controller enables its inactive Fibre Channel port and begins servicing I/O requests on that loop. When the failed controller is replaced, it automatically restarts if the *Duplex Fault Signals* have been implemented and enabled in the controller parameters. If not, a relinquish controller command is issued to the controller to allow the replaced controller to start up and resume normal I/O processing.

The active/inactive configuration requires that each system drive's affinity be mapped only to a single port. System drive affinity should not be mapped to inactive ports, the firmware transparently adjusts system drive affinity if a failover occurs.

➢ See also “Active Fibre Port” on page 109, “Failed Controller” on page 118, “Failover and Failback” on page 35, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Inactive Fibre Port” on page 123, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, “Surviving Controller” on page 148, and “System Drive Affinity” on page 64.

Failover Topology, Master/Slave

SCOPE: This feature was implemented in firmware version 5.0 and applies to the Mylex Fibre Channel controllers. The master/slave topology is a legacy topology and is not recommended for use with firmware versions greater than 5.4.

The master/slave failover topology requires alternate path software and fully redundant paths to the system drives in order to redirect I/O in case of a cable and/or host bus adapter failure. This topology is shown in the following illustration.

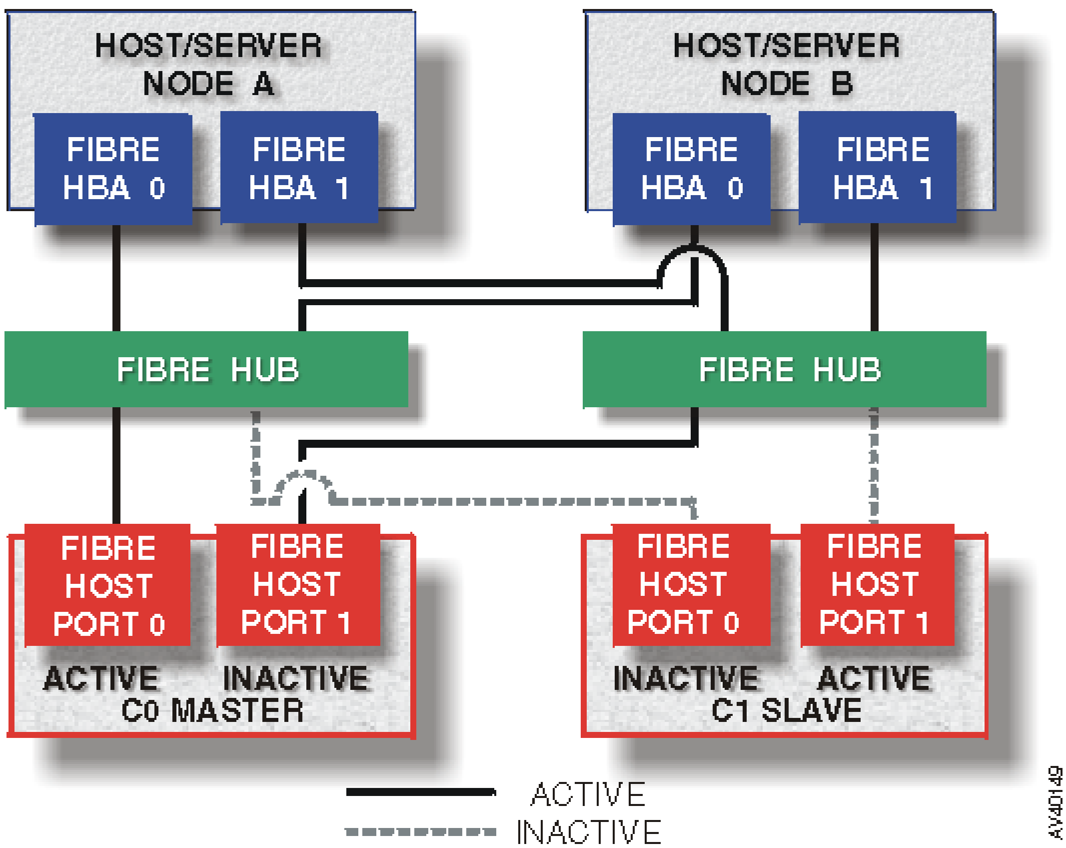

Dual-Active Controller Master/Slave Topology

This configuration is dependent upon the fibre HBA driver to handle the failure of a cable or HBA. The controller pair still handle a controller failure transparent to any connected hosts. During normal operation, only one controller has both fibre ports active and available to service requests from the host. The inactive controller’s ports serve as failover ports for the partner controller. All write cache data is mirrored onto the other controller. If the active controller fails, the surviving controller activates both fibre host ports, impersonating the corresponding ports from the failed controller by using the same loop ID and begins servicing requests.

The master/slave configuration requires that system drive affinity be mapped only to the master controller. Changing between master/slave and active/inactive configurations will require changes to system drive LUN mapping. The Fibre Channel hard loop IDs must be the same for both ports.

➢ See also “Active Controller” on page 109, “Active Fibre Port” on page 109, “Failed Controller” on page 118, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Hard Loop ID” on page 120, “Host Bus Adapter (HBA)” on page 120, “Inactive Fibre Port” on page 123, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, and “System Drive Affinity” on page 64.

Failover Topology, Multiple Target ID

SCOPE: This feature was implemented in firmware version 6.0 and is supported only on the FFx, FF2, and FFx2 controllers.

The multiple target ID topology provides the controller with the ability to function as multiple target ports on a single arbitrated loop by acquiring and using multiple loop IDs and multiple WWNs. The Multiple Target ID feature of a Fibre Channel controller is based on the concept of a virtual port. This topology is shown in the following illustration.

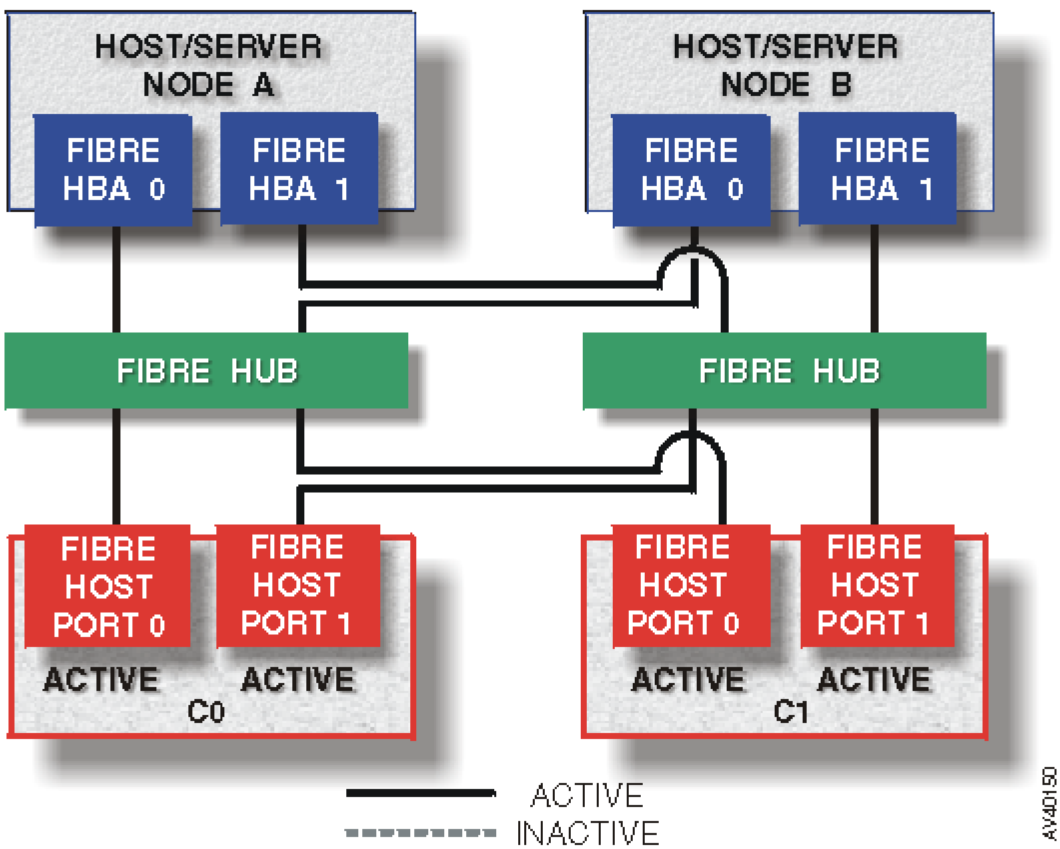

Dual-Active Controller Multiple Target ID Topology

In the event of a controller failure, the port on the surviving controller enables a virtual port, impersonating the port from the failed controller. This action is repeated for all host ports on a given controller. One drawback to a multiple target ID Fibre Channel topology is that only two fibre loops are available for hosts connection even though all four host ports are active. This limits the bandwidth capability of the controller.

The multiple target ID configuration requires that each system drive's affinity is mapped only to a single port.

When using the FFx2 controller, the alternate path software options must be set to *Loop Only.*

➢ See also "Active Fibre Port" on page 109, "Alternate Path Software" on page 7, "Bandwidth" on page 111, "Configuration Strategies" on page 94, "Controller Parameters" on page 69, "Dual-Active Controller Configuration" on page 32, "Failed Controller" on page 118, "Failover and Failback" on page 35, "Failover Topology, Inactive Port" on page 37, "Failover Topology, Master/Slave" on page 38, "Failover Topology, Multiport" on page 40, "Fibre Channel" on page 43, "Soft Addressing Detection" on page 61, "Hard Loop ID" on page 120, "Host Port" on page 121, "Inactive Fibre Port" on page 123, "Loop Address (Loop ID)" on page 125, "Primary Port" on page 137, "SANmapping™" on page 56, "Surviving Controller" on page 148, "Virtual Port" on page 151, and "World Wide Name (WWN)" on page 151.

Failover Topology, Multiport

SCOPE: This feature was implemented in firmware version 5.0 and is supported on all Mylex Fibre Channel controllers.

The multiport topology requires alternate path software to redirect communication during a component failure. This topology is shown in the following illustration.

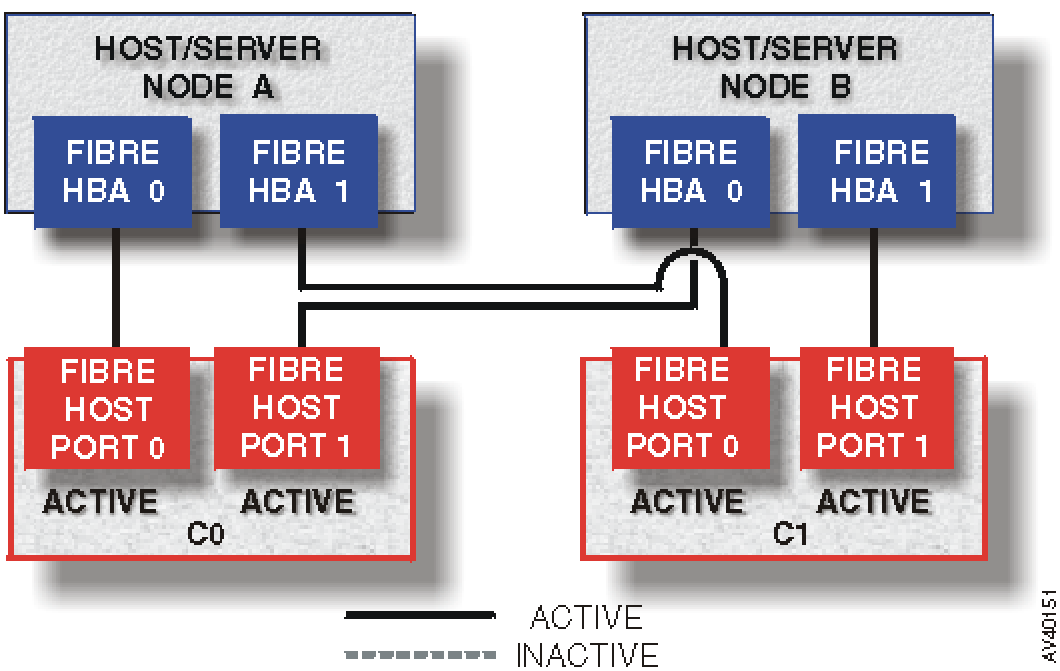

Dual-Active Controller Multiport Topology

The alternate path software handles I/O redirections in the event of Fibre Channel HBA or cable failure. In the event of a path failure, the software redirects I/O through the alternate path or HBA. The software is aware that the same device (Mylex controller) is accessible through multiple paths or HBAs in the same host. The software recognizes and uses a unique ID on each system drive to manage the multiple paths. The use of alternate path software eliminates the need for hubs.

All ports are active in this configuration. Both ports on controller 0 use the same WWN (node name), but vary the port name by reporting a unique port number. Both ports on controller 1 use the same naming convention, but utilize a WWN that is unique from controller 0. This topology has a distinct performance advantage over the multiple target ID topology, in that all four controller host ports may be connected to individual fibre loops, allowing for greater host connection bandwidth. Controller failover/failback is not transparent to the host using the multiport topology.

The multiport configuration requires that system drive affinity be mapped to all controller/host ports, the default setting when a new configuration is created.

When using the FFx2 controller, the alternate path software options must be set for *Loop Preferred*. This enables the switch to force point-to-point connections.

➢ See also “Active Controller” on page 109, “Active Fibre Port” on page 109, “Alternate Path Software” on page 7, “Failed Controller” on page 118, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Fibre Channel” on page 43, “Hard Loop ID” on page 120, “Host Bus Adapter (HBA)” on page 120, “Inactive Fibre Port” on page 123, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, and “System Drive Affinity” on page 64.

Fault Management

SCOPE: Support for AEMI monitoring was implemented in firmware versions prior to 3.2 and is supported on the DACSX controller. Support for SAF-TE monitoring was implemented in firmware version 3.2 and is supported on the DACSX, DACSS, DACSF, and DACFL controllers. Support for SES monitoring was implemented in firmware version 5.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

Controller fault management is implemented using one of the following:

- The Array Enclosure Management Interface (AEMI) (DACSX only)
- SCSI Accessed Fault-Tolerant Enclosure (SAF-TE) interface
- SCSI-3 Enclosure Services (SES) device interface

The following table lists the controllers on which each of the fault monitoring systems are implemented.

Fault Monitoring System Applications

	AEMI	SAF-TE	SES
DACSX	X	X	
DACSS		X	
DACSF		X	
DACFL		X	
DACFF (FF)			X
FFx			X
FF2			X
FFx2			X

AEMI, SAF-TE and SES monitor the same devices: fans, power supplies, temperature sensors and UPSs, as well as disk drive status.

In order for any of the monitoring systems to be enabled, the *Operational Fault Management* controller parameter must be enabled. If *Operational Fault Management* is disabled, SAF-TE and SES fault management are disabled.

Regardless of whether the system is monitored by SAF-TE or SES, fault conditions for fans and power supplies are reported for up to six units. Out-of-limit temperature conditions are reported as a single failure.

Drive locate is a fault management feature that allows the user to flash LEDs of individual disk drives and/or LUNs for 30 seconds. This feature is used to locate failed disk drives or critical LUNs. The locate feature is also used to locate physical devices, RAID devices, COD groups or the entire enclosure.

On a DACSX, if the controller detects a SAF-TE device during power-up and *Operational Fault Management* is enabled, both SAF-TE and AEMI input signals are monitored for fan, power supply, and out-of-limit temperature conditions. If no SAF-TE device is found on power-up and *Operational Fault Management* is enabled, the controller uses AEMI for fault management.

NOTE: Note that AEMI Fan0 and SAF-TE Fan0 are reported as a single fault on Fan0. Power supply faults are reported in the same manner.

➢ See also “AEMI” on page 6, “Channel” on page 112, “Configuration on Disk (COD)” on page 18, “Controller Parameters” on page 69, “Environmental Device” on page 117, “Operational Fault Management” on page 82, “Physical Device” on page 133, “SAF-TE” on page 55, and “SES” on page 58.

FF2

A Mylex disk array controller with Fibre Channel host interface and Fibre Channel disk drive interface. The FF2 supports two host channels and four drive channels. The FF2 utilizes firmware version 6.1 through 7.7. The FF2 utilizes fibre protocol processors that provide dual loop support.

ATTENTION: **Do not install an FF2 controller in a dual-active configuration with a DACFF (FF) controller. The current configuration and data may be lost.**

The FF2 controller appears similar to the DACFF (FF) controller, however, the two controllers are not interchangeable. A dual-active controller configuration must consist of two identical controllers. A DACFF (FF) system can be upgraded to an FF2 system following the procedure described under “Upgrading a DACFF (FF) System to an FF2 System” on page 102.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “DACFF (FF)” on page 24, “Drive Channel” on page 27, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, and “Host Channel” on page 121.

FFx

A Mylex disk array controller with Fibre Channel host interface and Fibre Channel disk drive interface. The FFx supports one host channel and two drive channels. The FFx controller is designed to be compatible with housings and connectors that meet the 3.5-inch Small Form Factor

(SFF) committee standards in width and height, allowing the FFx to be installed into a drive slot. The FFx utilizes firmware versions 6.0 through 7.7.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, and “Host Channel” on page 121.

FFx2

A Mylex disk array controller with Fibre Channel host interface and Fibre Channel disk drive interface. The FFx2 controller is the second generation of Mylex controllers designed to be compatible with housings and connectors that meet the 3.5-inch Small Form Factor (SFF) committee standards in width and height. The FFx2 supports two host channels and two drive channels. The FFx2 also supports 2 Gbps transfer speed across the host interface and the drive channel interface. The FFx2 utilizes firmware version 8.0.

➢ See also “Channel” on page 112, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, and “Host Channel” on page 121.

Fibre Channel

SCOPE: This feature was implemented in firmware version 4.0 and is supported on the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

A data transfer interface technology that allows for high-speed I/O and networking functionality in a single connectivity technology. The Fibre Channel standard supports several topologies, including Fibre Channel point-to-point, Fibre Channel fabric (generic switching topology), and Fibre Channel arbitrated loop (FC-AL).

Fibre Channel Arbitrated Loop (FC-AL)

A connection between multiple nodes. The node ports are interconnected so as to form a closed loop. Information is directed around the loop until the final destination is reached. The information is routed through and repeated by intermediate ports on the loop. Node ports that perform this loop technology, routing and repeating, are known as Node Loop Ports (NL_Ports). All the ports on the loop share the loop. Access to the loop is obtained through an arbitration process.

Fibre Channel Fabric

A Fibre Channel topology that uses multiple switches to route frames between nodes in a Fibre Channel network. The routing of frames is transparent to the nodes or devices.

Fibre Channel Point-to-Point

The simplest Fibre Channel connection. This provides a direct connection between the transmit output of one node and the receive input of a second node. A second connection is provided between the opposite connectors in order to complete the signals. The physical connection between the two nodes is called a link. No switches, loops, or fabric elements are needed.

➢ See also “Controller Specifications” on page 1, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Frame” on page 119, “Link” on page 124, “Microsoft Cluster Server Configurations (MSCS)” on page 47, “Node” on page 129, “Node Loop Port (NL_Port)” on page 129, and “Node Port (N_Port)” on page 130.

Firmware Header Information

SCOPE: The firmware header information varies with each firmware release.

The firmware header contains three fields that modify the behavior of the firmware: Type, Misc and Debug. The Type field is used to control the system configuration. The Misc field is used to control special features. The Debug field is used for development only.

The Type field defines the system configuration, for example, Type 5 supports dual-active controller configurations. The Type field is displayed in the upper right hand corner of the LCD and in the banner across the top of the VT100 screen. The following table defines Types 0 through 7 and lists the products with which each type can be used.

Firmware Types

Firmware Type	Description	Product
0	Single host, no UNIX support, no dual-active controller support	DACSX
1	Dual-host, no UNIX support, no dual-active controller support	DACSX DACSS, DACSF, DACFL, DACFF (FF), FF2
2	Single host, UNIX support, no dual-active controller support	DACSX
3	Dual-host, UNIX support, no dual-active controller support	DACSX, DACSS, DACSF, DACFL, DACFF (FF), FF2
4	Dual-active controller support, single host, no UNIX support	DACSX
5	Dual-active controller support, dual-host, no UNIX support	DACSX, DACSS, DACSF, DACFL, DACFF (FF), FFx, FF2, FFx2
6	Dual-active controller support, single host, UNIX support	DACSX
7	Dual-active controller support, dual-host, UNIX support	DACSX, DACSS, DACSF, DACFL, DACFF (FF), FFx, FF2, FFx2

NOTE: The FFx and FFx2 do not support firmware types 0,1,2, 3, 4, or 6. Single controller support is configured by the *Force Simplex* controller parameter.

Single host and dual host support is ignored by controllers with Fibre Channel host support (DACSF, DACFL, DACFF, FFx, FF2, and FFx2).

Dual-host configurations of the DACSX require a daughterboard.

The Misc field is used by firmware to enable various alarm features, enable dual-active fault signals, control LCD displays, and implement other version specific modifications.

The header fields can be customized using RAIDfx. Refer to the *RAIDfx Manager User Guide*. The RAIDfx help menus also provide definitions of the available settings. The Fwzap utility is also available to customize this header information.

➢ See also “Dual-Active Controller Configuration” on page 32, “Duplex” on page 116, “Embedded Configuration Utility” on page 35, “Firmware Versions” on page 95, “Global Array Manager (GAM)” on page 45, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, “Simplex” on page 60, and “UNIX Support and Firmware Types” on page 64.

Firmware Variables

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This feature was modified in firmware version 7.7 to add the WWN Source variable and is supported on the FFx, FF2, and FFx2 controllers.

A section added to the firmware image to set variables used by the controller. These variables are initialized once during the image creation process and can be modified using a utility to set values based on customer requirements. These values are used by firmware only and software applications have no visibility to these values. The following variables are currently supported.

- Controller 0's Fibre ID Value—This value defines the fibre loop ID used by Controller 0 on the drive channel.
- Controller 1's Fibre ID Value—This value defines the fibre loop ID used by Controller 1 on the drive channel.
- Online Spare Polling Interval—This value determines the number of seconds between online spare polling intervals. The default value is set to 60 seconds. A value of 0 disables online spare polling.
- WWN Source—This value determines the source from which the controllers' WWN is determined. The choices are 0=Unique WWN, 1=Common WWN, and 2=Enclosure WWN.
- Vendor Name—This variable sets the vendor name.
- Vendor Model—This variable sets the vendor model.

➢ See also "Hard Loop ID" on page 120, "Loop Address (Loop ID)" on page 125, "Online Spare Polling" on page 51, and "World Wide Name (WWN) Assignments" on page 65.

Global Array Manager (GAM)

SCOPE: This software was supported in firmware versions prior to 3.2 and is supported on all controllers.

Global Array Manager (GAM) is a Mylex software utility that provides configuration capabilities to the Mylex controllers. GAM is a GUI supported interface that allows the user to monitor, manage, maintain and configure Mylex controllers and the physical devices connected to the controllers. GAM has two components—the GAM Server and the GAM Client.

Refer to the *Global Array Manager Installation Guide and User Manual* for complete installation and user instructions.

GAM Server

The Server collects and disseminates information on disk array or host bus adapter (HBA) subsystem status and resource utilization. The Server executes the management instructions specified by the Client. The Server functions on Windows NT, Windows 2000, Netware, Unixware, Linux, SCO, and Solaris x86 operating systems.

GAM Client

Configuration functions are easily performed using RAID Assist™, a wizard-like utility in the Client that simplifies the process of setting up or reconfiguring a disk array. The Client organizes the information collected by the server through a graphical display. Errors and events are recorded in a log, and if a problem is serious enough to warrant immediate attention, operators can be alerted via pop-up windows, pagers, fax, or email. The Client also manages or performs maintenance on individual disk arrays and disk drives. This maintenance can include removing physical devices from operation in a functioning disk array, rebuilding disk drives, selecting spare disk drives, and initiating a consistency check on redundant array systems. The Client functions on Windows 95, Windows 98, Windows NT, and Windows 2000 operating systems.

➢ See also “Configuration on Disk (COD)” on page 18, “Configuration Tools and Utilities” on page 21, “Embedded Configuration Utility” on page 35, “Firmware Versions” on page 95, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, “Upgrading a Configuration Created with RAIDfx” on page 101, and “VT100” on page 65.

Implied Reserve and Release

SCOPE: This feature was implemented in firmware version 6.1, and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

Instead of relying on the host to explicitly reserve a system drive prior to accessing it, the controllers can implicitly reserve (lock) the system drive, or portions of the system drive for a host with every host access request. The implicit reservations for a portion of the system drive are referred to as locks to distinguish them from SCSI device reservations. The SCSI RESERVE and RELEASE commands and associated flushing and cache invalidation are still supported, but they are no longer required to support data consistency.

The implied reserve and release feature requires a controller to reserve a section of a system drive in order to perform updates to that system drive. An explicit reservation would be performed through the use of a system drive reserve command. An implicit reservation would be acquired when a write operation requires that the controller obtain the reservation. An implicit reservation occurs when an alternate active controller already owns the reservation or when the controller is required to perform a read operation and the alternate controller owns the reservation. This ensures that the alternate controller’s cache contains no dirty data that is not synchronized with the drive.

The reservation process is synchronized between all controllers to ensure the reservation table is coherent. Obtaining a reservation currently owned by an alternate active controller requires the alternate controller flushes and invalidates all cache data that is associated with the system drive. Releasing a reservation is not required, but may be performed by a system drive release command.

➢ See also “Cache” on page 111, “Cache Coherency” on page 17, “Cache Flush” on page 112, and “Host / Server” on page 121.

Intelligent Cache Mirroring

SCOPE: This feature was implemented in firmware version 5.4 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

This feature reduces the amount of data being transferred from one controller to the other via the cache copy operation. This improvement reduces the load on the channels used for cache copies. This feature also works in conjunction with parallel host writes to improve the overall effects of cache copies by coalescing numerous host writes to a write-back system drive into a single cache copy operation.

➢ See also “Cache” on page 111, “Channel” on page 112, “Controller Parameters” on page 69, “Device Combing” on page 73, “Parallel Host Writes” on page 51, and “Write-Back Cache” on page 151.

Liquid Crystal Display (LCD)

SCOPE: This feature was supported on the DACSX, DACSS, DACSF, DACFL, and DACFF (FF) controllers using firmware versions 3.2 through 5.0.

The optional front panel of the controller enclosure provides the user with an operational interface with a keypad and the LCD provides a monitor display for error and status messages. Limited support for the LCD functionality was provided with firmware versions 5.0 and greater. As of firmware version 7.0, the LCD front panel provides status messages only and the front panel keypad utility is no longer functional.

➢ See also “Alarm Signal” on page 6, “Configuration Tools and Utilities” on page 21, “Embedded Configuration Utility” on page 35, “Firmware Header Information” on page 44, “Global Array Manager (GAM)” on page 45, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, and “VT100” on page 65.

LUN Mapping

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on the DACSX, DACSS, DACSF, DACFL, and DACFF (FF) controllers. This feature was incorporated into SANmapping™ with firmware version 6.0.

A method whereby a LUN ID is assigned to a system drive, allowing a LUN to be made accessible through one or more host ports. The LUN assignments are per host port and are independent of the assignments on other host ports. A system drive may be assigned only one LUN per host port.

By not assigning a LUN to a system drive on a particular host port, that system drive is made inaccessible to that host port.

➢ See also “Host Port” on page 121, “Logical Unit Number (LUN)” on page 125, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, “System Drive Affinity” on page 64, and “System Drives” on page 149.

Microsoft Cluster Server Configurations (MSCS)

SCOPE: This feature was implemented in firmware version 4.2. The DACSX, DACSF, DACFL, and FFx controllers have been MSCS certified.

Microsoft Windows NT Enterprise Edition Cluster Server (MSCS) recommends the following configurations for use in MSCS clusters:

Configurations for FFx (Single Host Port)

Simplex Shared Loop. The two server nodes and the controller’s host port are attached to the same fibre loop. All system drives have affinity to the single host port.

Active-Active Shared Loop. Both controllers’ host ports are connected in a loop with both servers’ host bus adapters (HBAs). System drive affinities are set to either controller 0 or controller 1. This topology takes advantage of transparent failover. A controller failure will not cause a host failure.

NOTE: A loop failure causes both servers to fail.

Configurations for DACSX, DACSF, and DACFL (Dual Host Ports)

Simplex Point-to-Point. In a system with a simplex controller and two host ports, each host port is directly connected to a server HBA. The system drive affinities are set so all system drives have affinity to both ports.

Reset propagation is enabled for DACSX, and optional for the other controllers.

Active-Active Point-to-Point. One server is directly connected to controller 0 port 0 (C0P0); the second server is connected to controller 1 port 1 (C1P1). All system drives have affinity to all ports. It is the lowest cost active-active topology, but does not take advantage of transparent failover. In addition, a controller failure will cause an MS cluster server to fail. This topology provides a no-single-point-of-failure solution.

Active-Active Shared Loop. All four host ports, two on each controller, and both servers' HBAs are attached to a hub forming a single loop. System drive affinities are set to one of the four host ports. This topology uses transparent failover.

Active-Active Dual Loop. One MS cluster server is connected via a hub in a loop to C0P0 and C1P0. The second MS cluster server is connected via a hub in a loop to C0P1 and C1P1. Set affinities for all system drives so that each system drive is attached exactly once to each loop. This topology provides a no-single-point-of-failure solution.

Reserve / Release Support

The FFx controller complies with SCSI-3 standards for reservations in devices with multiple service delivery ports. If an initiator attached to a port reserves a system drive, access to the system drive is blocked for other initiators attached to the same port, and for all initiators attached to other ports. Multiple port configurations include simplex with two host ports and dual-active with two active host ports. SCSI reserve and release commands are used by MSCS to manage access to disk resources shared by the cluster.

Reset Coordination

The FFx controller complies with SCSI-3 standards for response to reset events (for example, TARGET RESET and LOGICAL UNIT RESET) for devices with multiple service delivery ports. Reset events are coordinated among host ports that share affinities to system drives. For example, if a reset event is detected by controller 0 port 0, that port is reset. All system drives with affinity to controller 0 port 0 process the reset event. If those system drives have affinity to controller 0 port 1, controller 0 port 1 will also process the reset event. To an initiator attached to controller 0 port 1, it appears that the shared system drives have received a reset event from another initiator.

➢ See also "Active Fibre Port" on page 109, "Cache Coherency" on page 17, "Configuration Strategies" on page 94, "Controller Parameters" on page 69, "Fibre Channel Arbitrated Loop (FC-AL)" on page 43, "Host Bus Adapter (HBA)" on page 120, "Hub" on page 122, "Implied Reserve and Release" on page 46, "Inactive Fibre Port" on page 123, "LUN Mapping" on page 47"Programmable LUN Mapping" on page 52, "SANmapping™" on page 56, "Simplex" on page 60, and "System Drive Affinity" on page 64.

Multiple Target ID (MTID)

SCOPE: This feature was implemented for SCSI controllers in firmware versions prior to 3.2 and for Fibre Channel controllers in firmware version 6.1. Multiple target IDs are supported on the DACSX, DACSS, FFx, FF2, and FFx2 controllers.

This feature allows a controller to respond to I/O requests directed to the partner controller, using a secondary ID, when the partner controller has failed. When a controller fails, the surviving controller takes over responsibility for servicing I/O requests on its secondary ID and continues to service its primary ID.

For the DACSX and DACSS controllers, the primary ID of the primary controller is acquired through the back plane connections. These signals are also available to the secondary controller. The secondary target ID of the primary controller is the primary target ID of the secondary controller, and the secondary target ID of the secondary controller is the primary target ID of the Primary controller.

➢ See also “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Failover Topology, Multiple Target ID” on page 39, “I/O” on page 123, “Primary ID” on page 136, “SCSI” on page 145, “Secondary ID” on page 146, “Surviving Controller” on page 148, and “Target ID (TID)” on page 149.

Mylex Online RAID Expansion (MORE)

SCOPE: This feature was implemented in firmware version 3.3 and is supported on all controllers.

The Mylex Online RAID Expansion operation adds capacity to an existing RAID set while the controller is online with the host. For example, a system using a five-disk-drive RAID set can add another disk drive to create a six-disk-drive RAID set. The MORE operation can be performed on all RAID levels except JBOD.

During the RAID set expansion process, which includes re-striping data from the old (smaller) RAID set to the new (expanded) RAID set, the controller continues to service host I/O requests.

MORE is supported in the simplex mode of operation only. One controller in a dual-active controller system must be disabled (failed over). An attempt to perform a MORE operation on a dual-active controller system is rejected.

The MORE operation provides two options for configuring the added capacity: Add System Drive and Enlarge System Drive.

The Add System Drive option adds a new system drive using the increased capacity rather than increasing the size of the existing system drive(s). The system drive's data is striped across a larger number of physical disk drives than before the expansion operation. The Add System Drive option is rejected if the maximum of 8 or 32 (firmware 7.0 or greater) system drives is defined. When the MORE operation is complete, the new system drive partition needs to be created and formatted according to the operating system in use.

The Enlarge System Drive option appends the added capacity to the system drive specified in the MORE operation request. Enlarge System Drive requires that only one system drive be defined using the set of physical disk drives. If more than one system drive exists on the same set of disk drives, the MORE operation is rejected. Creating and formatting the added capacity requires the current partition be deleted and the new, larger system drive be recreated. The Add System Drive option is rejected if the maximum of 8 or 32 (firmware 7.0 or greater) system drives is defined.

NOTE: Mylex recommends backing up all data on the system drive that is enlarged prior to performing the MORE operation. Following the MORE operation, the user must delete, recreate, and reformat the system drive partition(s) according to the operating system in use. The data can be saved back to the system drive following the operating system format process.

MORE Features

The following features define the MORE function:

- One to fourteen drives can be added to a RAID set at one time. The maximum number of physical disk drives in the new system drives must not exceed 8 or 16 (firmware 7.0 or greater).
- The minimum number of disk drives in the source RAID set is two, the maximum is 7 or 15 (firmware 7.0 or greater).
- The disk drive(s) being added must be configured as an online spare device(s).
- The disk drive(s) being added must not already be part of an array.
- The capacity of each of the added disk drives must be greater than or equal to that of the smallest disk drive in the RAID set.
- Adding capacity across spanned system drives is not supported. A MORE request to a system drive with more than one span is rejected.
- Using the Add System Drive option, a RAID set consisting of more than one system drive will process a MORE request on all system drives with one migration request, processing the system drives one at a time until all have been processed. Allowing one migration service to process all system drives on the same set results in one new extra capacity region being created when the migration is done. The new system drive will be added at the end. If 8 or 32 (firmware 7.0 or greater) system drives are defined prior to the migration request, the MORE request will be rejected.
- Using the Enlarge System Drive option, the system drive specified will be enlarged as well as having its data striped across a larger number of physical disk drives. For this option, the system drive specified must be the only system drive on the pack of disk drives being migrated.
- Add capacity input parameters and execution parameters are kept in NVRAM and COD. The process is restarted automatically by the firmware after a power failure.
- In the event of a disk drive failure, the process continues to completion in online critical mode. After migration is completed in online critical mode, then if the Automatic Rebuild controller parameter is enabled, and if an online spare disk drive is available, a rebuild starts immediately after migration is complete, restoring the online critical system drive back to an online optimal state. If however, the disk drive failure occurs during the initialization of the newly created system drive (Add System Drive option), the migration operation stops immediately and is considered complete. At this time the rebuild operation begins if applicable. The migration process is aborted in the event of a disk drive failure during the adding of capacity to a RAID 0 set or when a two-disk-drive failure occurs.
- In the event that a MORE process must be terminated (for example, two disk drives fail and cannot be recovered), the controller automatically aborts the migration process.
- MORE, Initialize, Rebuild, and Consistency Check are mutually exclusive operations. Only one process can run at a time.
- Data is not lost if the RAID controller becomes disabled (for example, powered off by mistake or a power supply goes bad). The process resumes after power up. Full recovery is done after a power failure and the operation continues with no data loss.
- No configuration update commands issued from the host are allowed during the MORE process.
- Write-back is disabled during the process but resumes at the end of the operation.
- The following parameters do not change as a result of a MORE operation: existing system drive write policy, LUN affinity/LUN mapping, cache size and all other controller parameters.
- Removing and replacing a controller during migration results in data loss. A hardware error which renders the controller unusable or the NVRAM unreadable also results in data loss during migration.
- Add capacity will not start if the target RAID set is in an online critical state.
- The system drive to be expanded must be online optimal.

➢ See also “Background Initialization” on page 11, “Conservative Cache Mode” on page 22, “Consistency Check” on page 23, “Controller Parameters” on page 69, “Critical” on page 113, “Initialize System Drive” on page 123, “I/O” on page 123, “Liquid Crystal Display (LCD)” on page 47, “Online” on page 131, “Online Critical” on page 131, “Physical Disk Drive” on page 135, “Programmable LUN Mapping” on page 52, “RAID Levels” on page 137, “Single Controller Mode” on page 146, “Spanning Drive Packs” on page 62, “Striping” on page 147, “System Drive Affinity” on page 64, “System Drive Format” on page 148, “System Drives” on page 149, and “Write-Back Cache” on page 151.

Online Spare Polling

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

The firmware polls every online spare disk drive periodically as a means of determining the health of that online spare disk drive. This feature is performed every *n* seconds, where n is configured via the firmware variables section.

➢ See also “Automatic Rebuild” on page 10, “Device States” on page 114, “Failed Disk Drive” on page 118, “Firmware Variables” on page 45, “Offline Failed” on page 131, “Online Optimal” on page 132, “Online Spare” on page 132, and “Rebuild” on page 144

Parallel Host Writes

SCOPE: This feature was implemented in firmware version 5.4 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

Allows multiple small writes to the same cache line to proceed in parallel resulting in improved performance. Firmware allows all non-overlapping writes to the same cache line to proceed immediately with no delay based on other host writes to the same cache line.

➢ See also “Cache Line Size” on page 18 and “Host / Server” on page 121.

Parity Check

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers. Parity Check is also referred to as consistency check.

The Parity Check function is used to verify the integrity of data on a system drive. It verifies that mirror or parity information matches the stored data on the redundant arrays (RAID 1, RAID 3, RAID 5, or RAID 0+1). If the parity block information is inconsistent with the data blocks, the controller has the ability to correct the inconsistencies.

A parity check continues through a controller failure (failover). The failed controller or replacement controller is held in reset until the parity check completes. If a relinquish controller command is executed, the parity check completes before relinquishing the controller (failback).

➢ See also “Consistency Check” on page 23, “Initialize System Drive” on page 123, “Mirroring” on page 129, “Mylex Online RAID Expansion (MORE)” on page 49, “Parity” on page 133, “RAID Levels” on page 137, “Redundant Array” on page 145, and “Striping” on page 147.

Predictive Failure Analysis™ (PFA)

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

Some IBM physical disk drives have the ability to send predictive failure analysis sense to the host. This information is an indication of an error or detected problem with the physical disk drive and the drive should be replaced. Mylex controllers recognize the PFA information and set the device state to offline failed.

➢ See also “Host / Server” on page 121, “Physical Disk Drive” on page 135, and “Self-Monitoring, Analysis and Reporting Technology (SMART)” on page 60.

Programmable LUN Mapping

SCOPE: This feature was implemented in firmware version 3.3 and is supported on the DACSX, DACSS, DACSF, DACFL, and DACFF (FF) controllers. This feature was incorporated into SANmapping™ with firmware version 6.0.

Programmable LUN mapping adds the ability to specify flexible LUN-to-system-drive mappings. This feature allows assignment of any LUN ID (even multiple LUN IDs) to any system drive on each port or to configure without specifying LUN-to-system-drive assignments using the default mapping algorithm. This feature permits the user to assign ownership of a system drive to any combination of controller/host ports.

System drives with all affinity are mapped to a LUN ID on every controller/host port. The LUN ID assignments for each controller/host port always start with LUN 0. The LUN ID mapping algorithm is illustrated in the following table.

LUN ID Mapping Algorithm

System Drive	Affinity Configuration (Controller:Host Port)	Controller 0: Host Port 0 LUN ID	Controller 0: Host Port 1 LUN ID	Controller 1: Host Port 0 LUN ID	Controller 1: Host Port 1 LUN ID
0	C0:P0	0			
1	All Affinity	1	0	0	0
2	C1:P1				1
3	C0:P0	2			
4	C0:P1		1		
5	C1:P0			1	
6	C0:P0	3			
7	All Affinity	4	2	2	2

Blank entries in the table indicate system drives which are inaccessible at that controller/host port target ID. Any LUN ID higher than the highest number in a column is an invalid LUN ID for that controller/host port/target ID.

In the event of a failover, the surviving controller assumes the partner’s SCSI target IDs on all available host ports. The LUN ID mapping for each target ID is not changed, but all of the target IDs are controlled by the surviving controller.

Programmable LUN Mapping and System Drive Affinity work together to define how the host accesses the available storage space. These features are incorporated and controlled through the SANmapping feature available with firmware version 6.0.

➢ See also “Controller/Host Port” on page 112, “Dual-Active Controller Configuration” on page 32, “Host Port” on page 121, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “SANmapping™” on page 56, “Simplex” on page 60, “System Drive Affinity” on page 64, “System Drives” on page 149, and “Target ID (TID)” on page 149.

Protected Boot ROM (PBR)

SCOPE: This feature was implemented in firmware version 7.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controller.

Protected Boot ROM (PBR) allows recovery of a controller if the firmware download process is interrupted.

The PBR image is stored in a section of the controller’s flash PROM. During controller initialization, the PBR image verifies that the firmware image is intact before transferring control to it. If the firmware image is found to be damaged or incomplete, the PBR image partially initializes the controller to allow the host to download a new firmware image.

➢ See also “Firmware Versions” on page 95 “Global Array Manager (GAM)” on page 45, and “SANArray Manager™ (SAM)” on page 56.

RAIDfx

SCOPE: RAIDfx is supported with firmware versions 3.2 through 6.3 and not supported with firmware versions 7.0 and greater.

RAIDfx is a Mylex software utility that provides configuration capabilities to the Mylex controllers. RAIDfx is used to configure and manage disk arrays connected to one or more Mylex controllers. The software provides the same functions as the LCD/VT100 menu-driven utility. RAIDfx has additional capabilities beyond those functions available through the LCD/VT100.

The RAIDfx utility requires a host PC connected to the array system. RAIDfx operates through a DOS ASPI system or Windows NT operating system, both via SCSI connections to the controller.

Refer to the *RAIDfx Manager User Guide* for complete installation and user instructions.

➢ See also “Configuration Tools and Utilities” on page 21, “Firmware Versions” on page 95, “Global Array Manager (GAM)” on page 45, “Liquid Crystal Display (LCD)” on page 47, “SANArray Manager™ (SAM)” on page 56, and “VT100” on page 65.

Replacement Controller Behavior with Forced Simplex Enabled

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

Often Mylex controllers are shipped from the factory with the *Forced Simplex* controller parameter enabled rather than releasing simplex-only firmware (type 0, 1, 2, or 3). This parameter produces the following behaviors when upgrading to a dual-active controller system or when replacing a failed controller.

Forced Simplex Behavior During Controller Replacement

Existing Controller	Replacement or Additional Controller	Replacement Method	Result
Dual-Active	Forced Simplex	Hot (Power On)	Forced simplex changes to dual-active and the configuration is updated.
Dual-Active	Forced Simplex	Cold (Power Off)	Forced simplex changes to dual-active and the configuration is updated. The controllers reset once after power is supplied.
Dual-Active	Dual-Active	Either Hot or Cold	There is no change to the configuration.
Forced Simplex	Upgrade to Dual-Active with Forced Simplex Controller	Hot (Power On)	The additional controller is held in reset until the controllers restart. Both controllers operate in dual-active mode and the Force Simplex controller parameter is automatically cleared.
Forced Simplex	Upgrade to Dual-Active with Forced Simplex Controller	Cold (Power Off)	Both controllers start dual-active.

➢ See also “Cold Swap” on page 112, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Force Simplex” on page 78, “Hot Plugging” on page 121, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “Replacement Controller” on page 145, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, “Reset Controllers” on page 54, “SANmapping™” on page 56, “Surviving Controller” on page 148, “System Drive Affinity” on page 64, and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

Reset Controllers

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers. This feature has evolved through the various firmware versions and releases. The number of features, operations, and parameter changes that require a controller reset vary with the firmware version.

Firmware version 4.2 allows some controller parameter changes to take effect immediately, a controller reset is not required to change the parameter. Refer to “Controller Parameters” on page 69 for detailed information about each controller parameter.

Firmware version 7.2 allows changes to the configuration to take effect immediately, a controller reset is not required when adding or configuring system drives or logical devices.

This operation performs a *warm* power cycle and reinitialization of the controllers. A controller reset is required after changing some controller parameters and is recommended after making any changes to the configuration.

After Configuration Changes

When deleting system drives or an entire configuration, each system drive is quiesced before being deleted. This feature ensures no stale data remains in the cache after the system drive or entire configuration is deleted and a new system drive or configuration is created. A controller reset is no longer required when deleting or creating new system drives or configurations.The quiescence process is performed on a system-drive basis including the steps listed below.

1. Block all new host I/O by returning a busy status.
2. Wait for the active I/O to complete.
3. Flush and validate the cache.
4. Invalidate the mirrored cache on the partner controller.
5. Delete the system drive.

➢ See also “Cache” on page 111, “Cache Flush” on page 112, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “I/O” on page 123, “Partner Controller” on page 133, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, and “Replacing a Failed Controller in Simplex Systems” on page 99.

SAF-TE

SCOPE: This feature was implemented in firmware version 3.2 and is supported on the DACSX, DACSS, DACSF, and DACFL controllers.

A SCSI Accessed Fault-Tolerant Enclosure (SAF-TE) provides monitoring of disk drives, power supplies, fans, and temperature in the enclosure.

During controller initialization, each SCSI bus is scanned for SAF-TE devices. If a SAF-TE device is detected, the controller polls the SAF-TE device every ten seconds for the following conditions:

- Disk drive insertion or removal
- Power supply malfunctioning
- Fan malfunctioning
- Enclosure temperature out of range

The controller also updates the SAF-TE device with array disk drive status every ten seconds.

Host software can access the SAF-TE device directly by issuing Mylex vendor-unique pass-through direct commands. The software can pass the following CDBs to the SAF-TE device:

- Test Unit Ready
- Request Sense
- Inquiry
- Read Buffer
- Write Buffer

Fan and power supply failures are reported to the configuration tools and utilities as failures on fan or power supply units 0 through 5. If the SAF-TE device reports a fan or power supply failure on a unit number greater than 5, it is reported as a failure on unit 0.

➢ See also “AEMI” on page 6, “Bus” on page 111, “Conservative Cache Mode” on page 22, “Drive State Management” on page 31, “Environmental Device” on page 117, “Fault Management” on page 41, “Host / Server” on page 121, and “SES” on page 58.

SANArray Manager™ (SAM)

SCOPE: This version of software is supported with firmware versions 7.4 and greater on the FFx, FF2, and FFx2 controllers.

SANArray Manager (SAM) and Global Array Manager (GAM) are components of a Mylex software utility that provides configuration capabilities to the Mylex controllers. SAM/GAM are a GUI supported interface that allows the user to monitor, manage, maintain and configure Mylex controllers and the physical devices connected to the controllers. SAM/GAM has two components—the GAM Server and the SAM Client. The Server functions on Windows NT, Netware, and several UNIX operating systems. The Client functions on Windows NT operating systems.

Refer to the *SANArray Manager Client Software for Mylex External Disk Array Controllers Installation Guide and User Manual* and the *Global Array Manager Server Installation Guide and User Manual* for complete installation and user instructions.

GAM Server

The Server collects and disseminates information on disk array or host bus adapter (HBA) subsystem status and resource utilization. The Server executes the management instructions specified by the Client. The Server functions on Windows NT, Windows 2000, Netware, Unixware, Linux, SCO, and Solaris x86 operating systems.

SAM Client

Configuration functions are easily performed using RAID Assist™, a wizard-like utility in the Client that simplifies the process of setting up or reconfiguring a disk array. The Client organizes the information collected by the server through a graphical display. Errors and events are recorded in a log, and if a problem is serious enough to warrant immediate attention, operators can be alerted via pop-up windows, pagers, fax, or email. The Client also manages or performs maintenance on individual disk arrays and disk drives. This maintenance can include removing physical devices from operation in a functioning disk array, rebuilding disk drives, selecting spare disk drives, and initiating a consistency check on redundant array systems. The Client functions on Windows 95, Windows 98, Windows NT, and Windows 2000 operating systems.

➢ See also “Configuration on Disk (COD)” on page 18, “Configuration Tools and Utilities” on page 21, “Firmware Versions” on page 95, “Global Array Manager (GAM)” on page 45, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, “Upgrading a Configuration Created with RAIDfx” on page 101, and “VT100” on page 65.

SANmapping™

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

The SANmapping feature restricts host access to configured system drives similarly to the Programmable LUN Mapping feature. Programmable LUN Mapping is incorporated into the SANmapping feature. SANmapping is intended for use in configurations in which multiple host computers attach to one or more Mylex controllers or a Storage Area Network (SAN) configuration. The host computers are attached to the controller(s) through a Fibre Channel arbitrated loop or through a switch. An example of Fibre Channel arbitrated loop configuration is shown in the following illustration.

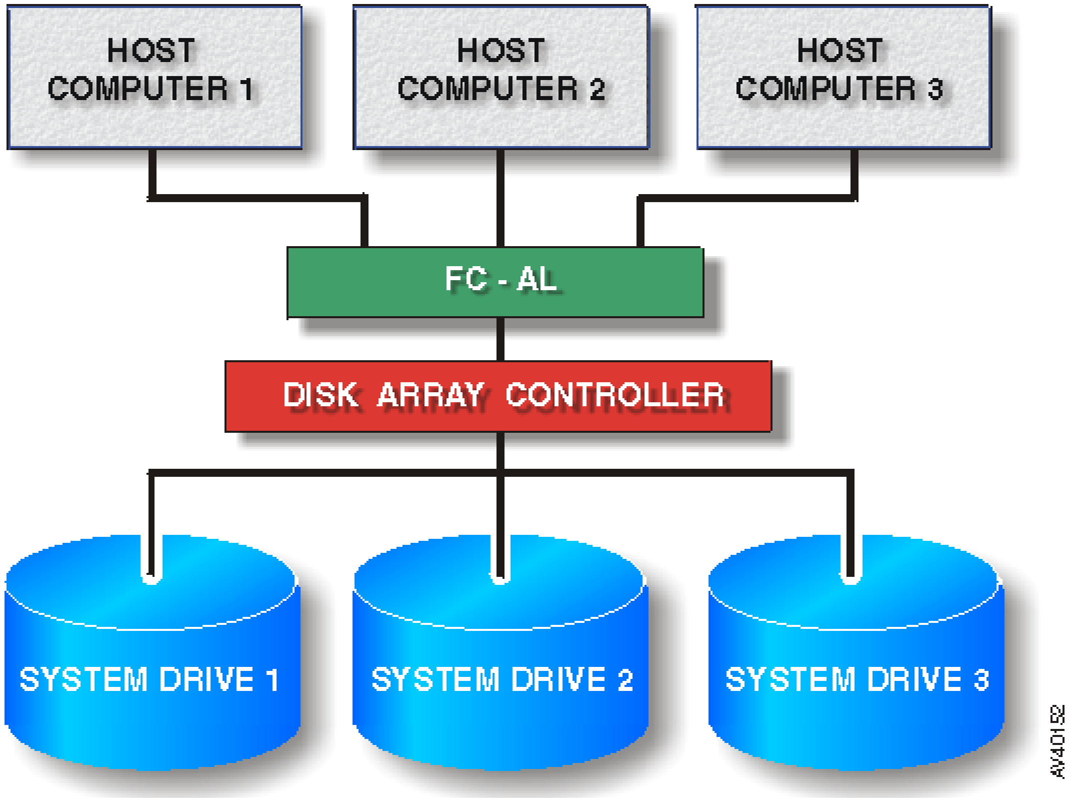

Fibre Channel Arbitrated Loop System Configuration

In this example, the storage available to the controller has been configured into three separate system drives. Without SANmapping, each host computer (1 through 3) has complete access to all three system drives. When the host systems start, operating systems such as Windows NT attempt to access all of these system drives. As a result, any data on the array can be accessed by any of the Windows NT systems. In some cases, Windows NT will automatically write an identifying *signature* to these storage areas. This results in a data corruption if another host system has stored data there.

When a new configuration is created, the system drive affinity defaults to all controller/host ports and all hosts. By utilizing SANmapping, each system drive's affinity can be reconfigured to limit accessibility to a certain controller/host port or single host computer and to change LUN mapping assignments. Using the example in the above illustration, the following table shows how SANmapping is used to limit the system drive access to only one host computer on one controller/host port. The simplest configuration would allow Host Computer 1 access to System Drive 1 only, Host Computer 2 access to System Drive 2 only, and Host Computer 3 access to System Drive 3 only.

SANmapping Configuration Data

	System Drive 1	**System Drive 2**	**System Drive 3**
Affinity to C0P0	Y	Y	Y
LUN ID Number	0	1	2
Visibility	Y	Y	Y
Host 1 Enabled	Y	N	N
Host 2 Enabled	N	Y	N
Host 3 Enabled	N	N	Y

The controller uses a table of WWNs to determine access to a specific system drive. If a host sends a new command to the controller, the controller validates the WWN, LUN, and controller/host port prior to servicing the command. If the WWN, LUN, and port information are valid for the system drive, the requested command completes normally.

There are three exceptions to the response to commands when the WWN, LUN, and port combination are not valid:

- If the request is an INQUIRY command, the controller returns the Inquiry data with the peripheral qualifier set to indicate that the target is capable of supporting the specified device type on this LUN, but no device is currently connected to that LUN.
- If the request is a REPORT LUNS command, and the addressed LUN is 0, the controller completes the command normally, reporting only the LUNs accessible by the host requesting the command.
- If the request is a Mylex direct command or pass-through command, the command is processed normally by the controller. This allows a controller that is not configured to be reconfigured to operate correctly with the attached hosts.

➢ See also “Array” on page 111, “Configuration Strategies” on page 94, “Configuration Tools and Utilities” on page 21, “Controller/Host Port” on page 112, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, “Host Channel” on page 121, “Host / Server” on page 121, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “Storage Area Network (SAN)” on page 147, “System Drives” on page 149, “System Drive Affinity” on page 64, “World Wide Name (WWN)” on page 151, and “World Wide Name (WWN) Table” on page 66.

SES

SCOPE: This feature was implemented in firmware version 5.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers. Enhancements to SES features have been added in firmware version 6.1, 6.2, 7.0, and 7.4. The SES Enhancements, N-Way SES, and Adding SES Enclosures features (added in firmware 7.4) are supported on the FFx, FF2, and FFx2 controllers.

SCSI-3 Enclosure Services (SES) provides support for disk drives, power supply, temperature, door lock, alarms, UPS, and enclosure services controller electronics. The SES process polls each SES device once every 10 seconds.

SES can be implemented by two methods: access can be through an enclosure services device or through a non-enclosure services device. When access is through an enclosure services device, an SES device is present on the fibre loop and SCSI commands are sent directly to the device. When access is through a non-enclosure services device, SCSI commands are sent to a direct access storage device and passed through to the SES device. Mylex firmware supports access through both methods.

During controller initialization each device that is attached to each fibre loop is interrogated and the inquiry data is stored. After the interrogation and if SES devices are detected, the SES process is started.

The SES process polls every ten seconds updating:

- Disk drive insertion status
- Power supply status
- Cooling element status
- Temperature status
- UPS status

If a UPS is connected to any of the enclosures monitored by SES and a UPS reports ac failure or two minute warning, conservative cache is enabled and all system drives are switched to write-through cache.

Host software can access the SES device to retrieve enclosure information.

Failures on fans or power supplies numbered greater than 5 are reported to the configuration tools and utilities as failures on unit 0. There is only one over temperature indicator.

Enhancements

The following SES enhancements have been added.

- SES Firmware Download—The controller firmware supports pass-through commands that are used to download new SES firmware to the SES processor. This allows a host utility the ability to communicate SCSI commands directly to a device that is visible only to the controller.
- Enclosure Controller Page
    - Power Supply Element—The enclosure services data for reporting power supply element information is expanded to present information for ac fail, dc fail, temperature warning and over temperature failure.
    - Temperature Sensor Element—The enclosure services data for reporting temperature sensor element information is expanded to present information for over temperature (OT) failure, under temperature (UT) failure, and UT warning.
    - Set Temperature Element
- Prepare Disk Drive For Removal—This command is used to prepare a disk drive for removal. Upon receipt of this command from a host utility application, the controller will command a disk drive to spin down and then return status to the host after completion. The command is only accepted for failed disk drives and disk drives not part of a system drive configuration. This command is used to remove a disk drive from an enclosure. An SES enclosure is required for this command to work. The enclosure’s capabilities are used to spin down the disk drive and mark it as removed.
- Disk Drive Rebuilding—The controller provides additional information to the SES processor to indicate that a disk drive is being rebuilt. The rebuild/remap bit is used to indicate that a disk drive is in the process of being rebuilt.
- Predicted Drive Failure—The controller provides additional information to the SES processor to indicate that a disk drive was killed because of a predicted fault indication received from the disk drive.

N-Way SES

This features allows for more than two SES devices to be used in an enclosure. The firmware only keeps track of a primary and secondary SES device for management purposes. If a disk drive to which an SES device is mapped is removed, the firmware scans all other drives in the enclosure attempting to find a replacement SES device. In previous firmware versions, when an SES device was removed, it could only be replaced by a new SES device in the same enclosure slot.

➢ See also “Adding SES Enclosures” on page 91, “AEMI” on page 6, “Conservative Cache Mode” on page 22, “Conservative Cache Mode Parameter” on page 72, “Drive State Management” on page 31, “Environmental Device” on page 117, “Fault Management” on page 41, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, “SAF-TE” on page 55, “SCSI” on page 145, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Self-Monitoring, Analysis and Reporting Technology (SMART)

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

The Self-Monitoring, Analysis and Reporting Technology (SMART) standard provides physical disk drives with the ability to predict or anticipate the failure of a disk drive and send sense information to the host. This allows the user to back up data prior to the drive's failure. Mylex controllers recognize the information and set the device state to offline failed.

➢ See also "Failed Disk Drive" on page 118, "Host / Server" on page 121, "Physical Disk Drive" on page 135, and "Predictive Failure Analysis™ (PFA)" on page 52.

Simplex

SCOPE: The simplex configuration is a legacy configuration supported by all Mylex controllers. Performing a MORE (Mylex Online RAID Expansion) operation requires a dual-active controller system be in simplex (failed over) mode.

A disk array system having only one controller or capable of having dual-active controllers but only one controller is present or active. A simplex system does not provide fault tolerance in the event of a controller failure. The single controller handles all I/O requests. Simplex controller firmware type is 0, 1, 2 or 3. The firmware type is displayed in the firmware header.

NOTE: The FFx and FFx2 controllers do not support firmware types 0,1,2, or 3. Single controller (Simplex) support is configured using the *Force Simplex* controller parameter.

A single controller connects to a single host bus adapter (HBA) and all the disk drives in the array. Disk drives are fault tolerant when a RAID level providing redundancy is used. In the simplex configuration, however, if the controller or HBA should fail, the data will not be accessible until the failure is corrected.

Simplex controllers supporting two host channels can be utilized in a dual-host configuration using either a single server or two servers. In a single server system, the two host channels may be connected to two HBAs within a single server. A single HBA failure can occur without losing access to the disk array. This requires alternate path software to direct all I/O to the surviving HBA. A two server system physically isolates the two HBAs while allowing the systems to share the array by connecting to different host channels on the controller. This provides fault tolerance in case one server in the system fails. It also ensures compatibility between systems that are not designed to allow two hosts to share the SCSI bus.

ATTENTION: **If two systems independently access the same volume of data, and the operating system does not support file locking, data corruption may occur. To avoid this, create two or more volumes (or LUNs) and configure each volume to be accessed by one system only.**

➢ See also "Active Controller" on page 109, "Alternate Path Software" on page 7, "Array" on page 111, "Configuration Strategies" on page 94, "Controller Parameters" on page 69, "Dual-Active Controller Configuration" on page 32, "Failed Controller" on page 118, "Fault Tolerance" on page 118, "Firmware Header Information" on page 44, "Force Simplex" on page 78, "Mylex Online RAID Expansion (MORE)" on page 49, "Replacing a Failed

Controller in Simplex Systems" on page 99, "Simplex Fibre Channel System" on page 146, "Single Controller Mode" on page 146, and "Upgrading From Simplex to Dual-Active Configurations" on page 103.

Single Point of Failure (SPOF)

SCOPE: This feature was implemented in firmware versions 5.4 and is supported on the DACFL controller.

This feature benefits systems with drive channels routed from the primary enclosure to an external expansion enclosure. The expansion enclosure must have a unique path to each controller and active termination implemented on each drive channel that allows access through an alternate path.

SPOF is designed to accommodate failure of one of the SCSI expansion cables. This functionality causes a controller to failover when the external SCSI cable is removed or damaged. A SCSI cable failure results in drive failure or failure of the entire drive channel and sends a disk error to the controller. When a disk error occurs, the detecting or reporting controller queries the partner controller to access that drive. If the partner controller can successfully issue an inquiry command to the drive, the partner controller holds the reporting controller in reset (issues a kill partner). Failing the reporting controller allows access to the disk drives to failover to the partner controller for access by the host system. Recovery from this condition requires isolating the problem, for example broken cable, loose cable, or faulty controller, then replacing the failed component.

If the partner controller cannot access the drive, it sends a message to the reporting controller that the drive cannot be accessed. The reporting controller then processes a normal drive failure. This results in failing the disk drive. Recovery from this condition requires removing, replacing, and rebuilding the disk drive.

If the reporting controller is *not* operating in a dual-active configuration, the disk drives and the entire drive channel are marked as failed. Access to the external enclosure and associated disk drives is lost. Recovery from this condition requires investigating all components for failures, including cables, controller, enclosure, and disk drives.

➢ See also "Alternate Path" on page 110, "Drive Channel" on page 27, "Drive Channel Failover and Failback" on page 28, "Dual-Active Controller Configuration" on page 32, "Failed Disk Drive" on page 118, "Failover and Failback" on page 35, "Host / Server" on page 121, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, "Replacing a Failed Controller in Simplex Systems" on page 99, "SCSI Cabling" on page 99, and "SCSI Termination" on page 100

Soft Addressing Detection

SCOPE: This feature was implemented in firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

Soft Addressing occurs during Fibre Channel loop initialization (LIP) when a conflict in device loop IDs is detected; for example, two or more drives have the same physical or hard loop ID, and the Fibre Channel loop protocol chip assigns unique loop IDs for the drives in conflict.

➢ See also "Hard Loop ID" on page 120, and "Loop Initialization Primitive (LIP)" on page 126.

Spanning Drive Packs

SCOPE: This feature was implemented in firmware versions prior to 3.2 and is supported on all controllers. The feature was modified with firmware version 7.0 to allow 16 disk drives per pack.

Spanning drive packs provides the ability to configure multiple or parts of multiple drive packs as one system drive. This allows more than 8 or 16 (firmware version 7.0 or greater) physical disk drives to be configured as a system drive. A system drive can span up to 4 or 16 (firmware version 7.0 or greater) drive packs. The controller creates the spanned system drive during the array configuration process. Data is striped across the drive packs of the spanned system drive.

An example configuration using firmware versions earlier than 7.0 and the maximum number of system drives, maximum number of spanned packs, and maximum number of physical disk drives per pack would require 256 physical disk drives (8 system drives x 4 spanned packs/system drive x 8 physical disk drives/pack).

An example configuration using firmware versions 7.0 and greater and the maximum number of system drives, maximum number of spanned packs, and maximum number of physical disk drives per pack would require 8192 physical disk drives (32 system drives x 16 spanned packs/system drive x 16 physical disk drives/pack). Mylex controllers are currently limited by the maximum number of disk drives supported per drive channel.

If using firmware versions earlier than 7.0 to create spanned system drives, take into account the following considerations and limitations.

- RAIDfx requires that all spans be derived from packs that have the same number of disk drives and capacity.
- RAIDfx prevents the spanning feature when drive packs are arranged so that a drive pack with a different number of physical disk drives is defined between two packs of equal numbers of physical disk drives.
- RAIDfx automatically spans disk packs. To prevent automatic spanning, completely finish configuring a drive pack by defining its RAID level and system drives, then initializing the system drives. Continue creating other drive packs that contain the same number of physical disk drives.
- Spanned system drives cannot be created using the LCD keypad, VT100 interface, or embedded configuration utility.
- SAM/GAM does allow spanning, but requires modification to an .ini file and prevents the use of other options available to the user.

Spanning across several drive packs is shown in the following illustration.

NOTE: Spanning and MORE operations are mutually exclusive. A request to perform a MORE operation will be rejected if any system drive is spanned across multiple drive packs.

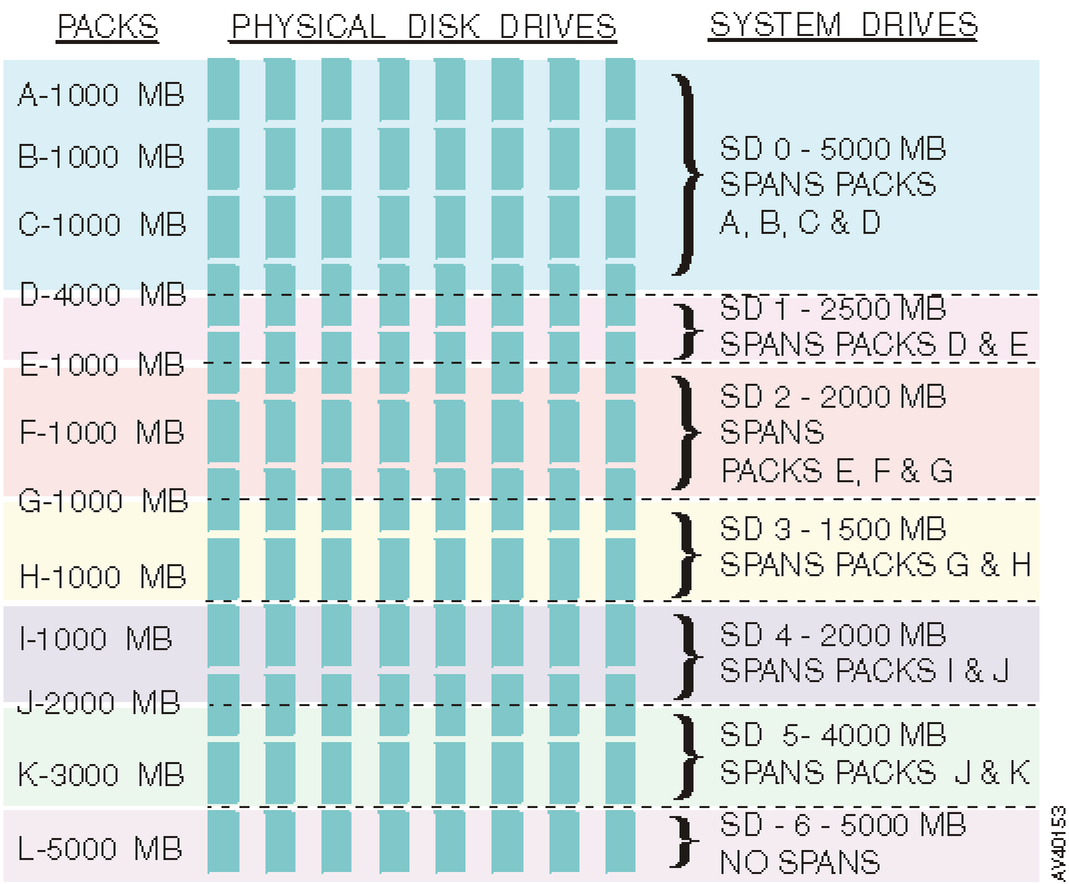

System Drive Spanning

➢ See also “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Packs / Groups” on page 29, “Embedded Configuration Utility” on page 35, “Global Array Manager (GAM)” on page 45, “Initialize System Drive” on page 123, “Liquid Crystal Display (LCD)” on page 47, “Mylex Online RAID Expansion (MORE)” on page 49, “Physical Disk Drive” on page 135, “RAID Levels” on page 137, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, “System Drives” on page 149, and “VT100” on page 65.

Standard Data Caching

SCOPE: This feature was implemented in firmware version 3.3 and is supported on all controllers.

Read cache is always enabled.

The controller supports either write-through or write-back data caching. Write caching is set independently for each system drive by the configuration utility. With write-through data caching enabled, write data is transferred to the disk drives before completion status is issued to the host initiator. With write-back caching enabled, a write completion status is issued to the host initiator after the data is stored in the controller’s cache, but before the data is transferred to the disk drives. In dual-active controller configurations with write-back caching enabled, the write data is always copied to the cache of the partner controller before completion status is issued to the host initiator.

➢ See also “Active Controller” on page 109, “Cache” on page 111, “Cache Flush” on page 112, “Configuration Tools and Utilities” on page 21, “Conservative Cache Mode” on page 22, “Dual-Active Controller Configuration” on page 32, “Host / Server” on page 121, “Partner Controller” on page 133, “System Drives” on page 149, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Stripe Size

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers. As of firmware version 7.0, stripe size is set at the time a configuration is created and is no longer a controller parameter.

Defines the amount of data, in kilobytes (1024 bytes), that a single data cache line can hold in cache memory. It also defines the amount of data written to a single disk drive. For any given RAID striping algorithm, the stripe size is the base unit size written to every disk drive in the RAID group. Currently supported stripe sizes for Mylex external RAID controllers are 8 KB, 16 KB, 32 KB, and 64 KB.

➢ See also “Cache” on page 111, “Cache Line Size” on page 18, and “Stripe Size Parameter” on page 86.

System Drive Affinity

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on the DACSX, DACSS, DACSF, DACFL, and DACFF (FF) controllers. This feature was incorporated into SANmapping™ with firmware version 6.0.

System Drive Affinity assigns system drive access to a single host port in the system. A system drive is assigned affinity to a particular controller/host port combination or it is assigned *all affinity*, in which case it is accessible via all controller/host port combinations.

Dual-active controller configurations may use System Drive Affinity to define affinity of each system drive to one, two, three or four of the possible controller/host ports. Simplex configurations with two host ports may use the System Drive Affinity feature to define affinity of each system drive to one or both host ports. System drives which are *not owned* by a controller/host port are not accessible.

System Drive Affinity and Programmable LUN Mapping work together to define how the host accesses the available storage space. These features are incorporated and controlled through the SANmapping feature available with firmware version 6.0.

➢ See also “Controller/Host Port” on page 112, “Dual-Active Controller Configuration” on page 32, “Host Port” on page 121, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, “Simplex” on page 60, “System Drives” on page 149, and “Target ID (TID)” on page 149.

UNIX Support and Firmware Types

SCOPE: This feature was implemented in firmware prior to version 3.2 and is supported on all controllers.

If the firmware type indicates the UNIX support bit is set, types 2, 3, 6, and 7, the firmware will return good status for a TEST UNIT READY command issued to LUN 0 when no system drives exist or if all system drives are offline. The firmware will also return a false 100 MB capacity for a READ CAPACITY command issued to LUN 0. The UNIX bit has no effect on an INQUIRY command. The UNIX bit is checked only when system drives mapped to LUN 0 are offline or there are no system drives mapped to LUN 0.

➢ See also “Firmware Header Information” on page 44, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, “System Drive Affinity” on page 64, and “System Drives” on page 149.

VT100

SCOPE: The VT100 interface has limited support with firmware versions prior to 7.0, but has been modified to support the features available with firmware versions 7.0 and greater.

Terminal-emulation interface that allows the user to access the controller through a PC without an additional utility. Refer to the *Terminal-Emulation (VT100) User Guide and LCD Front Panel Status Guide* for more information.

➢ See also “Configuration Tools and Utilities” on page 21, “Embedded Configuration Utility” on page 35, “Global Array Manager (GAM)” on page 45, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, and “SANArray Manager™ (SAM)” on page 56.

World Wide Name (WWN) Assignments

SCOPE: Unique WWN support was implemented in firmware version 7.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

This feature was modified in firmware version 7.4 to include common WWN support and is supported on the FFx, FF2, and FFx2 controllers.

This feature was modified in firmware version 7.7 to include enclosure WWN support and is supported on the FFx, FF2, and FFx2 controllers.

Mylex controllers configured in a dual-active controller system have the ability to present WWN information using a variety of methods. The method used is determined with a firmware variable. The default method uses unique WWNs for both controllers. A more robust method forces both controllers to assume the WWN information of the enclosure in which the controllers are installed. This method can be used only with enclosures that support this feature. If the firmware variable is set for enclosure WWN and the enclosure does not support this feature, common WWNs are used.

Unique World Wide Names

This method of assigning WWNs presents each controller or device to the host system with a unique WWN. Dual-active controllers have unique WWNs derived from the MAC addresses programmed into the controllers in the factory, for example, controller A and B.

NOTE: If either controller fails and is replaced, the host system must be restarted before recognizing the new controller. The host must be restarted whether the controller is replaced with power on (hot replace) or off (cold replace).

Common World Wide Names

This method of presenting WWN information allows dual-active controllers to share a common WWN derived from the MAC address of the primary controller, for example, controller A. If the primary controller fails and is hot replaced, the replacement controller assumes the WWN of the controller it replaced. The replacement controller uses this assumed WWN until it is restarted. After the replacement controller is restarted, a new common WWN is derived from the replacement controller's MAC address and used by both controllers, for example, controller D. If the primary controller is cold replaced, the new common WWN is used by both controllers.

NOTE: The host must be restarted before recognizing the new WWN whether the controller is hot replaced or cold replaced.

The controllers log into the Fibre Channel loop using the same node name but each controller has a unique port name, for example, Port 0 and Port 1.

Enclosure World Wide Name

This method derives the WWN from the enclosure. Both controllers in a dual-active controller pair report the same WWN as the enclosure. If the primary controller fails and is replaced, the replacement controller assumes the enclosure WWN. The replacement controller assumes the enclosure WWN whether the replacement occurs with power on or off. Additionally, restarting the host is not required in order to recognize the new controller.

➢ See also "Dual-Active Controller Configuration" on page 32, "Fibre Channel" on page 43, "Firmware Variables" on page 45, "Host / Server" on page 121, "Hot Plugging" on page 121, "Node Name" on page 130, "Port Name" on page 135, "Primary Controller" on page 136, "Replacement Controller" on page 145, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, "World Wide Name (WWN)" on page 151.

World Wide Name (WWN) Table

SCOPE: This feature was implemented in firmware version 6.0 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

This feature was modified with firmware version 7.7 to include a delete WWN from WWN table option.

This feature was modified with firmware version 8.0 to increase the number of supported host WWNs from 64 to 256.

ATTENTION: **All SANmapping assignments created using firmware version 8.0 are lost if the firmware is downgraded to a previous firmware version.**

A controller currently maintains a host WWN table until the configuration is cleared. Entries in the WWN table corresponding to hosts no longer logged into the controller remain consumed and unusable, potentially limiting the ability to provide storage access to newly added hosts. The WWN table was modified with firmware version 7.7 to enable the user to remove or delete unused host WWNs. The user can determine which host WWNs are obsolete and need to be removed from the WWN table.

As WWNs are removed and the WWN table updated, any WWNs following those deleted are moved up to fill the vacancies in the WWN table. The SAN map uses the indices of the WWN table entries to specify hosts that have access to a specific system drive. When the WWN table entries are deleted, the indices for any entries following those WWNs that were deleted change.

➢ See also “Controller Specifications” on page 1, “Configuration on Disk (COD)” on page 18, “SANmapping™” on page 56, “World Wide Name (WWN) Assignments” on page 65.

CONTROLLER PARAMETERS

This section describes the controller parameters. Each Mylex controller is shipped from the factory with initial settings that have been found to work well in a majority of applications and environments. These settings are listed as the controller parameter settings and vary depending on product and user requirements. Some parameters are product or configuration dependent and do not have a recommended default setting provided. User requirements are not always the same as the suggested default settings, so you may want to modify certain settings. Additionally, if you are going from a simplex configuration to a dual-active controller configuration, certain controller parameters must be changed to accommodate the new dual-active controller configuration. The following list of controller parameters is detailed in the following paragraphs.

A thorough understanding of the parameters and settings is strongly recommended before modifying the current settings and creating a configuration.

NOTE: Not all parameters are supported on every controller or through every configuration utility. Refer to the specific controller parameter for detailed information about limitations.

Automatic Failback or Autorestore	Automatic Restart (Reboot) on Failure Parameter	Automatic Rebuild Management
Block Size	Blocking Factor	Cache Line Size Parameter
Coalescing	Command Tagging	Conservative Cache Mode Parameter
Controller Name	Controller Present/Fault Select	Controller Present/Fault Signals
Data Bus Width	Device Combing	Disable BUSY Status During Failback
Disable Check Condition for Invalid LUN	Disable Queue Full Status	Disable Wide Operation
Disconnect on First Command	Disk Startup Delay 1	Disk Startup Delay 2
Disk Startup Mode	Disk Startup Number of Devices	Disk Write Through Verify
Disks Per Spin	Duplex Fault Signals	Duplex Fault Signals on Channel 4
Elevator	Enable Background Initialization	Failover Topologies
Force Simplex	Frame Size	Generate Debug Dump
Hard Loop IDs	Host (SCSI) Bus Reset Delay	Host Transfer Width
Maximum IOPs	Mode Select SRA Enable	No Pause on Controller Not Ready
Node Name Retention	On Queue Full Give Busy	Operational Fault Management
Override Multiport Reset	PCI Latency Control	Queue Limit
RAID5 Algorithm	Read Ahead	Read/Write Control
Reassign Restricted to One Block	Rebuild and Check Consistency Rate	Reset Propagation

SAF-TE Data for UPS Support	SCSI Transfer Rate or Mega-transfers/sec	Serial Control
Serial Port Baud Rate	Serial Port Debug Type	Serial Port Usage
Simplex - No Reset	Smart Large Host Transfers	Spin-up Disk Delay
Spin-up Option	Spin-up Sequence Delay	Stripe Size Parameter
Super Read Ahead	Transfer Speed	Transfer Width
True Verify	Use of UPS	Vendor Unique Direct Command
Vendor Unique Pass-Through Command	Vendor Unique TUR	Write Through Verify

➢ See also “Automatic Rebuild” on page 10, “Configuration Strategies” on page 94, “Configuration Tools and Utilities” on page 21, “Conservative Cache Mode” on page 22, “Dual-Active Controller Configuration” on page 32, “Embedded Configuration Utility” on page 35, “Fibre Channel” on page 43, “Global Array Manager (GAM)” on page 45, “Hard Loop ID” on page 120, “Liquid Crystal Display (LCD)” on page 47, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, “Simplex” on page 60, “Stripe Size” on page 64, “Upgrading From Simplex to Dual-Active Configurations” on page 103, and “VT100” on page 65.

Automatic Failback or Autorestore

Default=Disabled. When enabled in a dual-active controller system, Automatic Failback or Autorestore allows automatic recovery of a partner controller when a replacement is inserted. If you enable this option, you must also enable the *Controller Present/Fault Signals* or the *Duplex Fault Signals* parameter. A controller reset is required before this parameter takes effect.

➢ See also “Controller Present/Fault Signals” on page 73, “Duplex Fault Signals” on page 77, “Dual-Active Controller Configuration” on page 32, “Partner Controller” on page 133, “Replacement Controller” on page 145, “Replacement Controller Behavior with Forced Simplex Enabled” on page 53, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, and “Reset Controllers” on page 54.

Automatic Restart (Reboot) on Failure Parameter

SCOPE: This controller parameter was implemented with firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This parameter controls the behavior of the Automatic Restart on Failure feature. This parameter sets the following values:

- Maximum number of times a controller attempts automatic restart to recover from firmware detected errors and
- Minimum time interval that the controller must operate before refreshing the number of restart attempts.

When a controller reaches the maximum restart attempts, the automatic restart feature becomes disabled until the value is refreshed. Any subsequent firmware detected errors require manual intervention to recover the controller. Possible values for the maximum automatic restart attempts are 0 to 15 attempts.

Possible values for the time interval between refreshing the number of restart attempts are 0 to 15, based on the values in following table.

Required Operational Time Intervals

Possible Value	Time Interval
0	Infinity (never reset automatically)
1	3 minutes
2	5 minutes
3	15 minutes
4	30 minutes
5	60 minutes
6	90 minutes
7	2 hours
8	4 hours
9	8 hours
10	12 hours
11	16 hours
12	24 hours
13	2 days
14	4 days
15	7 days

When the automatic restart feature is disabled, manual intervention is required to recover from fatal firmware detected errors. Manual intervention may involve physically removing and replacing the failed controller. This parameter takes effect immediately, without resetting the controllers.

➢ See also “Replacing a Failed Controller in Existing Duplex Systems” on page 97 and “Replacing a Failed Controller in Simplex Systems” on page 99.

Automatic Rebuild Management

Default=Enabled. The Automatic Rebuild Management function allows the controller to take autonomous actions when a disk drive fails and a configured online spare disk drive is present.

The Automatic Rebuild Management function works in conjunction with *Operational Fault Management* and features in AEMI, SAF-TE and SES certified disk array enclosures to detect the removal of a failed disk drive. The Automatic Rebuild Management function also performs an automatic rebuild after a failed disk drive is removed and a replacement disk drive is installed into a redundant (fault tolerant) array (RAID 1, RAID 3, RAID 5, and RAID 0+1).

Without this enabled, a host must issue the rebuild command. A controller reset is required before this parameter takes effect.

➢ See also “AEMI” on page 6, “Automatic Rebuild” on page 10, “Failed Disk Drive” on page 118, “Fault Management” on page 41, “Fault Tolerance” on page 118, “Operational Fault Management” on page 82, “RAID Levels” on page 137, “Redundant Array” on page 145, “Reset Controllers” on page 54, “SAF-TE” on page 55, and “SES” on page 58.

Block Size

SCOPE: The Block Size controller parameter is a legacy option and cannot be changed.

Default=512 (not changeable). The Block Size parameter indicates that the logical block sizes of the system drives are 512 bytes.

Blocking Factor

ATTENTION: **Changing this value after data is written to the disk drives results in data loss.**

SCOPE: The Blocking Factor controller parameter was implemented with firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This value multiplied by 8 KB is the stripe size used by the controller to allocate cache lines in memory and data striping on the disk. Possible values are 1, 2, 4, and 8; which result in stripe size values of 8 KB 16 KB, 32 KB, and 64 KB, respectively.

This value is set only at the time a new configuration is created.

Cache Line Size Parameter

SCOPE: The Cache Line Size controller parameter is a legacy command available only through the SAM/GAM configuration utility.

Default=Stripe Size. The Cache Line Size parameter is set in conjunction with Stripe Size and represents the size of the data "chunk" that will be read or written at one time. The cache line size must be set to the same value as the stripe size.

➢ See also "Cache Line Size" on page 18, "Global Array Manager (GAM)" on page 45, "SANArray Manager™ (SAM)" on page 56, "Stripe Size Parameter" on page 86, and "Stripe Size" on page 64.

Coalescing

SCOPE: This parameter is referred to as Device Combing in firmware versions 7.0 or greater. Refer to "Device Combing" on page 73 for more information. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Command Tagging

SCOPE: This command is not supported with firmware versions 4.2 or greater.

Default=Enabled for all drive channels. Command Tagging controls SCSI command tag queuing support for each drive channel. This function normally remains enabled. Disable only when using older SCSI technology that does not support command tag queuing. This parameter is only available through the RAIDfx configuration utility.

➢ See also "Controller Specifications" on page 1, "Drive Channel" on page 27, "Firmware Versions" on page 95, "RAIDfx" on page 53, and "SCSI Cabling" on page 99.

Conservative Cache Mode Parameter

Default=Disabled. The Conservative Cache Mode function is provided to allow a controller an extra degree of data safety when operating in a failed-over state. This function switches write-back caching to write-through operation after a controller failure occurs. When the failed controller is replaced, write-back caching operations resume. This parameter takes effect immediately, without resetting the controllers.

There is a performance loss during Conservative Cache Mode. Enabling this option has no effect during normal operation.

➢ See also “Conservative Cache Mode” on page 22, “Failover and Failback” on page 35, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Controller Name

SCOPE: The Controller Name controller parameter was implemented with firmware version 7.0. This command is not supported by SAM/GAM or VT100 at this time.

The Controller Name parameter allows the user to assign names to the controllers. This parameter takes effect immediately, without resetting the controllers.

Controller Present/Fault Select

SCOPE: This parameter is only supported on the DACSX controller.

Default=A. This function allows the user to select between two sets of backplane signals, A or B, to use when detecting the presence of a partner controller. This option has no effect in a simplex environment.

➢ See also “DACSX” on page 25, “Dual-Active Controller Configuration” on page 32, “Partner Controller” on page 133, and “Simplex” on page 60.

Controller Present/Fault Signals

SCOPE: This parameter is referred to as Duplex Fault Signals in firmware versions 7.0 or greater. Refer to “Duplex Fault Signals” on page 77 for more information. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Data Bus Width

SCOPE: This parameter is ignored by controllers supporting Fibre Channel drive channels and is not supported with firmware versions 7.0 or greater.

Default=Disabled. Enabling this option prevents the controller from negotiating for Wide SCSI transfers. This function should be enabled only when connecting Wide SCSI (16 bit) disk drives to the controller using a Narrow SCSI (8 bit) cable.

➢ See also “Controller Specifications” on page 1, “SCSI Cabling” on page 99, and “Transfer Width” on page 88.

Device Combing

SCOPE: This parameter is referred to as Coalescing in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

The Device Combing function provides device queuing coalescing optimization. The function enables data traffic coalescing (combining of address adjacent I/Os) on the traffic of each device. This joins the data from adjacent I/Os into a single I/O to improve performance. If disk drive queues are deep, the disk drives take care of this activity and controller coalescing has very little effect. If the queue limit of the disk drives is decreased (Queue limit = 2) and coalescing is enabled, the controller coalesces data. Coalescing enabled does not degrade performance even with a larger Queue limit. This parameter takes effect immediately, without resetting the controllers.

➢ See also “Physical Disk Drive” on page 135, and “Queue Limit” on page 83.

Disable BUSY Status During Failback

SCOPE: This parameter is available only with firmware version 7.0 or greater.

Default=Disabled. The Disable BUSY Status During Failback function allows the controller to disregard new requests without returning a BUSY status. If enabled, during failback, the surviving controller ignores all new requests and does not return any status. If disabled, the surviving controller returns a BUSY status to new commands received from the host during the cache flush operation. A controller reset is required before this parameter takes effect.

Disable Check Condition for Invalid LUN

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Enabled. When enabled, disables Check Condition for an invalid LUN. This effects the handling of the INQUIRY command when the referenced LUN is invalid. If enabled: the INQUIRY command returns data with the peripheral qualifier indicating that the peripheral device is not connected. If disabled: the INQUIRY command will be failed with a Check Condition of Illegal Request, LUN Not Supported. A controller reset is required before this parameter takes effect.

➢ See also “Logical Unit Number (LUN)” on page 125 and “UNIX Support and Firmware Types” on page 64.

Disable Queue Full Status

SCOPE: This parameter is referred to as On Queue Full Give Busy in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Disabled. Enabling this parameter sets the controller to return a busy status when a queue full condition is detected. Disabling this parameter sets the controller to return a queue full status. When a command is received and the controller detects a queue full condition, it normally returns Queue Full status. This parameter is intended to help hosts that are confused by queue full. A controller reset is required before this parameter takes effect.

➢ See also “Host / Server” on page 121.

Disable Wide Operation

SCOPE: This parameter is ignored by controllers supporting Fibre Channel host channels and has no effect. This parameter is referred to as Host Transfer Width in firmware versions 7.0 or greater. Refer to “Host Transfer Width” on page 80 for more information.

Disconnect on First Command

SCOPE: This parameter is no longer supported with firmware versions 4.2 or greater.
This parameter is available only through the RAIDfx configuration utility.

Default=Disabled. This option enables Disconnect on First Command for the drive channel SCSI busses. A controller reset is required before this parameter takes effect.

➢ See also "Bus" on page 111, "Controller Specifications" on page 1, "Drive Channel" on page 27, and "SCSI Cabling" on page 99.

Disk Startup Delay 1

SCOPE: This parameter is referred to as Spin-up Disk Delay in firmware versions earlier than 7.0.

Default=6. The Disk Startup Delay 1 function varies with the selection of the *Disk Startup Mode* parameter. If AUTOSPIN is selected, this function specifies the number of seconds between physical device spin-up cycles. If PWRSPIN is selected, this function sets the number of seconds before the controller issues start up commands. This value should be set to the device's power-on-to-spin time. Possible settings are 0 to 225 seconds. This parameter takes effect immediately, without resetting the controllers.

➢ See also "Disk Startup Delay 2" on page 75, "Disk Startup Mode" on page 75, and "Disk Startup Number of Devices" on page 76.

Disk Startup Delay 2

SCOPE: This parameter is referred to as Spin-up Sequence Delay in firmware versions earlier than 7.0.

Default=0. For AUTOSPIN mode, the Spin-up Sequence Delay function sets the number of seconds for a SCSI ID-based motor spin delay. The device has power immediately, but waits n seconds multiplied by its SCSI ID before starting its motor. This parameter is normally set to 0 for AUTOSPIN mode. For PWRSPIN mode, the Disk Startup Delay 2 function specifies the number of seconds between physical device start-up cycles. The value should be set to the enclosure's power delay between powering banks of devices.

Possible settings are 0 to 90 seconds for firmware versions 6.3 or earlier. Possible settings are 0 to 225 seconds for firmware versions 7.0 or greater. A controller reset is required before this parameter takes effect.

➢ See also "Disk Startup Delay 1" on page 75, "Disk Startup Mode" on page 75, and "Disk Startup Number of Devices" on page 76.

Disk Startup Mode

SCOPE: This parameter is referred to as Spin-up Option in firmware versions earlier than 7.0.

Default=Autospin (0). The Disk Startup Mode function controls how the disk drives in the array are started (spun up). There are three different startup modes that may be selected by the user. This parameter takes effect immediately, without resetting the controllers.

- AUTOSPIN (0)—The AUTOSPIN mode issues start commands to all devices automatically. This mode waits the amount of time specified in Disk Startup Delay 1, issues Start Unit commands to the devices as specified in Disk Startup Number of Devices, then waits the specified Disk Startup Delay 1 again. This cycle repeats until all devices have been issued Start Unit commands. This mode proceeds with a sequence delay, specified by the Disk Startup

Delay 2 parameter, while the drives become ready. Ready is equal to the *on power* spin-up mode. The sequence delay will not exceed a maximum of 75 seconds. The sequence delay is normally 0 when drives are jumpered to spin immediately on power-up, but can be set to a number of seconds for a target ID-based motor spin delay (where the drive has power immediately, but waits n seconds multiplied by its target ID before starting its motor).

- PWRSPIN (1)—Devices spin on power application. This mode is designed for systems where drives are powered on in sequence by the drive enclosure. This mode waits the amount of time specified in Disk Startup Delay 1, after which the first bank of devices is expected to be ready. The first bank of drives are then checked. This wait-check cycle repeats for each subsequent bank of drives. Disk Startup Number of Devices is ignored for this mode. The Disk Startup Delay 1 is set to the drive power-on-to-spin time and the Disk Startup Delay 2 is set to the enclosure delay between powering banks of drives. This mode assumes all drives with the same target ID are in a bank.
- WSSUSPIN (2)—The controller waits for the host to issue a SSU (Start/Stop Unit) command then performs the Autospin mode described above. This mode causes the drive initialization to stall until the host sends the controller a start unit command, then proceeds with the Autospin mode.

NOTE: As of firmware version 7.4, the WSSUSPIN startup option is obsolete and treated as Autospin.

The following table details the relationship between Disk Startup Mode, Disk Startup Number of Devices, Disk Startup Delay 1, and Disk Startup Delay 2.

Disk Spin-up Option Parameters

Disk Startup Mode	=	AUTOSPIN	PWRSPIN	WSSUSPIN
Disk Startup Number of Devices	=	#Devices/spin	Undefined	#Devices/spin
Disk Startup Delay 1	=	Device spin wait	Initial delay	Device spin wait
Disk Startup Delay 2	=	(0)	Sequence delay	(0)

➢ See also “Disk Startup Delay 1” on page 75, “Disk Startup Delay 2” on page 75, and “Disk Startup Number of Devices” on page 76.

Disk Startup Number of Devices

SCOPE: This parameter is referred to as Disks per Spin in firmware versions earlier than 7.0.

Default=2. This option specifies the number of physical disk drives to be spun up at one time. The controller waits the amount of time determined by Disk Startup Delay1 before spinning up the next group of disk drives. Possible settings are 1 through 8. This parameter takes effect immediately, without resetting the controllers.

This parameter is ignored if PWRSPIN mode is selected.

➢ See also “Disk Startup Delay 1” on page 75, “Disk Startup Delay 2” on page 75, and “Disk Startup Mode” on page 75.

Disk Write Through Verify

SCOPE: This parameter is referred to as Write Through Verify in firmware versions earlier than 7.0.

Default=Disabled. The Disk Write Through Verify function enables Force Unit Access for reads and writes during error handling. Force Unit Access bypasses the caches and forces all reads and writes directly to or from the disk. A controller reset is required before this parameter takes effect.

NOTE: For some devices, enabling Force Unit Access reduces sequential write performance by up to 86%.

Disks Per Spin

SCOPE: This parameter is referred to as Disk Startup Number of Devices in firmware versions 7.0 or greater. Refer to “Disk Startup Number of Devices” on page 76 for more information.

Duplex Fault Signals

SCOPE: This parameter is referred to as Controller Present/Fault Signals in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater. This parameter is not supported on the FFx or FFx2 controllers.

Default=Enabled. The Duplex Fault Signals function is provided to inform a controller that certain signals should be used to detect the presence or absence of a partner controller. A controller reset is required before this parameter takes effect.

NOTE: If you have enabled *Automatic Failback or Autorestore*, select Enabled for this function also. This parameter is necessary for hot plugging controllers and automatic failback.

➢ See also “Automatic Failback or Autorestore” on page 70, “Dual-Active Controller Configuration” on page 32, “Failover and Failback” on page 35, “Hot Plugging” on page 121, “Partner Controller” on page 133, and “Reset Controllers” on page 54.

Duplex Fault Signals on Channel 4

SCOPE: This is a legacy command. This parameter is supported only on the DACSX controller. This parameter is not applicable to fibre host controllers and has no effect.

Default=Enabled. This function informs a controller that certain signals should be used to detect the presence or absence of a partner controller. A controller reset is required before this parameter takes effect.

NOTE: If you have enabled *Automatic Failback or Autorestore*, select Enabled for this function also. This parameter is necessary for hot plugging controllers and automatic failback.

Elevator

SCOPE: This parameter is no longer supported with firmware versions 4.2 and greater. This parameter is not available through the LCD/VT100 configuration utility.

Default=Disabled. This option attempts to improve performance by keeping the heads moving in the same direction as long as possible when accessing the disk drives.

Enable Background Initialization

SCOPE: This parameter was implemented with firmware version 7.7 and is supported on the FFx, FF2, and FFx2 controllers

Default=Enabled. Enabling the Background Initialization controller parameter allows the controller to perform a background initialization automatically when an uninitialized system drive is detected. Background initialization makes the system drive instantly available for host read and write access. If this parameter is disabled, foreground initialization must be performed. Foreground initialization blocks access from the host until the system drive has completed the initialization process. Depending on the size of the system drive, foreground initialization may take several hours.

➢ “Background Initialization” on page 11, “Initialize System Drive” on page 123, “System Drives” on page 149

Failover Topologies

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 and greater. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

This option sets the Fibre Channel host port topology. The choices are:

- Inactive Port—One active and one inactive host port per controller. The active port is for normal traffic while the inactive port is for the partner’s traffic when it is failed over. This topology is only recommended for the DACSF, DACFL, and DACFF (FF) controllers.
- Multiport—All host ports are active and connected to individual fibre loops. This topology requires alternate path software and does not support transparent controller failover/failback.
- Clustering—This topology is not currently supported.
- Multiple Target ID—Provides the controller with the ability to function as multiple target ports on a single arbitrated loop. During failover, the surviving controller enables a virtual port, impersonating the ports from the failed controller. This topology is only available on the FFx, FF2, and FFx2.
- Master-Slave— Master/Slave topology requires one active controller while the other controller remains inactive. If the active controller fails in the Master/Slave topology, the surviving controller joins the loop and assumes the responsibilities for the failed controller. This topology is not recommended for use with firmware versions greater than 5.4.

A controller reset is required before this parameter takes effect.

➢ See also “Active Fibre Port” on page 109, “Controller Specifications” on page 1, “Failed Controller” on page 118, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Host / Server” on page 121, and “Inactive Fibre Port” on page 123.

Force Simplex

Default=Disabled. The Force Simplex function allows duplex (dual-active controller) firmware to function in a simplex (single controller) environment. The Force Simplex function forces the duplex firmware to ignore some of the dual-active operations. A controller reset is required before this parameter takes effect.

ATTENTION: **Do not enable Force Simplex unless it is required. Disabling Force Simplex and returning to a dual-active controller mode requires that each controller be reconfigured independently of the other controller. This is accomplished by removing one controller from the system, reconfiguring the remaining one, then swapping the controllers and reconfiguring the second controller, then reinstalling the first controller.**

NOTE: The FFx and FFx2 controllers automatically upgrade from simplex to duplex upon detecting a partner controller during startup.

➢ See also “Dual-Active Controller Configuration” on page 32, “Duplex” on page 116, “Mylex Online RAID Expansion (MORE)” on page 49, “Replacement Controller Behavior with Forced Simplex Enabled” on page 53, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, and “Simplex” on page 60.

Frame Size

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 or greater. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=2KB. This option sets the host Fibre Channel data frame size to either 512 bytes, 1 KB or 2 KB. A frame size of 2 KB is recommended since it provides the largest packet transfers and the most throughput. A controller reset is required before this parameter takes effect.

➢ See also “Controller Specifications” on page 1, “Fibre Channel” on page 43, “Frame” on page 119, and “Host / Server” on page 121.

Generate Debug Dump

SCOPE: This parameter was implemented with firmware version 7.4 and is supported on the FFx, FF2, and FFx2 controllers.

This option generates debug dump information during a controller abort. During the abort sequence, data is written to a selected disk drive. This data can be retrieved by the host using the rddump.exe utility. Data is written to the first 4 MB of the 16 MB RDA.

➢ See also “Debug Dump” on page 25.

Hard Loop IDs

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 and greater. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

This option allows you to enable or disable use of the controller/host port and, if enabled, to define the hard loop ID. Specifying the hard loop ID means that the same ID will always be requested. The valid range for loop IDs is from 0 to 125. A controller reset is required before this parameter takes effect.

Fibre Channel arbitrated loop nodes acquire loop IDs in this order: Previous (LI_PA), Hard (LI_HA), Soft (LI_SA). Any soft ID is used in the next LIPA cycle.

➢ See also “Controller/Host Port” on page 112, “Controller Specifications” on page 1, “Fibre Channel” on page 43, “Hard Loop ID” on page 120, and “Host / Server” on page 121.

Host (SCSI) Bus Reset Delay

SCOPE: The Host (SCSI) Bus Reset Delay function is supported only on controllers with a SCSI host interface, DACSX and DACSS. This parameter is not applicable to fibre host controllers and has no effect.

Default=No Reset (0). The Host (SCSI) Bus Reset Delay parameter allows a controller to reset the host in failover and failback situations. If No Reset (0) is set, no SCSI bus reset is generated on the host channel(s) after a failover or failback occurs. If Reset Delayed (1 to 14) is set, a SCSI bus reset is generated on the host channel(s) approximately 1-14 seconds after a failover or failback occurs. If Reset Immediately (15) is set, a SCSI bus reset is generated immediately on the host channel with no delay. Generally, this parameter should be set to 0. A controller reset is required before this parameter takes effect.

NOTE: If you are using the Solaris™ operating system, set this parameter to 5.

➢ See also “Bus” on page 111, “Controller Specifications” on page 1, “DACSS” on page 25, “DACSX” on page 25, “Failover and Failback” on page 35, and “Reset Controllers” on page 54.

Host Transfer Width

SCOPE: This parameter is ignored by controllers supporting Fibre Channel host channels and has no effect. This parameter is referred to as Disable Wide Operation in firmware versions earlier than 7.0.

Default=Disabled. Disabling this option disables wide operations on all host channels. The controller will not negotiate wider than this. Narrowing the host data path without need results in performance loss.

This is a legacy command. This parameter is not applicable to fibre host controllers and has no effect.

➢ See also “Controller Specifications” on page 1, “Host Channel” on page 121, and “SCSI Cabling” on page 99.

Maximum IOPs

SCOPE: This parameter is no longer supported with firmware versions 4.2 and greater. This parameter is not available through the LCD/VT100 configuration utility.

Default=244. This option sets the maximum number of IOP structures allowed for controller data traffic operations, limiting the Queue Tag Limit value. Values of 64 through 244 can be defined. A controller reset is required before this parameter takes effect.

➢ See also “Queue Limit” on page 83.

Mode Select SRA Enable

SCOPE: This parameter is no longer supported.

Default=Disabled. This parameter allows the user to enable Super Read Ahead through the controller's mode pages.

No Pause on Controller Not Ready

SCOPE: This parameter is referred to as No Pause for Host Commands in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Disabled. This parameter turns the pause off or on for certain commands when the controller is not ready. Normally, when the controller is starting up, certain commands encounter a pause. This only happens when the controller has not reached STARTUP COMPLETE. The pause lasts one second. The commands affected are: PREFETCH, READ/WRITE, READ/WRITE EXTENDED, TUR, VERIFY, and WRITE VERIFY. A controller reset is required before this parameter takes effect.

Node Name Retention

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 and greater.This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

ATTENTION: **If the host uses node names to locate logical devices, or system drives, this option must be enabled, or access to data may be lost.**

Default=Disabled. This option disables/enables a failed controller's node name to be retained through a controller failure. When disabled, each controller shares its node name with its partner controller through failover; however, when failback occurs, the replacement controller uses its own node name. When enabled, each controller shares its node name with its partner controller and those names are used through all phases of failover and failback. A controller reset is required before this parameter takes effect.

Not having this bit set has serious ramifications if the controllers are connected to a host that uses node names to locate logical devices.

➢ See also "Controller Specifications" on page 1, "Failed Controller" on page 118, "Failover and Failback" on page 35, "Fibre Channel" on page 43, "Host / Server" on page 121, "Node Name" on page 130, "Partner Controller" on page 133, and "Replacement Controller" on page 145.

On Queue Full Give Busy

SCOPE: This parameter is referred to as Disable Queue Full Status in firmware version 7.0 and greater. Refer to "Disable Queue Full Status" on page 74 for more information. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Operational Fault Management

Default=Enabled. The Operational Fault Management function allows the controller to take autonomous actions when a failure occurs. Actions that the Operational Fault Management function monitors and reports include drive failures, background activity completion status, and enclosure events. This function should remain enabled during normal controller operation. Do not disable this function unless specifically instructed to do so as part of a trouble-shooting diagnostic activity. The Operational Fault Management function works in conjunction with *Automatic Rebuild Management* and features in AEMI, SAF-TE and SES certified disk array enclosures to detect the removal of a failed disk drive.

Without Operational Fault Management enabled, a host program or operator must handle all failure cases. This parameter takes effect immediately, without resetting the controllers.

➢ See also “AEMI” on page 6, “Automatic Rebuild Management” on page 71, “Failed Disk Drive” on page 118, “Fault Management” on page 41, “SAF-TE” on page 55, and “SES” on page 58.

Override Multiport Reset

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 and greater. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

The Override Multiport Reset restricts internal resets to ports that have logical devices reserved through that port. If enabled, an internal reset is executed by a port only if a logical device has been reserved through that port. If disabled, internal resets are not qualified by logical device reservations. This parameter takes effect immediately, without resetting the controllers.

When this bit is set, only the Reset Event receiving port and the reserved port are reset.

➢ See also “Cache Coherency” on page 17 and “Implied Reserve and Release” on page 46.

PCI Latency Control

SCOPE: This parameter is only supported on controllers with Fibre Channel host support using firmware versions 4.0 and greater. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater. This parameter is not supported on the FFx or FFx2 controllers.

Default=Short. The PCI Latency Control function allows adjustment of the Fibre Channel processor’s use of the PCI bus. This function controls the amount of data each processor can burst across the primary bus before relinquishing bus ownership to the next device. PCI Latency Control takes effect only when both ports are active and are arbitrating for the bus. PCI Latency Control allows the integrator to tune the controller’s operating parameters for specific applications. For maximum throughput, LONG is recommended and is equivalent to the time necessary to transfer 1024 bytes; MEDIUM is equivalent to 684 bytes and SHORT is equivalent to 512 bytes. A controller reset is required before this parameter takes effect.

➢ See also “Controller Specifications” on page 1, “Fibre Channel” on page 43, “Host / Server” on page 121, and “Port” on page 135.

Queue Limit

SCOPE: This parameter is referred to as Queue Tag Limit in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=32. This option specifies the maximum allowed queue depth for tagged commands to all attached disk drives. This value is further limited to the disk drive's own tag limit. Allowed values are 1 to 255. A setting of 1 is similar to no tags. If using device combing, set the queue tag limit to 2. This parameter takes effect immediately, without resetting the controllers.

➢ See also "Device Combing" on page 73.

RAID5 Algorithm

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Right Asymmetric. The alternative setting sets the RAID 3 and 5 algorithm to left symmetric. This parameter determines the method for distributing parity for a RAID 3 and 5 array.

This value is set only at the time a new configuration is created.

ATTENTION: **Data loss will occur after you change the RAID 5 Algorithm on a controller with existing system drives if data is not backed up.**

➢ See also "Parity" on page 133, "RAID Levels" on page 137, and "Striping" on page 147.

Read Ahead

Default=Enabled. The Read Ahead function improves data retrieval performance by allowing the controller to read into cache a full stripe of data at a time. This greatly improves the percentage of cache hits. For small transfers, the read ahead algorithm helps with performance. For example, if the stripe size is set to 8 KB and the host requests 1 KB of data, when this function is enabled the controller will read ahead the full 8 KB. When the host requests the next 1 KB block, that data will already be in the controller's cache. This function should remain enabled during normal controller operation. A controller reset is required before this parameter takes effect.

➢ See also "Blocking Factor" on page 72, "Cache" on page 111, and "Stripe Size" on page 64.

Read/Write Control

This physical device parameter changes the configured write policy. Write-Through writes data to the device before returning completion status to the host. Write-Back returns a completion status to the host as soon as the cache receives the data. The target device receives the data at a more appropriate time. This parameter takes effect immediately, without resetting the controllers.

➢ See also "Conservative Cache Mode" on page 22, "Standard Data Caching" on page 63, "Write-Back Cache" on page 151, and "Write-Through Cache" on page 151.

Reassign Restricted to One Block

Default=Disabled. The Reassign Restricted to One Block function limits reassigning failures to the single failed block. This parameter is limited to recovered errors and medium errors. If the sense on the error does not indicate one of these errors, then this setting does not apply. When the Reassign Restricted to One Block is disabled, all reassigns are for the entire I/O process, possibly a large number of blocks where not all have failed. A controller reset is required before this parameter takes effect.

See also "Disk Media Error Management" on page 26.

Rebuild and Check Consistency Rate

Default=50. The Rebuild and Consistency Check Rate parameter approximates one-half of the percentage of available rebuild cycles to be used. CPU utilization is always shared with data traffic. This function also determines the amount of resources the controller devotes to Consistency Check and MORE operations. Integer values from 0-50 can be defined. Entering a value of 50 means that all of the resources that can be dedicated for the operations are utilized. For low priority and high array performance, specify a lower value. For high priority and reduced array performance, select 50. This parameter takes effect immediately, without resetting the controllers.

➢ See also "Automatic Rebuild" on page 10, "Background Initialization" on page 11, "Consistency Check" on page 23, "Failed Disk Drive" on page 118, "Mylex Online RAID Expansion (MORE)" on page 49, and "Rebuild" on page 144.

Reset Propagation

SCOPE: This parameter applies only to controllers with SCSI host interfaces, including the DACSX and DACSS controllers. This parameter is not applicable to fibre host controllers and has no effect. This parameter is only available through the RAIDfx configuration utility using firmware versions earlier than 7.0.

Default=Disabled. The Reset Propagation function allows a port to issue an internal reset without causing a reset event to occur on its attached interface. If enabled, a port that issues an internal reset propagates the reset by causing a reset event to occur on its attached interface. If disabled, a port will not cause a reset event on its attached interface as part of issuing an internal reset.

This option permits incoming SCSI bus resets to be propagated out other host ports. Only those ports that are defined accessible through the LUN mapping feature (that is, only those ports that access the affected system drive) are affected. This feature is required for configurations utilizing Microsoft Cluster Server with SCSI host connections.

➢ See also "Cache Coherency" on page 17, "Controller/Host Port" on page 112, "Controller Specifications" on page 1, "Host / Server" on page 121, "Implied Reserve and Release" on page 46, "LUN Mapping" on page 47, "Microsoft Cluster Server Configurations (MSCS)" on page 47, "Programmable LUN Mapping" on page 52, "SANmapping™" on page 56, and "System Drive Affinity" on page 64.

SAF-TE Data for UPS Support

SCOPE: This parameter is referred to as Use of UPS in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Enabled (0). The SAF-TE Data for UPS Support function provides UPS monitoring via the SAF-TE vendor unique bytes described in the SAF-TE specification. The current state of all SAF-TE inputs can also be determined via SAF-TE pass-through commands.

Enabled (0)—UPS monitoring is enabled if it is also supported by the system enclosure.

Disabled (1)—UPS monitoring is disabled.

UPS monitoring currently supports the following input/outputs.

- AC FAIL (input)—The UPS has detected a loss of ac power. The controller switches to conservative cache mode for this event.
- BAT LOW (input)—The UPS has detected that its battery power is now limited. The controller switches to conservative cache mode for this event.
- Shutdown (output)—If AC FAIL and BAT LOW are active and the cache has been flushed, the controller issues a shutdown signal to the UPS.

This parameter takes effect immediately, without resetting the controllers.

➢ See also “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, and “SAF-TE” on page 55.

SCSI Transfer Rate or Mega-transfers/sec

SCOPE: This parameter is ignored by controllers supporting Fibre Channel drive channels. This parameter is referred to as Transfer Speed in firmware versions 7.0 and greater. Refer to “Transfer Speed” on page 87 for more information.

Serial Control

SCOPE: The Serial Control parameter is a read-only function that reports serial port parameters such as data bits, stop bits, and parity. This field is not selectable and appears for information purposes only.

Each parameter (data bits, stop bits, and parity) is available only through the RAIDfx configuration utility using firmware versions earlier than 7.0.

Serial Port Baud Rate

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=19200. This option sets the serial port baud rate for VT100 or debug use. Options are 2400, 4800, 9600, and 19200. A controller reset is required before this parameter takes effect.

➢ See also “Embedded Configuration Utility” on page 35 and “VT100” on page 65.

Serial Port Debug Type

SCOPE: This option is used for engineering purposes only.

This option configures the serial port for debug output and input. An external terminal device is required. This mode may be forced via a jumper that overrides this setting. A controller reset is required before this parameter takes effect.

NOTE: Setting this mode of operation during normal traffic results in a performance loss due to activity on the serial port.

➢ See also “Embedded Configuration Utility” on page 35 and “VT100” on page 65.

Serial Port Usage

NOTE: The debug port is for development use only.

Default=Debug. This parameter sets serial channel A to be used as either the SLP/VT100 port or the debug port. A controller reset is required before this parameter takes effect.

➢ See also “Embedded Configuration Utility” on page 35 and “VT100” on page 65.

Simplex - No Reset

SCOPE: This parameter is intended for use in simplex environments only.

Default=Disabled. Enabling this function inhibits the reset signal normally sent from one controller to another controller in a dual-active controller system. A controller reset is required before this parameter takes effect.

➢ See also “Dual-Active Controller Configuration” on page 32 and “Simplex” on page 60.

Smart Large Host Transfers

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Enabled. The Smart Large Host Transfers function allows selection of Coalesce (fewer disconnects on large transfers) or As Available (more disconnects) on host data transfers. This function is most effective on SCSI, but has some benefit on Fibre Channel systems. The Smart Large Host Transfers function takes effect for transfers larger than the stripe size. This parameter takes effect immediately, without resetting the controllers.

➢ See also “Device Combing” on page 73 and “Stripe Size” on page 64.

Spin-up Disk Delay

SCOPE: This parameter is referred to as Disk Startup Delay 1 in firmware versions 7.0 and greater. Refer to “Disk Startup Delay 1” on page 75 for more information.

Spin-up Option

SCOPE: This parameter is referred to as Disk Startup Mode in firmware versions 7.0 and greater. Refer to “Disk Startup Mode” on page 75 for more information.

Spin-up Sequence Delay

SCOPE: This parameter is referred to as Disk Startup Delay 2 in firmware versions 7.0 and greater. Refer to “Disk Startup Delay 2” on page 75 for more information.

Stripe Size Parameter

SCOPE: This parameter is not supported with firmware version 7.0 and greater. If using firmware 7.0 or greater, the stripe size must be set while creating the configuration.

Default=8 KB. The Stripe Size function is used to tune the controller performance for a specific environment or application. Generally, stripe size optimization is as follows:

- Smaller stripe sizes provide better performance for random I/O.
- Larger stripe sizes provide better performance for sequential transfers.

The default setting is 8 KB (optimum random I/O performance and reduced sequential throughput). Changing the stripe size to 16 KB, 32 KB, or 64 KB alters the way data is written on the disk drives. Once stripe size is set, you cannot change the stripe size without first copying all of the data off of the array, changing the stripe size, reinitializing the system drives, and copying the saved data back onto the array. In a large capacity array, this operation could take several hours. Be sure to define the stripe size before configuring the array. A controller reset is required before this parameter takes effect.

ATTENTION: **Data loss will occur after changing the stripe size on a controller with existing system drives. Always back up all data before making a stripe size change.**

ATTENTION: **You must reset the controllers and re-initialize the system drives after a new stripe size is saved.**

➢ See also “Cache Line Size” on page 18, “Cache Line Size Parameter” on page 72, and “Stripe Size” on page 64.

Super Read Ahead

Default=Disabled. This function increases performance for applications that must access large blocks of sequential data. This option only improves performance for large sequential read operations and has no effect on write operations. This function incorporates intelligent data request monitoring to track data requests by the host. With Super Read Ahead enabled, the controller detects requests for data that are stored in sequence on the drives. It reads the data into cache so that the cache remains at least one request ahead of the host. This function should remain disabled during normal controller operation. A controller reset is required before this parameter takes effect.

Transfer Speed

SCOPE: This parameter is ignored by controllers supporting Fibre Channel drive channels. This parameter is referred to as SCSI Transfer Rate in firmware versions earlier than 7.0.

Default=80 Mhz. The Transfer Speed function sets the maximum transfers rate for each device. The possible settings are 0 (asynchronous), 5, 10, 20, 40, 80 (LVD), and 1000 (fibre). This setting produces 10 MB/sec transfers for Fast SCSI if 8-bit SCSI disk drives are used, or 20 MB/sec transfers for Fast and Wide SCSI disk drives. The 20 MHz setting allows 20 MB/sec transfers for 8-bit Wide Ultra SCSI devices, or 40 MB/sec transfers for Ultra Wide devices. This parameter takes effect immediately, without resetting the controllers.

NOTE: The default setting should be changed only if problems are encountered in communicating with a device. Do not change the default setting unless you are doing so as part of a trouble-shooting activity. Problems communicating with a disk drive can be caused by several conditions; for example, improper termination, wrong target ID setting, SCSI cable is too long, faulty equipment, and so on.

➢ See also “Controller Specifications” on page 1, “Drive Channel” on page 27, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

Transfer Width

SCOPE: This parameter is ignored by controllers supporting Fibre Channel drive channels.

Default=16 bits. The Transfer Width function is a physical device parameter and determines the maximum data transfer width size, in bits, for the drive channels. The possible settings are serial, 8 bits, and 16 bits. This parameter takes effect immediately, without resetting the controllers.

True Verify

SCOPE: This parameter is referred to as True Verification of Data in firmware versions earlier than 7.0. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Disabled. If this option is disabled, the Verify command returns a status without data checking. If this option is enabled, the Verify command checks data before returning a status. Enabling this parameter compromises performance. A controller reset is required before this parameter takes effect.

Use of UPS

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is referred to as SAF-TE DATA for UPS Support in firmware versions 7.0 and greater. Refer to “SAF-TE Data for UPS Support” on page 84 for more information

Vendor Unique Direct Command

SCOPE: This parameter is not supported by firmware version 7.0 and greater.

This option sets the direct command opcode to 0x20. This parameter takes effect immediately, without resetting the controllers.

Vendor Unique Pass-Through Command

SCOPE: This parameter is not supported by firmware version 7.0 and greater.

This option sets the direct command opcode to 0x21. This parameter takes effect immediately, without resetting the controllers.

Vendor Unique TUR

SCOPE: This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

Default=Disabled. The Vendor Unique Test Unit Ready (TUR) function enables a vendor unique TUR response sent to an offline LUN. If disabled, a hard error status is returned. If enabled, a Not Ready status is returned. A controller reset is required before this parameter takes effect.

➢ See also “Logical Unit Number (LUN)” on page 125 and “UNIX Support and Firmware Types” on page 64.

Write Through Verify

SCOPE: This parameter is referred to as Disk Write Through Verify in firmware versions 7.0 and greater. Refer to “Disk Write Through Verify” on page 76 for more information. This parameter is not available through the LCD/VT100 configuration utility using firmware versions earlier than 7.0. This parameter is available through the VT100 and embedded configuration utility using firmware versions 7.0 or greater.

FREQUENTLY USED PROCEDURES

This section describes strategies for configuring an array using the Mylex controllers and instructions for frequently used procedures. The information applies to all Mylex controllers, unless otherwise noted. Refer to the appropriate installation guide for more detailed information on installing and using any of the controllers.

Adding SES Enclosures

This feature allows users to add disk enclosures (one or more) to a configured system while the system continues to operate. After the enclosure or enclosures have been added to the system, this feature starts the SES monitoring process for the new enclosure. The user can then configure the additional disk capacity without restarting the system.

Additional enclosures are added to the configured system using the following procedure.

1. Check for ID conflicts. Each enclosure and disk drive must have a unique ID.
2. Resolve any ID conflicts.
3. Connect drive channels from the existing system to the new enclosure or enclosures.
4. Supply power to the new enclosure or enclosures. This causes a LIP on the drive channel notifying the controller that new disk drives have been added to the fibre loop.
5. Wait for the controller to supply power to the disk drives. All disk drives must have completed the spin-up process before proceeding.
6. Issue the scan for additional enclosures. This may be an option incorporated in the configuration utility, or issued as a direct SCSI command.
7. After the SES process completes polling the loop, the new enclosures and disk drives are ready for configuration.
8. If the new enclosure(s) are not detected, remove and insert a disk drive froma new enclosure. Removing and inserting a disk drive generates a LIP, and the new disk drives will join the existing fibre loop.
9. Issue the scan for additional enclosures a second time.

➢ See also “AEMI” on page 6, “Conservative Cache Mode” on page 22, “Conservative Cache Mode Parameter” on page 72, “Drive State Management” on page 31, “Environmental Device” on page 117, “Fault Management” on page 41, “Soft Addressing Detection” on page 61, “Firmware Versions” on page 95, “SAF-TE” on page 55, “SCSI” on page 145, “SES” on page 58, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

CAP Strategy for Selecting a RAID Level

Capacity, Availability, and Performance: these are the three benefits that all Mylex RAID solutions have in common. The terms, collectively known as "CAP," should characterize your expectations of the disk array subsystem.

Each RAID configuration has requirements based on capacity, availability, and performance. An optimal solution is achieved by prioritizing and combining these benefits. Creating a system with high availability (data redundancy) sacrifices the total available capacity by using some of the disk drives for redundancy.

The following guidelines will help with selecting which benefits are a priority and how to combine these features to create the optimal RAID configuration.

Configuring for Maximum Capacity

To see the relationship between the different RAID levels and effective capacities offered for the quantity *X* disk drives of *N* capacity, refer to the following table. As an example, computed capacities for six disk drives of size 2GB each are provided.

NOTE: The following table presents theoretical values for drive capacity. Actual drive capacity is less due to drive size truncation and configuration on disk information.

RAID Level Maximum Capacity

RAID Level	Effective Capacity	Example: (Capacity in GB)
0	X*N	6*2 = 12
1	2*N/2	2*2/2 = 2
3 and 5	(X-1)*N	(6-1)*2 = 10
0+1	(X*N)/2	(6*2)/2 = 6
JBOD	X*N	6*2 = 12

The greatest capacities are provided by RAID 0 and JBOD, with the entire capacity of all disk drives being used. These two solutions, however, provide no fault tolerance.

RAID 3 and RAID 5 give the next best capacity, followed by RAID 0+1 and RAID 1.

Configuring for Maximum Availability

When considering optimizing for availability, it is important to understand some of the terminology concerning the condition or state of array operation. These definitions are presented in the following table.

Array Operating Conditions/States

Array Condition	Meaning
Online Optimal (Normal)	The array is operating in a fault-tolerant mode, and can sustain a disk drive failure without data loss.
Online Critical (Critical)	The array is functioning and all data is available, but the array cannot sustain a second disk drive failure without potential data loss.

Array Operating Conditions/States (Continued)

Array Condition	Meaning
Online Critical and Rebuilding or Degraded	The array is functioning and all data is available, but the array cannot sustain a second disk drive failure without potential data loss. Additionally, a reconstruction or rebuild operation is taking place, reducing the performance of the array. The rebuild operation takes the array from a critical condition to an optimal condition.
Offline	The array is not functioning. If the array is configured with a redundant RAID level, two or more of its member disk drives are not online. If the array is configured as a RAID 0 or JBOD, one or more of its member disk drives are not online.
Not fault tolerant	No fault-tolerant RAID levels have been configured for any of the disk drives in the array.

An additional measure of fault tolerance (or improved availability) can be achieved using an online spare (hot spare or standby) disk drive. This disk drive is powered on but idle during normal array operation. If a failure occurs on a disk drive in a fault-tolerant set, the online spare disk drive takes over for the failed disk drive, and the array continues to function in a fully fault-tolerant mode after it completes its automatic rebuild cycle. This means that the array can suffer a second disk drive failure after rebuild and continue to function before any disk drives are replaced.

Impact of Controller Cache on Availability

Every Mylex controller has a disk cache. The physical memory is used to increase the performance of data retrieval and storage operations. The amount of disk cache varies with the controller model.

The controller may report to the operating system that the write is complete as soon as the controller receives the data. This is referred to as Write-Back (WB) cache. In a single-controller configuration, this will improve performance, but will expose the data to loss if a system crash or power failure occurs before the data in the cache is written to disk. To avoid this potential loss of data, use of an uninterruptable power supply (UPS) is recommended.

In dual-active controller systems, the data for write-back system drives is copied to the cache of the partner controller before the write completion is reported to the host initiator. During the time between writing the data to the partner and writing the data to the disk, the system is exposed to possible data loss should a power failure or system crash occur. An uninterruptable power supply is recommended.

RAID Levels and Availability

Refer to the following table for a summary of the different RAID levels offered by the Mylex controllers and the advantages and disadvantages of the RAID levels as they apply to availability.

RAID Levels and Availability

RAID Level	Availability (Fault tolerance)
0	No fault tolerance. Data is striped across a set of multiple disk drives. If a disk drive in the set ceases to function, all data contained on the set of disk drives is lost. This configuration is not recommended if fault tolerance is needed.
1	Mirrored fault tolerance. Data is written to one disk drive, and then the same data is written to another disk drive. If either disk drive fails, the other one in the pair is automatically used to store and retrieve the data.

RAID Levels and Availability (Continued)

RAID Level	Availability (Fault tolerance)
3 and 5	Striped fault tolerance. Data and parity are striped across a set of multiple (at least three) disk drives. If one of the disk drives fail, the data (or parity) information from the failed disk drive is computed from the information on the remaining disk drives.
0+1	Mirrored and striped fault tolerance. Mirror images of the data are striped across sets of disk drives
JBOD	"Just a Bunch of Disks." This configuration offers no redundancy and is not recommended for applications requiring fault tolerance.

Configuring for Maximum Performance

The relative performance advantages of each of the RAID levels is presented in the following table.

RAID Level Maximum Performance

RAID Level	Access profile characteristics
0	Excellent for all types of I/O activity
1	Excellent for write-intensive applications
3 and 5	Excellent for sequential or random reads and sequential writes
0+1	Excellent for write-intensive applications
JBOD	Mimics normal, individual disk drive performance characteristics

➢ See also "Array" on page 111, "Automatic Rebuild" on page 10, "Configuration on Disk (COD)" on page 18, "Configuration Strategies" on page 94, "Controller Specifications" on page 1, "Critical" on page 113, "Drive Distribution" on page 116, "Drive Packs / Groups" on page 29, "Drive Sizing" on page 30, "Dual-Active Controller Configuration" on page 32, "Fault Tolerance" on page 118, "Logical Device States" on page 124, "Mirroring" on page 129, "Offline" on page 130, "Online" on page 131, "Online Critical" on page 131, "Online Critical and Rebuilding" on page 132, "Online Optimal" on page 132, "Parity" on page 133, "RAID Guidelines" on page 97, "RAID Levels" on page 137, "Redundant Array" on page 145, "Spanning Drive Packs" on page 62, "Striping" on page 147, "System Drives" on page 149, "Write-Back Cache" on page 151, and "Write-Through Cache" on page 151.

Configuration Strategies

All Mylex controllers use RAID technology to provide protection against disk drive failures and improved performance, since disk drives are typically the weakest link in the data storage chain. Additionally, the controllers support configurations that employ redundant controllers and multiple hosts for improved fault tolerance and high performance data throughput. The user needs to determine requirements of the system prior to installing any of the individual components. A dual-active controller configuration provides protection in the event of a controller failure. A multiple host configuration provides protection in the event of a server or HBA failure. The user needs to determine termination and cabling requirements based upon the need to hot swap disk drives and hot plug controllers.

➢ See also "Alternate Path Software" on page 7, "CAP Strategy for Selecting a RAID Level" on page 92, "Controller Parameters" on page 69, "Controller Specifications" on page 1, "Drive Distribution" on page 116, "Dual-Active Controller Configuration" on page 32, "Failed Disk Drive" on page 118, "Fault Tolerance" on page 118, "Host Bus Adapter (HBA)" on page 120, "Hot Plugging" on page 121, "Hot Swapping" on page 122, "Hub" on page 122, "Microsoft

Cluster Server Configurations (MSCS)" on page 47, "RAID Guidelines" on page 97, "RAID Levels" on page 137, "Redundant Array" on page 145, "SCSI Cabling" on page 99, "SCSI Termination" on page 100, and "Storage Area Network (SAN)" on page 147.

Firmware Versions

ATTENTION: **Mylex recommends that you complete a full backup of your data before changing firmware levels, especially when moving back to a previous version.**

ATTENTION: **Downgrading firmware from 7.0 is not recommended. You need to backup all system data, clear the configuration, download the previous firmware version, recreate the configuration, and then replace the data.**

ATTENTION: **All SANmapping assignments created using firmware version 8.0 are lost if the firmware is downgraded to a previous firmware version.**

Firmware versions can be upgraded as needed. Refer to the following table for special instructions. RAIDfx (firmware versions prior to 7.0) or SAM/GAM can complete the firmware download process. Refer to the appropriate user guide for complete instructions.

Mylex does not recommend upgrading directly from firmware version 5.0 or 5.1 to firmware version 7.0. Upgrading to firmware version 7.0 can be performed using either of the following procedures:

- Upgrade the firmware to intermediate versions, following the recommendations in the Firmware Upgrade table. For example, upgrade from 5.0 or 5.1 to 5.4 or 6.1 then to 7.0, or
- Completely backup all data, delete the current configuration, upgrade the firmware to version 7.0, recreate the configuration, then restore the saved data to the new configuration.

Firmware Upgrade

Product	From FW Version	To FW Version	Notes
DACSX	3.2	4.2	Use normal download procedure.
	3.3	4.2	Use normal download procedure.
	4.2		This is the last release for DACSX. There is no upgrade available.
DACSS	4.4		The DACSS supports firmware version 4.4 only. There is no upgrade available.
DACSF	4.0 (Simplex)	4.1 (Dual-Active)	Use normal download procedure. Refer to "Upgrading From Simplex to Dual-Active Configurations" on page 103.
	4.1	4.2	Memory upgrade required. Install SIMMs in locations U35 and U36. Set hard loop IDs.[a]
	4.2		This is the last release for DACSF. There is no upgrade available.

Firmware Upgrade (Continued)

Product	From FW Version	To FW Version	Notes
DACFL	4.2	5.4	Use normal download procedure.
	4.2 or 5.4	6.3	Use normal download procedure.This is the last release for DACFL.
	6.3		This is the last release for DACFL. There is no upgrade available.
DACFF (FF)	5.0 (Simplex)	5.1 (Dual-Active)	Use normal download procedure. Refer to "Upgrading From Simplex to Dual-Active Configurations" on page 103.
	5.1	5.4	Use normal download procedure.
	5.1 or 5.4	6.1	Use normal download procedure.
	6.1	7.2	Use normal download procedure.
	7.2		This is the last release for DACFF (FF). There is no upgrade available.
FF2	6.1	7.2	Use normal download procedure.
	7.2	7.7	Use normal download procedure.
	7.7		This is the last release for FF2. There is no upgrade available.
FFx	6.0 (Simplex)	6.1 (Dual-Active)	Use normal download procedure. Refer to "Upgrading From Simplex to Dual-Active Configurations" on page 103.
	6.1	7.0	Use normal download procedure.
	7.0	7.4	Use normal download procedure.
	7.4	7.7	Use normal download procedure.
	7.7		There is no upgrade available at this time.
FFx2	8.0		There is no upgrade available at this time.

a. If you are operating under Windows NT, and are upgrading firmware on a DACSF, you must have hard loop IDs enabled prior to upgrading from 4.1 to 4.2. The alternative is to restart the host simultaneously with the controller restart. On restart, NT goes through a device discovery process and can properly find devices that have changed loop IDs.

➢ See also "Configuration on Disk (COD)" on page 18, "Controller Specifications" on page 1, "Global Array Manager (GAM)" on page 45, "Hard Loop ID" on page 120, "RAIDfx" on page 53, "SANArray Manager™ (SAM)" on page 56, and "Upgrading From Simplex to Dual-Active Configurations" on page 103.

Mode Page Parameters

Mylex recommends that certain mode page parameters be set to specific values as shown in the following table to ensure proper operations.

Recommended Mode Page Settings

Mode Page	Fields	Value	Bit/Byte
Page 01	AWRE	1	Bit 7 of Byte 2
Page 01	ARRE	1	Bit 6 of Byte 2
Page 08	WCE	0	Bit 2 of Byte 2
Page 19	RHA/DSA	48	Byte 3

Use of disk drives with Write Cache Enabled can promote a condition in which data loss can occur. Disk drives can report a deferred error when Write Cache is enabled. Firmware cannot recover from deferred errors and will immediately mark the disk drive as failed, resulting in possible data loss.

Mylex firmware requires hard loop IDs. If the Required Hard Address bit is not set appropriately, the disk drive address may change resulting in possible data loss.

ATTENTION: **To avoid data corruption, make sure Drive Mode Page 8 (Caching Page) specifies that the disk drives are not Write Cache Enabled and Drive Mode Page 19 specifies Required Hard Addressing.**

➢ See also "Default Mode Page Settings" on page 26, "Disk Media Error Management" on page 26 and "Hard Loop ID" on page 120.

RAID Guidelines

RAID stands for Redundant Array of Independent Disks. In a RAID system multiple disk drives are grouped into arrays. Each array is configured as a single system drive consisting of one or more disk drives.

A small, but important set of guidelines should be followed when connecting devices and configuring them to work with a Mylex controller. They are:

- Distribute the disk drives equally among all the drive channels on the controller. This will result in better performance. Mylex controllers have between two and six drive channels, depending on the model.
- A maximum of 8 or 16 (firmware version 7.0 or greater) devices can make up a drive pack.
- A drive pack can contain devices on any drive channel.
- If configuring for an online spare disk drive, make sure that the spare disk drive size is greater than or equal to the capacity of the largest disk drive in *all* the redundant drive packs.
- When replacing a failed disk drive, make sure that the replacement disk drive size is greater than or equal to the size of the failed disk drive in the affected drive pack.

➢ See also "Automatic Rebuild" on page 10, "CAP Strategy for Selecting a RAID Level" on page 92, "Controller Specifications" on page 1, "Drive Channel" on page 27, "Drive Packs / Groups" on page 29, "Failed Disk Drive" on page 118, "Hot Spare" on page 122, "Physical Device" on page 133, "Redundant Array" on page 145, "Replacement Drives" on page 145, "SCSI Cabling" on page 99, "SCSI Termination" on page 100.

Replacing a Failed Controller in Existing Duplex Systems

SCOPE: The Automatic Firmware Flash feature implemented in firmware version 8.0 ensures that dual-active controllers operate with the same firmware version, firmware type, and controller parameter settings.

To replace a failed controller cold (power off):

ATTENTION: **Make sure the replacement controller meets the requirements described under "Dual-Active Requirements" on page 33.**

ATTENTION: **If using firmware version 8.0 or greater, the Automatic Firmware Flash feature ensures that dual-active controllers operate with the same firmware version. One controller will flash firmware to the other controller. This firmware may not be the version used by the surviving controller. The correct**

version may need to be flashed to both controller after C-C nexus is established. Refer to “Automatic Firmware Flash” on page 7 for more information.

1. Power off the controller chassis.
2. Remove the failed controller.
3. Insert the replacement controller. The replacement controller Force Simplex parameter can be set to DIS or ENA. (This parameter is not supported by FFx and FFx2 controllers.)
4. Power on the controller.
5. Restart the host system.

The controllers accept the configuration on disk (COD) from the existing disk drives. The controllers may reset themselves during the COD update process.

To replace a failed controller hot (power on):

ATTENTION: **This may result in disk drive failure and possible data loss.**

ATTENTION: **Make sure the replacement controller meets the requirements described under “Dual-Active Requirements” on page 33. If the replacement controller does not match the surviving controller, the replacement controller is held in reset.**

ATTENTION: **If using firmware version 8.0 or greater, the Automatic Firmware Flash feature ensures that dual-active controllers operate with the same firmware version. One controller will flash firmware to the other controller. This firmware may not be the version used by the surviving controller. The correct version may need to be flashed to both controller after C-C nexus is established. Refer to “Automatic Firmware Flash” on page 7 for more information.**

1. Leave power on.
2. Remove the failed controller.
3. Install the replacement controller.

If the *Automatic Failback or Autorestore* parameter is enabled, the surviving controller detects the replacement controller and begins the failback process. If the *Automatic Failback or Autorestore* parameter is not enabled, a *Relinquish Partner* command must be issued. If the *Force Simplex* parameter is enabled, it is overridden and both controllers operate in dual-active controller mode. No further operator intervention is required.

➢ See also “Automatic Failback or Autorestore” on page 70, “Automatic Firmware Flash” on page 7, “Cold Swap” on page 112, “Configuration on Disk (COD)” on page 18, “Controller Parameters” on page 69, “Drive State Management” on page 31, “Dual-Active Controller Configuration” on page 32, “Duplex” on page 116, “Failed Controller” on page 118, “Failover and Failback” on page 35, “Force Simplex” on page 78, “Hot Plugging” on page 121, “Replacement Controller” on page 145, “Replacement Controller Behavior with Forced

Simplex Enabled” on page 53, “Reset Controllers” on page 54, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, “Surviving Controller” on page 148, and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

Replacing a Failed Controller in Simplex Systems

1. Power off the controller.
2. Remove the failed controller.
3. Insert the replacement controller.
4. Power on the controller.

The controller accepts the configuration on disk (COD) from the existing disk drives. The controller may reset itself during the COD update process.

➢ See also “Configuration on Disk (COD)” on page 18, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Replacement Controller” on page 145, “Reset Controllers” on page 54, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, and “Simplex” on page 60.

SCSI Cabling

SCOPE: SCSI Cabling requirements must be considered for the DACSX, DACSS, DACSF, and DACFL controllers.

When planning the cable requirements, be aware of the SCSI rules for maximum cable lengths to avoid performance problems. Maximum cable lengths are shown in following table.

Supported SCSI Formats and Cable Lengths

SCSI Trade Association Terms	Bus Speed, MB/Sec, Max.	Bus Width, bits	Max. Bus Length, Meters[a]			Max. Device Support (Including Controllers)
			Single-ended	Differential (HVD)	LVD	
SCSI-1[b]	5	8	6	25	[c]	8
Fast SCSI[b]	10	8	3	25	[c]	8
Fast Wide SCSI	20	16	3	25	[c]	16
Ultra SCSI[b]	20	8	1.5	25	[c]	8
Ultra SCSI[b]	20	8	3	-	-	4
Wide Ultra SCSI	40	16	-	25	[c]	16
Wide Ultra SCSI	40	16	1.5	-	-	8
Wide Ultra SCSI	40	16	3	-	-	4
Ultra2 SCSI[b, d]	40	8	[d]	25	12	8
Wide Ultra2 SCSI[d]	80	16	[d]	25	12	16
Ultra3 SCSI or Ultra 160[e]	160	16	[d]	[f]	12	16
Ultra 320	320	16	[d]	[f]	12	16

a. The listed maximum bus lengths may be exceeded in point-to-point and engineered applications.
b. Use of the word “Narrow,” preceding SCSI, Ultra SCSI, or Ultra2 SCSI is optional.
c. LVD was not defined in the original SCSI standards for this speed. If all devices on the bus support LVD, 12-meter operation is possible at this speed. If any device on the bus is single-ended only, the entire bus switches to single-ended mode and the distances in the single-ended column apply.
d. Single-ended is not defined for speeds beyond Ultra.
e. After Ultra2 all new speeds are wide only.
f. Differential (HVD) is not defined for speeds beyond Ultra2.

For the most recent information, go to the SCSI Trade Association web site at:

www.scsita.org

➢ See also “Channel” on page 112, “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Distribution” on page 116, “Dual-Active Controller Configuration” on page 32, “SCSI Termination” on page 100, and “Simplex” on page 60.

SCSI Termination

SCOPE: SCSI Termination requirements must be considered for the DACSX, DACSS, DACSF, and DACFL controllers.

Terminating a SCSI chain is accomplished by adding a terminator to each end of each SCSI bus or by terminating the devices closest to each of the two ends of each SCSI drive and host channel.

NOTE: It is better to terminate the ends of the SCSI bus itself than it is to terminate the end devices on the bus. This allows hot swapping of disk drives and hot plugging of controllers (in dual-active controller configurations) without affecting bus termination.

The following table describes termination guidelines for each configuration type.

Termination Guidelines

Configuration Type	Termination Settings
Simplex	If the controller is at one end of all of the disk SCSI busses, you may terminate the drive channels either by enabling termination on the distribution board or by supplying an external terminator. If you plan to hot swap the disk drives, do not enable termination on the disk drives; instead, terminate the far end of the disk SCSI busses with external terminators.
Dual-Active Controller	If the two controllers are at opposite ends of the disk SCSI busses, you may terminate the drive channels either by enabling termination on the distribution boards or by supplying external terminators. If either is at a location other than at the end of the disk SCSI busses, you must disable termination on that controller’s distribution board. If either controller is not used for disk bus termination, external terminators must be used at the end of the bus not terminated by a controller. If you plan to hot swap the controllers, do not enable termination on the distribution board; instead, use external terminators. Do not enable termination on the disk drives if you plan to hot swap the disk drives.

SCSI Termination for the DACSX Only

Only the DACSX controller has single-ended, drive channel SCSI bus terminations available. The other controllers do not provide on-board termination; however, the distribution board does contain SCSI bus terminators for the drive channels which may all be enabled or disabled. Refer to the *DACSX Installation Guide* for more information.

➢ See also “Channel” on page 112, “Configuration Strategies” on page 94, “DACSX” on page 25, “Dual-Active Controller Configuration” on page 32, “Hot Plugging” on page 121, “Hot Swapping” on page 122, “SCSI Cabling” on page 99, and “Simplex” on page 60.

Upgrading a Configuration Created with RAIDfx

ATTENTION: **Mylex recommends a complete data backup before performing any hardware or firmware upgrade.**

Previously, Mylex controllers supported an on-line configuration tool known as RAIDfx. With the release of firmware version 7.0, RAIDfx is no longer supported. RAIDfx does not support many of the features implemented with firmware version 7.0 or greater. RAIDfx only supports 16 physical devices per system, for example, 14 disk drives and two controllers.

Configurations created with RAIDfx using firmware version 6.0 and earlier, in most cases, can be upgraded to firmware version 7.0 and greater. Use the following procedure to upgrade an existing configuration created with RAIDfx to firmware version 7.0 or greater. Additional disk drives can be added following the firmware upgrade. Use SAM/GAM or the embedded configuration utility to edit, monitor, and create configurations.

ATTENTION: **The drive size truncation feature in firmware version 7.0 may effect the amount of available space on the physical disk drives.**

ATTENTION: **Known limitations exist when upgrading firmware version 6.x to 7.0. Refer to the firmware 7.0 release notes for a complete description.**

ATTENTION: **RAIDfx no longer functions with firmware version 7.0 or greater.**

1. Perform a complete data backup.
2. Start RAIDfx and verify the physical configuration to ensure a 16 MB area is reserved for controller information.
   a. Determine the physical size and usable space for the disk drives in the system using RAIDfx.
   b. From the Configuration menu select Device Information. (RAIDfx->Configuration ->View/Update->Define Packs->Device Information.)
   c. Divide the number of blocks returned in the physical size field by (1024 x 1024) and compare this value with the value returned in the usable space field.
   d. Subtract the usable space from the physical size. If 16 MB or more is available, there is adequate space to upgrade firmware.
3. Install SAM/GAM version 3.10 or greater per instructions in the Global Array Manager Installation Guide and User Manual and/or the SANArray Manager Client Software Installation Guide and User Manual.
4. Upgrade the firmware to version 7.0 or greater using SAM/GAM. Refer to the Firmware Upgrade table for the appropriate upgrade sequence. Refer to the *Global Array Manager Installation Guide and User Manual* or the *SANArray Manager Client Software for Mylex External Disk Array Controllers Installation Guide and User Manual* for instructions on downloading firmware.
5. Use SAM/GAM version 3.10 or greater or the embedded configuration utility to edit or monitor the configuration. Additional disk drives can be added at this time.

➢ See also “Configuration on Disk (COD)” on page 18, “Configuration Tools and Utilities” on page 21, “Controller Specifications” on page 1, “Drive Sizing” on page 30, “Embedded Configuration Utility” on page 35, “Firmware Versions” on page 95, “Global Array Manager (GAM)” on page 45, “Product Supported Features” on page 2, “RAIDfx” on page 53, “SANArray Manager™ (SAM)” on page 56, and “VT100” on page 65.

Upgrading a DACFF (FF) System to an FF2 System

The FF2 controller is very similar to the DACFF (FF) controller; however, the two controllers are not interchangeable. A dual-active controller configuration must consist of two identical controllers.

NOTE: Mylex recommends a complete data backup before performing any hardware or firmware upgrade.

ATTENTION: **Do not install an FF2 controller in a dual-active configuration with a DACFF (FF) controller. The current configuration and data may be lost.**

Upgrade from Simplex DACFF (FF) to Simplex FF2

1. Power off the controller.
2. Remove the DACFF (FF) controller.
3. Insert the FF2 controller.
4. Power on the controller.

The controller accepts the configuration on disk (COD) from the existing disk drives. The controller may reset itself during the COD update process.

Upgrade from Dual-Active DACFF (FF) to Dual-Active FF2

ATTENTION: **Both controllers in a dual-active controller configuration must be identical and meet the requirements described under “Dual-Active Requirements” on page 33. Both DACFF (FF) controllers must be replaced with FF2 controllers. The two models are not compatible or interchangeable.**

1. Power off both controllers.
2. Remove both DACFF (FF) controllers.
3. Insert both FF2 controllers.

ATTENTION: **Both controllers must be powered on at the same time in order to ensure configuration and data integrity.**

4. Power on the controllers simultaneously.

The controllers accept the configuration on disk (COD) from the existing disk drives. The controllers may reset themselves during the COD update process.

Upgrade from Simplex DACFF (FF) to Dual-Active FF2

This is a two-step procedure involving upgrade of the simplex system followed by the upgrade to a dual-active controller system.

ATTENTION: **Both controllers in a dual-active controller configuration must be identical and meet the requirements described under "Dual-Active Requirements" on page 33.**

1. Power off the existing DACFF (FF) controller.
2. Remove the DACFF (FF) controller.
3. Insert one FF2 controller.
4. Power on the controller.

The controller accepts the configuration on disk (COD) from the existing disk drives. The controller may reset itself during the COD update process.

This upgrades the system to a simplex FF2 controller, now it is possible to upgrade to a dual-active FF2 controller system. Follow the procedures described under "Upgrading From Simplex to Dual-Active Configurations" on page 103.

➢ See also "Automatic Failback or Autorestore" on page 70, "Controller Parameters" on page 69, "Controller Specifications" on page 1, "DACFF (FF)" on page 24, "Dual-Active Controller Configuration" on page 32, "FF2" on page 42, "Force Simplex" on page 78, "LUN Mapping" on page 47, "Programmable LUN Mapping" on page 52, "Replacement Controller" on page 145, "Replacing a Failed Controller in Existing Duplex Systems" on page 97, "Replacing a Failed Controller in Simplex Systems" on page 99, "SANmapping™" on page 56, "System Drive Affinity" on page 64, and "Upgrading From Simplex to Dual-Active Configurations" on page 103

Upgrading From Simplex to Dual-Active Configurations

ATTENTION: **This procedure may result in disk drive failure and possible data loss.**

ATTENTION: **Make sure both controllers meet the requirements described under "Dual-Active Requirements" on page 33.**

To Perform Hot-Plug Upgrade (Mylex recommended)

NOTE: *Automatic Failback or Autorestore* must be disabled for this scenario to function properly.

ATTENTION: **Both controllers must have the same firmware version and dual-active firmware type to perform a hot-plug upgrade.**

1. Leave power to the existing controller on.
2. Disable *Force Simplex* on the existing controller.
3. Reset the existing controller.

4. After the existing controller reaches STARTUP COMPLETE, install the additional controller. The new controller starts up and accepts the configuration on disk from the existing controller.
5. Change any LUN mapping assignments to take advantage of the additional controller/host ports available on the new controller. This step is not necessary, but performance can be optimized by load balancing the LUN mapping between the controller/host ports.
6. Reset the controllers and restart the host. LUN mapping changes are not accepted by the host without restarting.

To Perform Cold Upgrade (power off)

ATTENTION: **This may result in disk drive failure and possible data loss.**

1. Leave power to the existing controller on.
2. Download a dual-active firmware type (4,5,6, or 7), if necessary.
3. Disable *Force Simplex* on the existing controller.
4. Power off the existing controller.
5. Install the additional controller and power up.
6. Download a dual-active firmware type (4,5,6, or 7), if necessary. The firmware version and type must match the existing controller.
7. Disable *Force Simplex* on the existing controller.
8. Power off the additional controller.
9. Install the existing controller and power both controllers simultaneously.
10. Change any LUN mapping assignments to take advantage of the additional controller/host ports now available. This step is not necessary, but performance can be optimized by load balancing the LUN mapping between the controller/host ports.
11. Reset the controllers and restart the host. LUN mapping changes are not accepted by the host without restarting.

➢ See also “Automatic Failback or Autorestore” on page 70, “Configuration on Disk (COD)” on page 18, “Controller/Host Port” on page 112, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Failed Disk Drive” on page 118, “Failover and Failback” on page 35, “Firmware Header Information” on page 44, “Force Simplex” on page 78, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “Reset Controllers” on page 54, “SANmapping™” on page 56, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, “Simplex” on page 60, and “System Drive Affinity” on page 64.

TERMINOLOGY DEFINITIONS

This chapter provides an alphabetic listing of commonly used terminology and associated definitions. The information applies to all Mylex controllers, unless otherwise noted.

Acronym List

The following table presents a list of acronyms used throughout this document.

Acronyms Used in the Encyclopedia

Acronym	Term	Definition or Usage
AEMI	Array Enclosure Management	Refer to “AEMI” on page 6.
ac	alternating current	Refer to “Battery Backup Unit (BBU)” on page 12, “Conservative Cache Mode” on page 22, “Environmental Device” on page 117, “SAF-TE Data for UPS Support” on page 84, and “SES” on page 58.
AL_PA	Arbitrated Loop Physical Address	Refer to “Arbitrated Loop Physical Address (AL_PA)” on page 111.
ARRE	Automatic Read Reallocation Enable	Refer to “Mode Page Parameters” on page 96.
ASIC	Application Specific Integrated Circuit	Refer to “Battery Backup Unit (BBU)” on page 12.
AUTOSPIN	Automatic disk drive spin up	Refer to “Disk Startup Mode” on page 75.
AWRE	Automatic Write Reallocation Enable	Refer to “Mode Page Parameters” on page 96.
BAT	Battery	Refer to “SAF-TE Data for UPS Support” on page 84.
BBU	Battery Backup Unit	Refer to “Battery Backup Unit (BBU)” on page 12.
C0P0	Controller 0, Port 0	Refer to “Controller/Host Port” on page 112.
CAP	Capacity, Availability, Performance	Refer to “CAP Strategy for Selecting a RAID Level” on page 92.
C-C Nexus	Controller-Controller Nexus	Refer to “Controller-Controller Nexus (C-C Nexus)” on page 23.
CDB	Command Descriptor Block	Refer to “SAF-TE” on page 55.

Acronyms Used in the Encyclopedia

Acronym	Term	Definition or Usage
COD	Configuration on Disk	Refer to “Configuration on Disk (COD)” on page 18.
CPU	Central Processing Unit	Refer to “Rebuild and Check Consistency Rate” on page 84.
CRC	Cyclic Redundancy Check	Refer to “Cyclic Redundancy Check (CRC)” on page 113.
DACFF	Disk Array Controller Fibre-Fibre	Refer to “DACFF (FF)” on page 24.
DACFL	Disk Array Controller Fibre-LVD	Refer to “DACFL” on page 24.
DACSF	Disk Array Controller SCSI-Fibre	Refer to “DACSF” on page 24.
DACSS	Disk Array Controller SCSI-SCSI	Refer to “DACSS” on page 25.
DACSX	Disk Array Controller SCSI-SCSI	Refer to “DACSX” on page 25.
dc	direct current	Refer to “SES” on page 58.
DIMM	Dual In-line Memory Module	Refer to “Battery Backup Unit (BBU)” on page 12.
DIS	Disable	Refer to “Controller Parameters” on page 69.
DSA	Disable Soft Addressing	Refer to “Default Mode Page Settings” on page 26.
ECC	Error Correcting Code	Refer to “Controller Fault Indication” on page 23 and “Disk Media Error Management” on page 26.
ENA	Enable	Refer to “Controller Parameters” on page 69.
FC	Fibre Channel	“Fibre Channel” on page 43
FC-AL	Fibre Channel Arbitrated Loop	Refer to “Fibre Channel Arbitrated Loop (FC-AL)” on page 43.
FF2	Fibre-Fibre Dual Loop Controller	Refer to “FF2” on page 42.
FFx	SFF Fibre-Fibre Controller	Refer to “FFx” on page 42.
FFx2	SFF Fibre-Fibre Controller (2nd generation)	Refer to “FFx2” on page 43.
F0/FB	Failover/Failback	Refer to “Drive Channel Failover and Failback” on page 28 and “Failover and Failback” on page 35.
FW	Firmware	Refer to “Firmware Versions” on page 95.
GAM	Global Array Manager	Refer to “Global Array Manager (GAM)” on page 45.
GB	Gigabyte	Refer to “Gigabyte (GB)” on page 120.
Gbps	Gigabit per second	Refer to “FFx2” on page 43.
GUI	Graphic User Interface	Refer to “Global Array Manager (GAM)” on page 45 or “SANArray Manager™ (SAM)” on page 56.
HBA	Host Bus Adapter	Refer to “Host Bus Adapter (HBA)” on page 120.
HVD	High Voltage Differential	Refer to “SCSI Cabling” on page 99.
I/O	Input/Output	Refer to “I/O” on page 123.
IOP	I/O Process	Refer to “Maximum IOPs” on page 80.
JBOD	Just a Bunch of Disks	Refer to “JBOD” on page 123.

Acronyms Used in the Encyclopedia

Acronym	Term	Definition or Usage
KB	Kilobyte	Refer to “Kilobyte (KB)” on page 123.
L_Port	Loop Port	Refer to “Loop Port (L_Port)” on page 128.
LCD	Liquid Crystal Display	Refer to “Liquid Crystal Display (LCD)” on page 47.
LED	Light Emitting Diode	Refer to “Drive State Management” on page 31 or “Fault Management” on page 41.
LIFA	Loop Initialization Fabric Assigned	Refer to “Loop Initialization Fabric Assigned (LIFA)” on page 126.
LIHA	Loop Initialization Hard Assigned	Refer to “Loop Initialization Hard Assigned (LIHA)” on page 126.
LILP	Loop Initialization Loop Position	Refer to “Loop Initialization Loop Position (LILP)” on page 126.
LIM	Loop Initialization Master	Refer to “Loop Initialization Master (LIM)” on page 126.
LIP	Loop Initialization Primitive	Refer to “Loop Initialization Primitive (LIP)” on page 126.
LIPA	Loop Initialization Previously Acquired	Refer to “Loop Initialization Previously Acquired (LIPA)” on page 127.
LIRP	Loop Initialization Report Position	Refer to “Loop Initialization Report Position (LIRP)” on page 127.
LISA	Loop Initialization Soft Assigned	Refer to “Loop Initialization Soft Assigned (LISA)” on page 127.
LISM	Loop Initialization Select Master	Refer to “Loop Initialization Select Master (LISM)” on page 127.
Loop ID	Loop Address	Refer to “Loop Address (Loop ID)” on page 125.
LPB	Loop Port Bypass	Refer to “Loop Port Bypass (LPB)” on page 128.
LPE	Loop Port Enable	Refer to “Loop Port Enable (LPE)” on page 128.
LRC	Loop Redundancy Circuit or Loop Resiliency Circuit	Refer to “Dual-Active Controller Configuration” on page 32 and “Loop Redundancy Circuit (LRC)” on page 128.
LUN	Logical Unit Number	Refer to “Logical Unit Number (LUN)” on page 125.
LVD	Low Voltage Differential	Refer to “SCSI Cabling” on page 99.
MAC	Media Access Control	Refer to “Enclosure World Wide Name” on page 66.
MB	Megabyte	Refer to “Megabyte (MB)” on page 129.
MORE	Mylex Online RAID Expansion	Refer to “Mylex Online RAID Expansion (MORE)” on page 49.
MSCS	Microsoft Cluster Server	Refer to “Microsoft Cluster Server Configurations (MSCS)” on page 47.
MTID	Multiple Target ID	Refer to “Multiple Target ID (MTID)” on page 48.
NiCd	Nickel Cadmium	Refer to “Battery Backup Unit (BBU)” on page 12.

Acronyms Used in the Encyclopedia

Acronym	Term	Definition or Usage
NiMH	Nickel Metal Hydride	Refer to “Battery Backup Unit (BBU)” on page 12.
N_Port	Node Port	Refer to “Node Port (N_Port)” on page 130.
NL_Port	Node Loop Port	Refer to “Node Loop Port (NL_Port)” on page 129.
NVRAM	Non-Volatile RAM	Refer to “NVRAM” on page 130.
PBC	Port Bypass Circuit	Refer to “Port Bypass Circuit (PBC)” on page 135.
PBR	Protected Boot ROM	Refer to “Protected Boot ROM (PBR)” on page 53.
PCI	Peripheral Component Interconnect	Refer to “PCI Latency Control” on page 82.
PFA	Predictive Failure Analysis	Refer to “Predictive Failure Analysis™ (PFA)” on page 52.
PLPB	Prevent Loop Port Bypass	“Default Mode Page Settings” on page 26.
PROM	Programmable Read Only Memory	Refer to “Protected Boot ROM (PBR)” on page 53.
PWRSPIN	Spin up on power	Refer to “Disk Startup Mode” on page 75.
RAID	Redundant Array of Independent Disks	Refer to “CAP Strategy for Selecting a RAID Level” on page 92, “RAID Guidelines” on page 97, or “RAID Levels” on page 137.
RAM	Random Access Memory	Refer to “NVRAM” on page 130.
RDA	Reserved Disk Area	Refer to “Debug Dump” on page 25.
RHA	Require Hard Address	“Default Mode Page Settings” on page 26.
ROM	REad Only Memory	Refer to “Protected Boot ROM (PBR)” on page 53.
SAF-TE	SCSI Accessed Fault-Tolerant Enclosure	Refer to “SAF-TE” on page 55.
SAM	SANArray Manager™	Refer to “SANArray Manager™ (SAM)” on page 56.
SAN	Storage Area Network	Refer to “Storage Area Network (SAN)” on page 147.
SCO	Santa Cruz Operation (Unix)	Refer to “Global Array Manager (GAM)” on page 45 or “SANArray Manager™ (SAM)” on page 56.
SCSI	Small Computer System Interface	Refer to “SCSI” on page 145.
SES	SCSI-3 Enclosure Services	Refer to “SES” on page 58.
SFF	Small Form Factor	Refer to “FFx” on page 42.
SIOP	SCSI I/O Processor	Refer to “SIOP” on page 146.
SLP	Serial Line Protocol	Refer to “Serial Port Usage” on page 86.
SMART	Self-Monitoring, Analysis and Reporting Technology	Refer to “Self-Monitoring, Analysis and Reporting Technology (SMART)” on page 60.

Acronyms Used in the Encyclopedia

Acronym	Term	Definition or Usage
SPOF	Single Point of Failure	Refer to “Single Point of Failure (SPOF)” on page 61.
SRA	Super Read Ahead	Refer to “Super Read Ahead” on page 87.
SSU	Start/Stop Unit	Refer to “Disk Startup Mode” on page 75.
TB	Terabyte	Refer to “Terabyte (TB)” on page 150.
TID	Target ID	Refer to “Target ID (TID)” on page 149
TUR	Test Unit Ready	Refer to “Vendor Unique TUR” on page 88.
UPS	Uninterruptable Power Supply	Refer to “SAF-TE Data for UPS Support” on page 84.
WB	Write-Back Cache	Refer to “Write-Back Cache” on page 151.
WCE	Write Cache Enabled	Refer to “Mode Page Parameters” on page 96.
WT	Write-Through Cache	Refer to “Write-Through Cache” on page 151.
WWN	World Wide Name	Refer to “World Wide Name (WWN) Assignments” on page 65 or “World Wide Name (WWN)” on page 151.
XOR	Exclusive OR	Refer to “Rotated XOR Redundancy” on page 145.

Active Controller

A controller that functions properly and handles host I/O requests. In a dual-active controller configuration one or both controllers may be active.

➢ See also “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, and “Failover and Failback” on page 35.

Active Fibre Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

A Fibre Channel port connected to an arbitrated loop and capable of arbitrating the loop. The port may be opened by an initiator on the loop and process I/O requests. An active fibre port is also referred to as a *participating* port.

➢ See also “Failover and Failback” on page 35, “Failover Port” on page 118, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Inactive Fibre Port” on page 123, “Port” on page 135, “Primary Port” on page 137, and “Simplex Fibre Channel System” on page 146.

Address Identifier

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

A Fibre Channel 24-bit value used to identify the source or destination of a frame.

➢ See also“Fibre Channel” on page 43 and “Frame” on page 119.

Alternate Address

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

Dual-ported drives can be connected to two drive channels. When the controller firmware looks for drives, one of the drive ports is assigned a *preferred* (channel-target ID) address; the other is assigned an *alternate* address. When a drive channel fails, if that channel is associated with the preferred address, the drive is accessed through its alternate address. When the failed channel is restored, the drive is again accessed through its preferred address.

➢ See also “Drive Channel” on page 27, “Drive Channel Failover and Failback” on page 28, “Dual-Ported Drive Support” on page 34, “Loop Address (Loop ID)” on page 125, “Preferred Address” on page 135, and “Target ID (TID)” on page 149

Alternate Channel

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2 and FFx2 controllers.

At drive discovery time, a dual-ported drive may be found on two different controller drive channels. One of these channels is assigned the drive's preferred channel; the other is the drive’s alternate channel. The controller sends all commands to the drive on its preferred channel unless the preferred channel is not available (the channel has failed).

➢ See also “Channel” on page 112, “Drive Channel Failover and Failback” on page 28, “Dual-Ported Drive Support” on page 34, and “Preferred Channel” on page 136.

Alternate Path

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

In a dual-active configuration, one controller/host port combination serves as the primary data path and the secondary controller/host port combination serves as an alternate path for data transfer from a host to a LUN. If a failure occurs to the primary path, the alternate path assumes all activity for the failed path, maintaining access to the LUN.

➢ See also "Controller/Host Port" on page 112, "Dual-Active Controller Configuration" on page 32, "Host / Server" on page 121, "Logical Unit Number (LUN)" on page 125, "Primary Path" on page 137, and "SANmapping™" on page 56.

Arbitrated Loop Physical Address (AL_PA)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A unique, one-byte value assigned during loop initialization to each NL_Port on a Fibre Channel loop. There is a one-to-one mapping between loop IDs and AL_PAs.

➢ See also "Fibre Channel" on page 43, "Loop Address (Loop ID)" on page 125, "Node Loop Port (NL_Port)" on page 129, and "Node Port (N_Port)" on page 130.

Array

An array is created from multiple physical disk drives configured to behave as a single, independent disk drive.

➢ See also "CAP Strategy for Selecting a RAID Level" on page 92, "RAID Guidelines" on page 97, "RAID Levels" on page 137, and "System Drives" on page 149.

Asynchronous

Asynchronous I/O operations allow requests to overlap in time. Data can be accessed from different drives of an array at the same time.

➢ See also "Array" on page 111 and "I/O" on page 123.

Bandwidth

The maximum information carrying capacity of a system.

Bus

A physical connection between a host system and a disk array used to communicate data.

➢ See also "Array" on page 111, "Fibre Channel" on page 43, and "Host / Server" on page 121.

Cache

Cache is memory stored on the controller used to speed up data transfer to and from a disk drive or a partner controller.

➢ See also “Battery Backup Unit (BBU)” on page 12, “Cache Coherency” on page 17, “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, “Controller Parameters” on page 69, “Controller Specifications” on page 1, “Partner Controller” on page 133, “Standard Data Caching” on page 63, “Warm Start” on page 151, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Cache Flush

Cache flush is an operation where all unwritten data in a write-back cache are written to the target disk drives. This operation is necessary before powering down the system.

➢ See also “Battery Backup Unit (BBU)” on page 12, “Cache” on page 111, “Cache Coherency” on page 17, “Conservative Cache Mode” on page 22, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Partner Controller” on page 133, “Warm Start” on page 151, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151.

Channel

A path for the transfer of data and control information between the host or the physical disk drives and the controller. Mylex disk array controllers have one or two host channels and up to six drive channels. Depending on the type of controller and the firmware version, each drive channel can support from 14 up to 124 physical disk drives.

➢ See also “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Drive Distribution” on page 116, “Fibre Channel” on page 43, “Host Bus Adapter (HBA)” on page 120, “Host Channel” on page 121, “Physical Disk Drive” on page 135, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

Cold Swap

Replacing a physical device (disk drive or controller) while power is removed from the entire system.

➢ See also “Failed Disk Drive” on page 118, “Hot Plugging” on page 121, “Hot Swapping” on page 122, and “Replacing a Failed Controller in Existing Duplex Systems” on page 97.

Controller/Host Port

An identification address using the controller number and the host port number, for example C0P0 would be controller 0, host port 0. The controller/host port serves as the primary data path for data transfer from a host to a LUN.

➢ See also “Alternate Path” on page 110, “Drive Channel Failover and Failback” on page 28, “Failover and Failback” on page 35, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Host Channel” on page 121, “Host Port” on page 121, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, and “System Drive Affinity” on page 64.

Critical

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to online critical for a system drive or logical device. Refer to “Online Critical” on page 131 for more information.

The state of a system drive that has been configured with a redundant RAID level and one of its member disk drives is *not online*. Failure of another member disk drive will cause the system drive to go offline and will result in potential loss of data.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Dead” on page 113, “Device States” on page 114, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Logical Device States” on page 124, “Offline” on page 130, “Online” on page 131, “RAID Levels” on page 137, “Rebuild” on page 144, “Redundant Array” on page 145, “Replacement Drives” on page 145, and “System Drive State” on page 148.

Cyclic Redundancy Check (CRC)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

An error-correcting code used by Fibre Channel protocol.

➢ See also “Fibre Channel” on page 43 and “Frame” on page 119.

Data Redundancy

This technique stores extra data in addition to the original data to protect against disk failure. The redundant data is used to rebuild the original data in the event of disk failure. Data Redundancy is accomplished either through 100% mirroring of the data or through creation of parity data.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Configuration Strategies” on page 94, “Failed Disk Drive” on page 118 “Fault Tolerance” on page 118, “Mirroring” on page 129, “Parity” on page 133, “RAID Guidelines” on page 97, “RAID Levels” on page 137, “Redundant Array” on page 145, and “Striping” on page 147.

Data Transfer Capacity

The amount of data moved between devices, generally measured in Megabytes/sec. This is also referred to as data transfer rate.

➢ See also “Controller Parameters” on page 69 and “Transfer Speed” on page 87.

Dead

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to offline failed for both disk drives and system drives or physical devices and logical devices. Refer to “Offline Failed” on page 131 for more information.

The state of a physical disk drive that is not present, not powered on, or not operating properly and labeled DEAD by the controller.

➢ See also “Alarm Signal” on page 6, “Automatic Rebuild” on page 10, “Critical” on page 113, “Device States” on page 114, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Spare” on page 122, “Logical Device States” on page 124, “Online Spare” on page 132, “Physical Device States” on page 134, “Rebuild” on page 144, “Replacement Drives” on page 145, and “Standby” on page 147.

Device Access

The standard interface for host software is SCSI. Using the SCSI interface allows access to non-configured physical devices as well as logical devices. SCSI protocol is supported over Fibre Channel and parallel SCSI platforms.

➢ See also “Fibre Channel” on page 43, “Physical Device” on page 133, and “SCSI” on page 145.

Device States

SCOPE: The terms used by the firmware and configuration utilities have evolved through the various firmware versions and releases. The terms used by firmware versions 7.0 and greater are listed and the equivalent term for previous firmware versions are given in parenthesis.

Physical devices include physical disk drives and other devices attached to the controller through an enclosure.

Logical devices include configured system drives or logical drives using the physical disk drives attached to the controller.

Possible Physical Device states are:

- Unconfigured Offline (Offline or Unconfigured)—The device has been inserted, powered on into an unconfigured slot, or the configuration has just been cleared. The device is not part of a configuration.
- Commanded Offline (Offline)—The device has been commanded offline by the host. This physical device state is *not supported* at this time.
- Offline Failed (Dead)—The device has failed to operate properly.
- Offline Missing (Not supported prior to 7.0)—The device has been configured, but is not present at the time of system startup. This physical device state is *not supported* at this time.
- Offline Warm Spare (Not supported prior to 7.0)—This physical device state is *not supported* at this time.
- Online Optimal (Online)—The device is powered on, part of a configuration, and functioning normally.
- Online Critical (Not supported prior to 7.0)—The device is powered on but has sent PFA sense to the host, suggesting the device is not functioning normally and should be replaced. This physical device state is *not supported* at this time.
- Online Hot Spare (Standby or Hot Spare)—The device is part of a configuration and operating as a spare drive and will be used for automatic rebuild.
- Online Rebuild (Rebuild)—The device is powered on, part of a configuration, and in the process of being rebuilt.

Possible Logical Device states are:

- Offline Failed (Offline)—The logical device has suffered device failures that exceed the limit for the configured RAID type.

- Online Optimal (Online)—All devices associated with the configured logical device are powered on and functioning normally.
- Online Critical (Critical)—The logical device has suffered a device failure within the limit for the configured RAID type. If the logical device suffers another device failure, the state changes to offline failed.
- Online Critical and Rebuilding (Critical)—The logical device is in the process of rebuilding a failed device. If the logical device suffers another device failure, the state changes to offline failed.

➢ See also “Logical Device States” on page 124, “Offline Failed” on page 131, “Online Critical” on page 131, “Online Critical and Rebuilding” on page 132, “Online Optimal” on page 132, “Online Rebuild” on page 132, “Online Spare” on page 132, “Physical Device” on page 133, “Physical Device States” on page 134, and “Unconfigured Offline” on page 150.

Disk Loop

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel loop connecting multiple disk drives in an arbitrated loop topology.

➢ See also “Drive Channel Failover and Failback” on page 28 and “Fibre Channel” on page 43.

Disk Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The Fibre Channel interface between a disk loop and a disk drive.

➢ See also “Drive Channel Failover and Failback” on page 28 and “Fibre Channel” on page 43.

Disk Striping

The practice of dividing data into blocks and writing them across multiple disk drives for increased performance.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Data Redundancy” on page 113, “Initialize System Drive” on page 123, “Parity” on page 133, “RAID Levels” on page 137, “Redundant Array” on page 145, “Stripe Size” on page 64, and “Striping” on page 147.

Drive and Controller Addressing

SCOPE: The controller addressing requirements have evolved through the various firmware versions and releases. As technology has allowed for supporting greater number of devices, controller addressing has moved to accommodate for the greater number of devices.

Firmware version 3.2 first supported 16 SCSI IDs per drive channel, but reserved two IDs for the controllers.

Firmware version 5.0 first supported fibre loop topology and 16 drives per drive channel.

Firmware version 7.0 first supported 126 drives per drive channel, but reserved two IDs for the controllers.

In systems with SCSI disk busses, SCSI target ID 7 is reserved for controller 0, and in dual-active controller systems SCSI target ID 6 is reserved for controller 1. Disk drives and environmental devices may use any of the remaining SCSI target IDs; that is, 0-5 and 8-15 on dual-active controller systems or 0-6 and 8-15 on simplex systems.

In systems with fibre drive channels, hard addressing is required. Fibre loop ID 113 (71h) is reserved for controller 0, and in dual-active controller systems fibre loop ID 112 (70h) is reserved for controller 1. Disk drives may use any of the remaining fibre loop IDs; that is, 0-111 (0-6Fh) and 114-125 (72h-7Dh). A maximum of 16 disk drives per fibre drive channel is supported with firmware versions earlier than 7.0. A maximum of 124 disk drives per fibre drive channel is supported with firmware versions 7.0 and greater.

➢ See also “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Dual-Active Controller Configuration” on page 32, “Environmental Device” on page 117, “Fibre Channel” on page 43, “Hard Loop ID” on page 120, “Loop Address (Loop ID)” on page 125, “SCSI” on page 145, and “Target ID (TID)” on page 149.

Drive Distribution

Physical disk drives are connected to one or more of the drive channels on the Mylex controllers. When using Mylex controllers with two or more drive channels, it is recommended that disk drives be evenly distributed among the available channels for optimal performance. This will minimize the effect of any overhead incurred by bus arbitration.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Channel” on page 27, “Drive Packs / Groups” on page 29, “RAID Guidelines” on page 97, “RAID Levels” on page 137, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, “Spanning Drive Packs” on page 62, and “System Drives” on page 149.

Duplex

A disk array system with two active controllers handling host I/O requests, referred to as dual-active controllers. Duplex controller firmware type is either 4, 5, 6, or 7. The firmware type is displayed in the firmware header.

➢ See also “Active Controller” on page 109, “Configuration Strategies” on page 94, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Firmware Header Information” on page 44, “Firmware Versions” on page 95, “Mylex Online RAID Expansion (MORE)” on page 49, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, “Simplex” on page 60 and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

Environmental Device

SCOPE: Support for AEMI monitoring was implemented in firmware versions prior to 3.2 and is supported on the DACSX controller. Support for SAF-TE monitoring was implemented in firmware version 3.2 and is supported on the DACSX, DACSS, DACSF, and DACFL controllers. Support for SES monitoring was implemented in firmware version 5.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

AEMI, SAF-TE, or SES. Environmental monitoring devices that detect disk drive insertion or removal, power supply malfunction, fan malfunction, temperature extremes, and UPS ac failure.

➢ See also “AEMI” on page 6, “Fault Management” on page 41, “SAF-TE” on page 55, and “SES” on page 58.

External RAID Controller

The controller acts as a bridge between host channels and drive channels, and implements the RAID function. On the drive channels, the controller acts as an initiator. The host channel ports are implemented as target IDs, with logical drives presented as LUNs under the target IDs.

➢ See also “Channel” on page 112, “Drive Channel” on page 27, “Host Channel” on page 121, “Host / Server” on page 121, “Logical Drives” on page 125, and “Logical Unit Number (LUN)” on page 125.

External RAID Operation

External RAID Controllers use their host channel ports to connect to one or more host bus adapters in one or more host computer systems.

Depending on the model and configuration of the controller, there may be from one to four host ports. Each configured logical drive consumes a particular bus-target ID-LUN on the host system. Multiple sets of disk drives are connected to the controller via its drive channels. There may be as many as six drive channels, depending on the controller model and configuration.

In duplex, or dual-active mode, two identical External RAID Controllers are connected to the same host(s) via their host ports, to the same drives via their drive channels, and to each other via a special hardware link for heart beat sensing. Dual-active FFx and FFx2 controllers are not physically connected to each other, but communicate through a shared arbitrated loop with the drives. This forms a fault-tolerant controller system. The two controllers work together handling data traffic and mirror their write data in each other’s cache memory. If one of the controllers fails or otherwise becomes non-operational, the surviving controller takes over its responsibilities with no loss of data.

➢ See also “Bus” on page 111, “Cache” on page 111, “Channel” on page 112, “Configuration Coherency” on page 18, “Dual-Active Controller Configuration” on page 32, “Duplex” on page 116, “Failed Controller” on page 118, “Fault Tolerance” on page 118 “Host Bus Adapter (HBA)” on page 120, “Host Channel” on page 121, “Host / Server” on page 121, “Logical Drives” on page 125, “Logical Unit Number (LUN)” on page 125, “Surviving Controller” on page 148, and “Target ID (TID)” on page 149.

Failed Controller

A controller which has been determined to be malfunctioning by its partner. A failed controller is held in reset until replaced.

➢ See also “Active Controller” on page 109, “Automatic Restart (Reboot) on Failure” on page 8, “Debug Dump” on page 25, “Failover and Failback” on page 35, “Partner Controller” on page 133, and “Primary Controller” on page 136.

Failed Disk Drive

A physical disk drive that has failed to operate properly or has been labeled *Offline Failed* or *Dead* by the controller.

➢ See also “Automatic Rebuild” on page 10, “Critical” on page 113, “Dead” on page 113, “Hot Spare” on page 122, “Offline Failed” on page 131, “Online” on page 131, “Online Critical” on page 131, “Online Spare” on page 132, “Physical Device” on page 133, “Physical Device States” on page 134, “Physical Disk Drive” on page 135, “Rebuild” on page 144, “Replacement Drives” on page 145, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, and “Standby” on page 147.

Failover Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel port capable of assuming I/O requests for another (failed) port on the loop. During normal operation (other port is still functional), a failover port may be active or inactive. Failover ports assume the same loop ID and, optionally, the same node name from the failed port.

➢ See also “Active Fibre Port” on page 109, “Failover and Failback” on page 35, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Inactive Fibre Port” on page 123, “Loop Address (Loop ID)” on page 125, “Node Name” on page 130, and “Port” on page 135.

Fault Tolerance

Fault tolerance is the ability of a system to continue to perform its functions during the failure of one of its components. The system may function at a reduced performance level until the failure is corrected. Fault tolerance is implemented through the use of redundant components or through data redundancy.

RAID systems that support fault tolerance provide for complete data redundancy. RAID levels provide fault tolerance in the event of a physical disk drive failure. The data redundancy is created through mirroring the data on identical disk drives or drive groups, or through striping the data across several disk drives. In the event of a disk drive failure, the data can be reproduced from the duplicate disk drives or calculated from the information striped across the disk drives including the parity information.

Mylex controllers operating in a dual-active controller configuration also provide fault tolerance in the event of a controller failure. The surviving controller will assume the duties of the failed controller until a replacement controller becomes active and assumes the duties of the failed controller.

Mylex controllers also provide the ability to prevent data loss through the failure of a server or HBA used in a multiple host configuration.

➢ See also “Alternate Path Software” on page 7, “Automatic Rebuild” on page 10, “CAP Strategy for Selecting a RAID Level” on page 92, “Configuration Strategies” on page 94, “Data Redundancy” on page 113, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Failover and Failback” on page 35, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Host Bus Adapter (HBA)” on page 120, “Host / Server” on page 121, “Mirroring” on page 129, “Parity” on page 133, “RAID Levels” on page 137, “Redundant Array” on page 145, “Replacement Controller” on page 145, “Simplex” on page 60, “Striping” on page 147, and “Surviving Controller” on page 148.

Firmware Flash

Writing a new firmware image to the flash memory on the controller(s). The next time the controller(s) restart, the new firmware is loaded and executed.

➢ See also “Automatic Firmware Flash” on page 7 and “Firmware Versions” on page 95.

Firmware / Software Interface

A direct SCSI command set first implemented with firmware version 7.0. The old command set is supported, but using the older commands is highly discouraged and may result in data loss. Refer to the *Firmware/Software Interface OEM Reference Manual* for detailed information on the supported command set.

Frame

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A data structure used to transfer information from one Fibre Channel port to another. The frame structure includes Start of Frame, End of Frame, frame header information, Cyclic Redundancy Check, and can carry data of 0 to 2112 bytes which includes optional header information.

➢ See also “Cyclic Redundancy Check (CRC)” on page 113, “Fibre Channel” on page 43, and “Port” on page 135.

Full Duplex

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A mode of Fibre Channel communication that allows simultaneous transmission and reception of frames.

➢ See also “Fibre Channel” on page 43, “Frame” on page 119, and “Half Duplex” on page 120.

Gigabyte (GB)

Refers to the capacity of a disk drive or storage device. The value is calculated using two different measurement systems, binary and decimal. The binary value is used by engineers and developers when determining disk drive capacity and how the capacity is utilized. The decimal value is often used in marketing and sales information when referring to disk drive capacity.

The binary value is 2^{30} bytes or 1,073,741,824 bytes.

The decimal value is 10^9 bytes or 1,000,000,000 bytes.

➢ See also “Drive Sizing” on page 30, “Kilobyte (KB)” on page 123, “Megabyte (MB)” on page 129, and “Terabyte (TB)” on page 150.

Half Duplex

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A mode of Fibre Channel communication that allows transmission or reception of frames at any point in time, but not both.

➢ See also “Fibre Channel” on page 43, “Frame” on page 119, and “Full Duplex” on page 119.

Hard Loop ID

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A controller’s preferred loop ID as specified by the saved configuration. The controller attempts to acquire hard loop IDs during the LIHA (Loop Initialization Hard Address) phase of loop initialization.

➢ See also “Controller Parameters” on page 69, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Firmware Versions” on page 95, “Loop Address (Loop ID)” on page 125, and “Microsoft Cluster Server Configurations (MSCS)” on page 47.

Host Bus Adapter (HBA)

A device that connects a host system and the SCSI bus or fibre loop. The device usually performs the lower layers of the SCSI or Fibre Channel protocol and normally operates in the initiator role. One or more HBAs may be integrated into a host/server computer, cluster of host computers, or SAN.

➢ See also “Bus” on page 111, “Configuration Strategies” on page 94, “Dual-Active Controller Configuration” on page 32, “Fibre Channel” on page 43, “Host / Server” on page 121, “Microsoft Cluster Server Configurations (MSCS)” on page 47, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, and “Storage Area Network (SAN)” on page 147.

Host Channel

A path for the transfer of data and control information between the host and the controller. Mylex disk array controllers have access to one or two host channels.

➢ See also “Channel” on page 112, “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive Distribution” on page 116, “Fibre Channel” on page 43, “Host Bus Adapter (HBA)” on page 120, “Host / Server” on page 121, and “Physical Disk Drive” on page 135.

Host Port

The connection on the controller for access between the host channel and the host initiator, usually a host bus adapter (HBA). Mylex controllers have one or two host port connections.

➢ See also “Controller Specifications” on page 1, “Host Bus Adapter (HBA)” on page 120, “Host Channel” on page 121, and “Host / Server” on page 121.

Host / Server

The computer that communicates with the controller and disk array. This communication is controlled through the host bus adapter (HBA). More than one HBA may be integrated into the host/server computer. More than one host/server computer may have access to the controller and disk array, for example, a storage area network (SAN).

The HBA is sometimes referred to as the host while the computer is referred to as the server. Hosts or HBAs are considered initiators. Firmware supports up to 64 hosts when routed through a hub or switch.

➢ See also “Array” on page 111, “Configuration Strategies” on page 94, “Dual-Active Controller Configuration” on page 32, “External RAID Controller” on page 117, “Host Bus Adapter (HBA)” on page 120, “Microsoft Cluster Server Configurations (MSCS)” on page 47, and “Storage Area Network (SAN)” on page 147.

Hot Plugging

The action of removing a failed controller of a dual-active controller pair and inserting a functional replacement while system power is applied. This removal and insertion can occur while the surviving controller in a dual-active controller system is active. Hot plugging a controller requires that the controllers’ dual-active signals be isolated. Refer to the appropriate installation manual for details. Hot plug *does not* include the removal of a functioning controller.

➢ See also “Active Controller” on page 109, “Automatic Firmware Flash” on page 7, “Dual-Active Controller Configuration” on page 32, “Duplex” on page 116, “Failed Controller” on page 118, “Failover and Failback” on page 35, “Partner Controller” on page 133, “Replacement Controller” on page 145, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, “Replacing a Failed Controller in Simplex Systems” on page 99, “SCSI Cabling” on page 99, “SCSI Termination” on page 100, and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

Hot Spare

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to online spare for physical devices. Refer to “Online Spare” on page 132 for more information.

A physical disk drive not part of a system drive that the controller can use to automatically rebuild a critical system drive. The hot spare disk drive must have at least as much capacity as the largest disk drive in the array or the rebuild may not start.

➢ See also “Automatic Rebuild” on page 10, “Configuration Strategies” on page 94, “Dead” on page 113, “Drive Sizing” on page 30, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Swapping” on page 122, “Physical Device States” on page 134, “Rebuild” on page 144, “Replacement Drives” on page 145, and “Standby” on page 147.

Hot Swapping

The action of removing a failed physical disk drive and inserting a functioning physical disk drive while system power is applied. This removal and insertion can occur while the system is handling host I/O requests. Hot swapping online optimal disk drives forces the system drive into an online critical state. A redundant array can sustain one disk drive failure or removal and continue to process I/O requests while in an online critical state. If another disk drive fails or is removed, the system drive is forced into an offline failed state and data loss may occur. A replacement disk drive inserted into the removed disk drive’s slot is considered a replacement disk drive and given a drive state of online spare or standby. If *Automatic Rebuild* is enabled, the data is reconstructed onto the replacement disk drive. If using firmware versions earlier than 7.0, a disk drive can also be inserted into an empty, but configured slot and used as a hot spare and the data is reconstructed using the new disk drive.

➢ See also “Automatic Rebuild” on page 10, “Automatic Rebuild Management” on page 71, “Configuration Strategies” on page 94, “Dead” on page 113, “Drive Sizing” on page 30, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Spare” on page 122, “Operational Fault Management” on page 82, “Physical Device States” on page 134, “Rebuild” on page 144, “Redundant Array” on page 145, “Replacement Drives” on page 145, and “Standby” on page 147.

Hub

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel device that connects nodes into a logical loop. Hubs connect multiple drive channels to one or two host channels. The hub detects when a node has been inserted or failed and automatically adds the new node or removes the failed node while maintaining the loop.

One hub may be connected to another (cascaded) in order to increase arbitrated loop distances. Cascaded hubs allow for distances up to 500 meters between hubs or between a hub and a device.

➢ See also “Channel” on page 112, “Configuration Strategies” on page 94, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Microsoft Cluster Server Configurations (MSCS)” on page 47, and “Node” on page 129.

Inactive Fibre Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel port connected to an arbitrated loop that is not capable of arbitrating the loop. An inactive port cannot be opened and simply retransmits any frames it receives. Also referred to as a *non-participating* port.

➢ See also “Active Fibre Port” on page 109, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, “Fibre Channel” on page 43, “Frame” on page 119, and “Microsoft Cluster Server Configurations (MSCS)” on page 47.

I/O

Input/Output. Refers to disk reads and writes.

Initialize Controllers

During a power cycle or controller reset, each controller progresses through an initialization process. This process includes several checks and validations, which may include validating COD, checking hardware, starting the drives, and determining if there is a C-C nexus. The initialization or controller reset may take several minutes while the above procedures are performed.

➢ See also “Configuration on Disk (COD)” on page 18, “Consistency Check” on page 23, “Controller-Controller Nexus (C-C Nexus)” on page 23, “Disk Startup Mode” on page 75, “Parity” on page 133, and “Reset Controllers” on page 54.

Initialize System Drive

Also referred to as foreground initialization. During the initialization process all data is cleared and zeros are written to the disks.

➢ See also “Background Initialization” on page 11, “Physical Disk Drive” on page 135, and “System Drives” on page 149.

JBOD

Just a Bunch of Disks. A group of drives configured to appear as one logical device. This RAID level does not provide for data redundancy.

JBOD also refers to an enclosure of disk drives.

➢ See also “Data Redundancy” on page 113 and “RAID Levels” on page 137

Kilobyte (KB)

Refers to the capacity of a disk drive or storage device. The value is calculated using two different measurement systems, binary and decimal. The binary value is used by engineers and developers when determining disk drive capacity and how the capacity is utilized. The decimal value is often used in marketing and sales information when referring to disk drive capacity.

The binary value is 2^{10} bytes or 1,024 bytes.

The decimal value is 10^3 bytes or 1,000 bytes.

➢ See also “Drive Sizing” on page 30, “Gigabyte (GB)” on page 120, “Megabyte (MB)” on page 129, and “Terabyte (TB)” on page 150.

Link

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A means of communication between nodes in a Fibre Channel point-to-point topology. Links include two unidirectional fibre nodes transmitting in opposite directions and the associated transmitters and receivers that communicate between the nodes. A link is similar to the bus described in SCSI protocol.

➢ See also “Bus” on page 111, “Fibre Channel” on page 43, “Node” on page 129, and “SCSI” on page 145.

Logical Device States

The possible logical device (system drive) states are defined as *Online Optimal*, *Online Critical*, *Online Critical and Rebuilding*, and *Offline Failed*. The following illustration shows how a logical device changes from one state to another. Logical device states are dependent upon disk drive failure and rebuilding a failed disk drive.

Logical Device States

➢ See also "Automatic Rebuild" on page 10, "Device States" on page 114, "Drive State Management" on page 31, "Failed Disk Drive" on page 118, "Online Critical" on page 131, "Online Critical and Rebuilding" on page 132, "Online Optimal" on page 132, "Physical Device" on page 133, "Physical Device States" on page 134, "Redundant Array" on page 145, "System Drive State" on page 148, and "System Drives" on page 149.

Logical Drives

A logical drive is equivalent to a "system drive" or "logical device." Logical drives are presented to the operating system as available disk drives, each with a capacity specified by the controller.

➢ See also "System Drives" on page 149.

Logical Unit

SCSI defines a logical unit as a peripheral device, physical or virtual, that is addressed through a target. Disk storage space on one or more physical disk drives that appears to the host computer as one drive is a virtual logical unit. Sometimes logical units are referred to as system drives or logical drives. Mylex controllers operating with firmware versions earlier than 7.0 support up to 8 logical units. Mylex controllers operating with firmware versions 7.0 and greater support up to 32 logical units.

➢ See also "Logical Drives" on page 125, "Logical Unit Number (LUN)" on page 125, "LUN Mapping" on page 47, "Physical Disk Drive" on page 135, "Programmable LUN Mapping" on page 52, "SANmapping™" on page 56, "System Drive Affinity" on page 64, and "System Drives" on page 149.

Logical Unit Number (LUN)

A SCSI representation of a system drive on a given channel and target ID. This may be a single device or an array of devices configured to behave as a single device. Mylex controllers operating with firmware versions earlier than 7.0 support up to 8 logical units. Mylex controllers operating with firmware versions 7.0 and greater support up to 32 logical units.

➢ See also "Array" on page 111, "Channel" on page 112, "Logical Unit" on page 125, "LUN Mapping" on page 47, "Programmable LUN Mapping" on page 52, "SANmapping™" on page 56, "SCSI" on page 145, "System Drive Affinity" on page 64, "System Drives" on page 149, and "Target ID (TID)" on page 149.

Loop Address (Loop ID)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A unique ID determined at loop initialization for a node in a FC-AL topology. This is a seven-bit value numbered from 0 to 126 and represent the 127 legal AL_PA values on a loop.

➢ See also "Arbitrated Loop Physical Address (AL_PA)" on page 111, "Fibre Channel" on page 43, "Soft Addressing Detection" on page 61, and "Node" on page 129.

Loop Initialization Fabric Assigned (LIFA)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to gather all fabric assigned AL_PAs.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, and “Fibre Channel” on page 43.

Loop Initialization Hard Assigned (LIHA)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to gather all hard assigned AL_PAs.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, “Fibre Channel” on page 43, and “Hard Loop IDs” on page 79.

Loop Initialization Loop Position (LILP)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to inform all L_Ports of the relative positions of all participating L_Ports on the loop. These positions are relative to the LIM.

➢ See also “Fibre Channel” on page 43, “Loop Initialization Master (LIM)” on page 126, and “Loop Port (L_Port)” on page 128.

Loop Initialization Master (LIM)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The L_Port responsible for initializing the loop.

➢ See also “Fibre Channel” on page 43, “Loop Initialization Primitive (LIP)” on page 126, and “Loop Port (L_Port)” on page 128.

Loop Initialization Primitive (LIP)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The Loop Initialization Primitive (LIP) sequence is used by a L_Port to determine if it is part of an arbitrated loop or to recover from certain loop errors. A L_Port requests a LIP to acquire an AL_PA, to indicate a loop failure or invalid AL_PA, to reinitialize the loop or return the loop to a known state, to indicate that a loop failure has been detected, or to reset the NL_Port or to issue a vendor specific loop reset.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, “Fibre Channel” on page 43, “Loop Port (L_Port)” on page 128, “Node Loop Port (NL_Port)” on page 129, and “Node Port (N_Port)” on page 130.

Loop Initialization Previously Acquired (LIPA)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to gather all previously acquired AL_PAs.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, “Fibre Channel” on page 43, and “Previous Loop ID” on page 136.

Loop Initialization Report Position (LIRP)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to collect the relative positions of all participating L_Ports on the loop.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, “Fibre Channel” on page 43, and “Loop Port (L_Port)” on page 128.

Loop Initialization Soft Assigned (LISA)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to assign any remaining addresses as soft assigned AL_PAs.

➢ See also “Arbitrated Loop Physical Address (AL_PA)” on page 111, “Fibre Channel” on page 43, “Soft Addressing Detection” on page 61, and “Hard Loop IDs” on page 79.

Loop Initialization Select Master (LISM)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A loop initialization sequence used to select a LIM.

➢ See also “Fibre Channel” on page 43 and “Loop Initialization Master (LIM)” on page 126.

Loop Port (L_Port)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel port that supports an arbitrated loop topology.

➢ See also “Fibre Channel” on page 43 and “Port” on page 135.

Loop Port Bypass (LPB)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Loop Port Bypass (LPByx) is transmitted on a Fibre Channel arbitrated loop to bypass a L_Port and to activate the optional bypass circuit. A LPB can be used to diagnose the optional bypass circuit or for error recovery. After the L_Port is bypassed and the optional bypass circuit is activated, the L_Port only monitors the loop for a LPE and LIP.

A Loop Port Bypass all (LPBfx) is transmitted on a loop to bypass all L_Ports and activate the optional bypass circuits of all L_Ports. A LPBfx is used to verify that an operating loop is possible. It is also useful to bypass a non-participating L_Port.

➢ See also “Fibre Channel” on page 43, “Loop Initialization Primitive (LIP)” on page 126, “Loop Port (L_Port)” on page 128, “Loop Port Enable (LPE)” on page 128, “Port” on page 135, and “Port Bypass Circuit (PBC)” on page 135.

Loop Port Enable (LPE)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Loop Port Enable (LPEyx) is transmitted on a Fibre Channel arbitrated loop to enable a L_Port that was previously bypassed and deactivates the optional bypass circuit without a LIP. After the LPE is recognized, the bypassed L_Port may participate on the loop.

A Loop Port Enable (LPEfx) is transmitted on loop to deactivate all bypass circuits and to enable all L_Ports to participate on the loop.

➢ See also “Fibre Channel” on page 43, “Loop Initialization Primitive (LIP)” on page 126, “Loop Port (L_Port)” on page 128, “Loop Port Bypass (LPB)” on page 128, “Port” on page 135, and “Port Bypass Circuit (PBC)” on page 135.

Loop Redundancy Circuit (LRC)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

Also referred to as Loop Resiliency Circuit or Port Bypass Circuit (PBC). A hardware device that detects loop breakage and enables a loop bypass to recover loop operation. LRCs are required for every disk slot and for every disk loop expansion connections.

➢ See also “Fibre Channel” on page 43, “Port Bypass Circuit (PBC)” on page 135, and “Switch” on page 148.

Megabyte (MB)

Refers to the capacity of a disk drive or storage device. The value is calculated using two different measurement systems, binary and decimal. The binary value is used by engineers and developers when determining disk drive capacity and how the capacity is utilized. The decimal value is often used in marketing and sales information when referring to disk drive capacity.

The binary value is 2^{20} bytes or 1,048,576 bytes.

The decimal value is 10^6 bytes or 1,000,000 bytes.

➢ See also “Drive Sizing” on page 30, “Gigabyte (GB)” on page 120, “Kilobyte (KB)” on page 123, and “Terabyte (TB)” on page 150.

Mirroring

The 100% duplication of data on one disk drive to another disk drive. Each disk drive is the mirror image of the other; RAID level 1.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Consistency Check” on page 23, “RAID Levels” on page 137, and “Redundant Array” on page 145.

Name_Identifier

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Name_Identifier is a Fibre Channel protocol, 64-bit identifier with a 60-bit value preceded by a 4-bit Network_Address_ Authority_Identifier.

➢ See also “Fibre Channel” on page 43.

Node

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel entity (device) with one or more N_Ports or NL_Ports.

➢ See also “Fibre Channel” on page 43, “Node Loop Port (NL_Port)” on page 129, “Node Port (N_Port)” on page 130, and “Port” on page 135.

Node Loop Port (NL_Port)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The node loop port provides the mechanisms necessary to route information around the Fibre Channel arbitrated loop and to repeat the information until it arrives at its destination. Each NL_Port is identified with a port name.

➢ See also "Fibre Channel" on page 43, "Node" on page 129, "Node Port (N_Port)" on page 130, and "Port Name" on page 135.

Node Name

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A World Wide Name identifier associated with a node.

➢ See also "Fibre Channel" on page 43, "Node" on page 129, "World Wide Name (WWN)" on page 151, and "World Wide Name (WWN) Assignments" on page 65.

Node Port (N_Port)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The node port provides the mechanisms necessary to transport information units to or from another node. Each node port is identified with a port name.

➢ See also "Fibre Channel" on page 43, "Node" on page 129, "Node Loop Port (NL_Port)" on page 129, and "Port Name" on page 135.

NVRAM

Non-Volatile RAM. This is memory that maintains its contents across a power cycle. This memory is used to store controller configuration information.

➢ See also "Configuration on Disk (COD)" on page 18 and "Mylex Online RAID Expansion (MORE)" on page 49.

Offline

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to offline failed for a system drive or logical device. Refer to "Offline Failed" on page 131 for more information.

Refers to the state of a system drive. Offline does not apply to physical disk drives. A system drive is offline if it is configured with a redundant RAID level and two or more of its member disk drives are not online; or it is configured at RAID level 0 or JBOD and one or more of its member disk drives are not online.

➢ See also "Alarm Signal" on page 6, "Critical" on page 113, "Logical Device States" on page 124, "Online" on page 131, "System Drives" on page 149, and "System Drive State" on page 148.

Offline Failed

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was dead for physical devices and offline for system drives or logical devices. Refer to “Dead” on page 113 or “Offline” on page 130 for more information.

A device state that applies to physical devices and logical devices. An offline failed physical device is powered on but not functioning normally and should be replaced.

An offline failed logical device has suffered device failures that exceed the limit for the configured RAID type. In order to retrieve a logical device from the offline failed state, the failed physical devices must be returned to an online optimal state, either through a rebuild procedure or device state change.

➢ See also “Device States” on page 114, “Failed Disk Drive” on page 118, “Logical Device States” on page 124, “Physical Device” on page 133, “Physical Device States” on page 134, and “RAID Guidelines” on page 97.

Online

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to online optimal for a system drive or logical device. Refer to “Online Optimal” on page 132 for more information.

The state of a system drive or a physical disk drive. A system drive is online if all of its member disk drives are online. A physical disk drive is online if it is powered on, has been defined as a member of a drive pack, and is operating properly.

➢ See also “Alarm Signal” on page 6, “Critical” on page 113, “Drive Packs / Groups” on page 29, “Drive State Management” on page 31, “Offline” on page 130, “Physical Disk Drive” on page 135 “System Drives” on page 149, and “System Drive State” on page 148.

Online Critical

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was critical for system drives or logical devices. Refer to “Critical” on page 113 for more information.

A device state that applies to a logical device. A logical device is online critical when a device failure has occurred that is within the limit for the configured RAID type. If the failed device is replaced and data successfully rebuilt, the logical device state changes to online optimal. If the logical device suffers another device failure, the state changes to offline failed.

➢ See also “Device States” on page 114, “Failed Disk Drive” on page 118, “Logical Device States” on page 124, “Offline Failed” on page 131, “Online Optimal” on page 132, “RAID Levels” on page 137, and “Rebuild” on page 144.

Online Critical and Rebuilding

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was critical for a system drive or logical device. Refer to "Critical" on page 113 for more information.

A device state that applies to a logical device. The logical device has suffered the failure of one device and is in the process of rebuilding the failed device. If the rebuild procedure completes successfully, the state changes to online optimal. If the logical device suffers another device failure, the state changes to offline failed.

➢ See also "Device States" on page 114, "Failed Disk Drive" on page 118, "Logical Device States" on page 124, "Offline Failed" on page 131, "Online Optimal" on page 132, "Online Spare" on page 132, and "Rebuild" on page 144.

Online Optimal

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was online for physical disk drives and system drives or logical devices. Refer to "Online" on page 131 for more information.

A device state that applies to a logical device or a physical device. A logical device is online optimal if all member devices are powered on, configured into a RAID group, and functioning normally.

A physical device is online optimal if it is powered on, defined as a member of a RAID device, and functioning normally.

➢ See also "Device States" on page 114, "Logical Device States" on page 124, "Physical Device" on page 133, and "Physical Device States" on page 134.

Online Rebuild

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was rebuild for physical disk drives. Refer to "Rebuild" on page 144 for more information.

A device state that applies to a physical device that is powered on, part of a configuration, and in the process of being rebuilt.

➢ See also "Device States" on page 114, "Online Spare" on page 132, "Physical Device" on page 133, "Physical Device States" on page 134, and "Rebuild" on page 144.

Online Spare

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was standby or hot spare for physical disk drives. Refer to "Standby" on page 147 or "Hot Spare" on page 122 for more information.

A device state that applies to a physical device that is part of a configuration and operating as a spare drive in the event of another device failure and will be used for automatic rebuild.

➢ See also “Automatic Rebuild” on page 10, “Device States” on page 114, “Failed Disk Drive” on page 118, “Offline Failed” on page 131, “Online Optimal” on page 132, “Online Spare Polling” on page 51, and “Rebuild” on page 144.

Parity

A method of providing complete data redundancy while requiring only a fraction of the storage capacity of mirroring. The data and parity blocks are divided between the disk drives in such a way that if any single disk drive is removed (or fails), the data on it can be reconstructed using the data on the remaining disk drives. The parity data may exist on only one disk drive or be distributed between all disk drives in the RAID group.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Consistency Check” on page 23, “Data Redundancy” on page 113, “Mirroring” on page 129, “Parity Check” on page 51, “RAID Levels” on page 137, “Redundant Array” on page 145, and “Striping” on page 147.

Partner Controller

In a dual-active controller configuration, the partner controller is the controller that is not being accessed.

➢ See also “Active Controller” on page 109, “Automatic Restart (Reboot) on Failure” on page 8, “Cache Coherency” on page 17, “Configuration Coherency” on page 18, “Controller-Controller Nexus (C-C Nexus)” on page 23, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Debug Dump” on page 25, “Failed Controller” on page 118, “Multiple Target ID (MTID)” on page 48, “Primary Controller” on page 136, “Secondary Controller” on page 146, and “Surviving Controller” on page 148.

Physical Device

A hardware device (such as a disk drive) that is physically connected to the hardware. The physical device is addressed within the controller by channel number, target ID, and LUN.

➢ See also “Channel” on page 112, “Logical Unit Number (LUN)” on page 125, and “Target ID (TID)” on page 149.

Physical Device Access

All non-configured physical devices are accessible by host software; therefore, a physical disk drive is accessible if not configured. The controller behaves as a true SCSI host bus adapter with the following exceptions:

- Tagged commands are not allowed—the tag number is ignored.
- SCSI message control is not allowed—the host cannot send/receive SCSI messages.
- Host software cannot access the SCSI device in target mode.

The physical device addresses are:

- LUN ID
- Target ID
- Channel number (number of the physical channel to which the physical device is connected)

➢ See also “Channel” on page 112, “Controller Parameters” on page 69, “Host Bus Adapter (HBA)” on page 120, “Host / Server” on page 121, “Logical Unit Number (LUN)” on page 125, “Physical Device” on page 133, “SCSI” on page 145, and “Target ID (TID)” on page 149.

Physical Device States

The possible physical device states are defined as *Online Optimal*, *Online Rebuild*, *Online Spare*, *Online Critical*, *Commanded Offline*, *Offline Failed*, and *Unconfigured Offline*. Only Online Optimal, Online Rebuild, Online Spare, Offline Failed, and Unconfigured Offline are supported at this time. The following illustration shows how a physical device changes from one state to another. Physical device states are dependent upon disk drive failure, rebuilding a failed disk drive, and being part of a configuration.

Physical Device States

➢ See also “Automatic Rebuild” on page 10, “Device States” on page 114, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Logical Device States” on page 124, “Offline Failed” on page 131, “Online Optimal” on page 132, “Online Rebuild” on page 132, “Online Spare” on page 132, “Physical Device” on page 133, “Predictive Failure Analysis™ (PFA)” on page 52, “Rebuild” on page 144, “Redundant Array” on page 145, and “Unconfigured Offline” on page 150.

Physical Disk Drive

A single hard disk drive. Each physical disk drive is assigned a unique identification address.

➢ See also “Channel” on page 112, “Drive and Controller Addressing” on page 115, “Drive Packs / Groups” on page 29, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Physical Device States” on page 134“System Drives” on page 149, and “Target ID (TID)” on page 149.

Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The hardware connection within a node that performs data communication over the Fibre Channel link.

➢ See also “Fibre Channel” on page 43, “Link” on page 124, and “Node” on page 129.

Port Bypass Circuit (PBC)

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

Also referred to as a Loop Resiliency Circuit or Loop Redundancy Circuit (LRC). A circuit used in hubs and disk enclosures to automatically open or close the Fibre Channel loop to add or remove nodes (devices) on the loop.

➢ See also “Fibre Channel” on page 43, “Hub” on page 122, “Loop Redundancy Circuit (LRC)” on page 128, and “Node” on page 129

Port Name

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A unique 64-bit value identifier associated with a port. The port name is usually assigned to the port during manufacturing and does not change throughout the lifetime of the port.

➢ See also “Fibre Channel” on page 43, “Node Loop Port (NL_Port)” on page 129, “Node Port (N_Port)” on page 130, “Port” on page 135, and “World Wide Name (WWN) Assignments” on page 65.

Preferred Address

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFL, DACFF (FF), FFx, FF2, and FFx2 controllers.

Dual-ported drives can be connected to two drive channels. When the controller firmware looks for drives, one of the drive ports is assigned a *preferred* (channel-target ID) address; the other is assigned an *alternate* address. When a drive channel fails, if that channel is associated with the preferred address, the drive is accessed through its alternate address. When the failed channel is restored, the drive is again accessed through its preferred address.

➢ See also “Alternate Address” on page 110, “Drive Channel” on page 27, “Drive Channel Failover and Failback” on page 28, “Dual-Ported Drive Support” on page 34, “Loop Address (Loop ID)” on page 125, and “Target ID (TID)” on page 149

Preferred Channel

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

At drive discovery time, a dual-ported drive may be found on two different controller drive channels. One of these channels is assigned the drive's preferred channel; the other is the drive’s alternate channel. The controller sends all commands to the drive on its preferred channel unless the preferred channel is not available (the channel has failed or the drive is no longer accessible over its preferred port).

➢ See also “Alternate Channel” on page 110, “Drive Channel” on page 27, “Drive Channel Failover and Failback” on page 28, and “Dual-Ported Drive Support” on page 34.

Previous Loop ID

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The loop ID acquired during a prior loop initialization. The controller acquires previous loop IDs during the LIPA (Loop Initialization, Previous Address) phase of loop initialization.

➢ See also “Controller Parameters” on page 69, “Drive and Controller Addressing” on page 115, “Fibre Channel” on page 43, “Hard Loop ID” on page 120, “Loop Address (Loop ID)” on page 125, and “Loop Initialization Previously Acquired (LIPA)” on page 127.

Primary Controller

In a dual-active controller configuration the primary controller is the controller that is currently being accessed. Convention has assigned controller 0 as the primary controller.

➢ See also “Active Controller” on page 109, “Dual-Active Controller Configuration” on page 32, “Failover and Failback” on page 35, “Partner Controller” on page 133, and “Secondary Controller” on page 146.

Primary ID

The target ID set by hardware through the backplane connector.

➢ See also “Failover and Failback” on page 35, “Multiple Target ID (MTID)” on page 48, “Secondary ID” on page 146, and “Target ID (TID)” on page 149.

Primary Path

SCOPE: This term applies to Fibre Channel, dual-ported disk drives and is a part of the Drive Channel Failover and Failback feature. This feature was implemented in firmware version 6.0 and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

The primary path for data transfer from a host to a LUN. In a dual-active configuration, one controller/host port combination serves as the primary data path and the secondary controller/host port combination serves as an alternate path. If a failure occurs to the primary path, the alternate path assumes all activity for the failed path, maintaining access to the LUN.

➢ See also “Alternate Path” on page 110, “Alternate Path Software” on page 7, “Dual-Active Controller Configuration” on page 32, “Host / Server” on page 121, “Logical Unit Number (LUN)” on page 125, “LUN Mapping” on page 47, “Programmable LUN Mapping” on page 52, “SANmapping™” on page 56, and “System Drive Affinity” on page 64.

Primary Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

The active port on a controller during normal dual-active controller operation.

➢ See also “Active Fibre Port” on page 109, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Failover and Failback” on page 35, “Failover Port” on page 118, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, “Failover Topology, Multiport” on page 40, and “Inactive Fibre Port” on page 123.

RAID Levels

The Mylex controllers implement several different versions of the Berkeley RAID technology. Correct installation of the disk array and the controller requires a proper understanding of RAID technology and concepts.

An appropriate RAID level must be selected when system drives are defined or created. This decision is based on the following priorities:

- Disk capacity utilization (number of disk drives)
- Data redundancy (fault tolerance)
- Disk performance

The Mylex controllers make the RAID implementation and the disk drives’ physical configuration transparent to the host operating system. This means that the host operating logical drivers and software utilities are unchanged, regardless of the RAID level selected.

Berkeley RAID Levels

Mylex controllers support RAID 0, RAID 1, RAID 3, RAID 5, RAID 0 +1, and JBOD. Each RAID level is described in the following illustrations.

RAID 0 provides block striping which yields higher performance than individual disk drives. No redundancy or fault tolerance is provided, disk drive failure is not tolerated. RAID 0 requires a minimum of 2 disk drives and supports a maximum of 8 disk drives using firmware versions earlier than 7.0 or 16 disk drives using firmware versions 7.0 and greater.

RAID 1 configures 2 disk drives as a mirrored pair. All data is 100% duplicated on an equivalent disk drive. Fault tolerance is provided by complete duplication of data. RAID 1 is configured from 2 disk drives, no more, no less.

Mylex controllers configure RAID 3 and RAID 5 exactly the same. Data is striped and rotated across all disk drives in the drive pack. Parity protection is used for data redundancy and fault tolerance. RAID 3 and 5 require a minimum of 3 disk drives and support a maximum of 8 disk drives using firmware earlier than 7.0 or 16 disk drives using firmware 7.0 and greater.

RAID 0 + 1 is a combination of RAID levels 0 and 1. Mirror images of the data are striped and staggered across sets of disk drives. Data redundancy and fault tolerance are provided through mirroring. RAID 0 + 1 require a minimum of 3 disk drives and supports a maximum of 8 disk drives using firmware earlier than 7.0 or 16 disk drives using firmware 7.0 and greater.

JBOD, “Just a Bunch of Drives” treats each disk drive as a normal disk drive. Each disk drive operates independently. JBOD does not provide data redundancy or fault tolerance. JBOD requires a minimum of 1 disk drive and supports a maximum of 8 disk drives using firmware versions earlier than 7.0 or 32 disk drives using firmware version 7.0 and greater configured as independent system drives.

Spanned RAID Levels

Mylex controllers also support spanning across drive packs to create RAID 10, RAID 30 and RAID 50 configurations. Each RAID level is described in the following illustrations.

RAID 10 is created when more than two disk drives are configured as RAID 1 system drives. The disk drives are divided into drive packs of two disk drives. The second disk drive in each pack becomes a mirror image of the first disk drive. Spanning is used to stripe the data across up to 16 drive packs. A RAID 10 system drive must have drive packs with only two disk drives. A RAID 10 system drive can be configured from a minimum of four disk drives and a maximum of 32.

RAID 30 and RAID 50 are created when two or more drives packs are spanned to create RAID 3 or RAID 5 system drives. Each drive pack will maintain parity data for the disk drives in the drive pack. A RAID 30 and RAID 50 system drive must have a minimum of 3 disk drives per pack and will support a maximum of 16 disk drives per pack.

Mixed RAID Levels

Although a system drive may have only one RAID level, RAID levels can be mixed within a drive pack as shown in the following illustration.

RAID Levels Within a Drive Pack

The smaller system drive (B0) is assigned a RAID 5 level of operation, while the larger system drive (B1) is assigned a RAID 0+1 level of operation.

Remember that different RAID levels exhibit different performance characteristics for a particular application or environment. Mylex controllers allow for complete versatility in this regard, by allowing multiple, different RAID levels to be assigned to a drive pack.

➢ See also “Array” on page 111, “CAP Strategy for Selecting a RAID Level” on page 92, “Configuration Strategies” on page 94, “Data Redundancy” on page 113, “Drive Distribution” on page 116, “Drive Packs / Groups” on page 29, “Fault Tolerance” on page 118, “Firmware Versions” on page 95, “Host / Server” on page 121, “Mirroring” on page 129, “Parity” on page 133, “RAID Guidelines” on page 97, “Redundant Array” on page 145, “Spanning Drive Packs” on page 62, “Striping” on page 147, and “System Drives” on page 149.

Rebuild

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to online rebuild for a physical disk drive. Refer to “Online Rebuild” on page 132 for more information.

This term refers to the state of a physical disk drive that is in the process of being rebuilt.

Rebuild is also the process of regenerating and writing data to a replacement disk drive. During a rebuild process, the system drive operates in an online critical and rebuilding, or degraded state. The physical device is operating in an online rebuild state.

➢ See also “Automatic Rebuild” on page 10, “Automatic Rebuild Management” on page 71, “Controller Parameters” on page 69, “Critical” on page 113, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Spare” on page 122, “Logical Device States” on page 124, “Online Critical” on page 131, “Online Critical and Rebuilding” on page 132, “Online Rebuild” on page 132, “Online Spare” on page 132, “Operational Fault Management”

on page 82, “Physical Device States” on page 134, “Rebuild and Check Consistency Rate” on page 84, “Redundant Array” on page 145, “Replacement Drives” on page 145, and “Standby” on page 147.

Redundant Array

A RAID level that provides complete data redundancy. In the event of a disk drive failure or removal, the data can be reconstructed using the data on the remaining disk drives. RAID 1, RAID 3, RAID 5, and RAID 0+1 are redundant arrays.

➢ See also “Array” on page 111, “Automatic Rebuild” on page 10, “CAP Strategy for Selecting a RAID Level” on page 92, “Data Redundancy” on page 113, “Mirroring” on page 129, “Parity” on page 133, “RAID Levels” on page 137 and “Striping” on page 147.

Replacement Controller

A functioning controller that replaces a failed controller. In a dual-active controller pair, a failed controller can be replaced while power is supplied to the system (hot plugging). The new controller is referred to the replacement controller until it is fully operational.

➢ See also “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Failover and Failback” on page 35, “Hot Plugging” on page 121, “Replacing a Failed Controller in Simplex Systems” on page 99, “Replacing a Failed Controller in Existing Duplex Systems” on page 97, and “Surviving Controller” on page 148.

Replacement Drives

A physical disk drive that replaces a failed disk drive.

➢ See also “Automatic Rebuild” on page 10, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Spare” on page 122, “Online Rebuild” on page 132, “Physical Device States” on page 134, “Rebuild” on page 144, and “Standby” on page 147.

Rotated XOR Redundancy

This term (also known as “parity” or “consistency data”) refers to a method of providing complete data redundancy while requiring only a fraction of the storage capacity of mirroring. In a system configured under RAID 3 or RAID 5 (which require at least three disk drives), all data and parity blocks are divided between the drives in such a way that if any single drive is removed (or fails), the data on it can be reconstructed using the data on the remaining drives (XOR refers to the Boolean “Exclusive-OR” operator). In any RAID 3 or RAID 5 array, the capacity allocated to redundancy is the equivalent of one drive.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Data Redundancy” on page 113, “Failed Disk Drive” on page 118, “Mirroring” on page 129, “Parity” on page 133, and “RAID Levels” on page 137.

SCSI

Small Computer System Interface (SCSI) is a technological standard that defines connections between host computers and peripheral devices.

➢ See also “Fibre Channel” on page 43, “Host / Server” on page 121, “SCSI Cabling” on page 99, and “SCSI Termination” on page 100.

Secondary Controller

The partner to the primary controller. Convention has assigned controller 1 as the secondary controller.

➢ See also “Active Controller” on page 109, “Dual-Active Requirements” on page 33, “Failover and Failback” on page 35, “Partner Controller” on page 133, and “Primary Controller” on page 136.

Secondary ID

The second Target ID to which the controller will respond when the partner controller fails.

➢ See also “Failed Controller” on page 118, “Failover and Failback” on page 35, “Multiple Target ID (MTID)” on page 48, “Partner Controller” on page 133, “Primary ID” on page 136, and “Target ID (TID)” on page 149.

Simplex Fibre Channel System

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

In a simplex system, all Fibre Channel host ports are active and accept I/O requests.

➢ See also “Active Fibre Port” on page 109, “Fibre Channel” on page 43, “Host Port” on page 121, and “Simplex” on page 60.

Single Controller Mode

A single controller that is attached to a set of disk drives and offers RAID functionality without the fault tolerance of dual-active controller mode.

➢ See also “Configuration Strategies” on page 94, “Controller Parameters” on page 69, “Dual-Active Controller Configuration” on page 32, “Fault Tolerance” on page 118, “Force Simplex” on page 78, “Mylex Online RAID Expansion (MORE)” on page 49, “Replacing a Failed Controller in Simplex Systems” on page 99, “Simplex” on page 60, and “Upgrading From Simplex to Dual-Active Configurations” on page 103.

SIOP

SCSI I/O Processor.

Standby

SCOPE: This term is used by firmware versions earlier than 7.0. The term changed to online spare for physical devices. Refer to “Online Spare” on page 132 for more information.

A physical disk drive that is powered on, is able to operate properly, and has been defined as a hot spare or standby disk drive.

➢ See also “Automatic Rebuild” on page 10, “Drive State Management” on page 31, “Failed Disk Drive” on page 118, “Hot Spare” on page 122, “Hot Swapping” on page 122, and “Replacement Drives” on page 145.

Storage Area Network (SAN)

Configurations in which multiple host computers attach to one or more controllers. The host computers are attached to the controller(s) through a Fibre Channel arbitrated loop or through a switch.

➢ See also “Fibre Channel” on page 43, “Host Bus Adapter (HBA)” on page 120, “Host / Server” on page 121, and “SANmapping™” on page 56.

Striping

The practice of writing data to all disk drives in a RAID group. User data and/or parity is stored sequentially across all data disk drives in the RAID group. An example of striping with parity is shown in the following illustration.

See also “CAP Strategy for Selecting a RAID Level” on page 92, “Data Redundancy” on page 113, “Disk Striping” on page 115, “Drive Packs / Groups” on page 29, “Parity” on page 133, “RAID Levels” on page 137, “Redundant Array” on page 145, and “Stripe Size” on page 64.

Surviving Controller

A controller that has determined that its partner controller in a dual-active controller pair has failed and has assumed the duties of both controllers. An indication of a surviving controller is maintained in volatile storage. This controller is referred to as the surviving controller until the replacement controller is fully operational.

➢ See also “Active Controller” on page 109, “Automatic Restart (Reboot) on Failure” on page 8, “Debug Dump” on page 25, “Dual-Active Controller Configuration” on page 32, “Failed Controller” on page 118, “Failover and Failback” on page 35, “NVRAM” on page 130, “Partner Controller” on page 133, “Replacement Controller” on page 145, and “Replacing a Failed Controller in Existing Duplex Systems” on page 97.

Switch

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A Fibre Channel device used for connecting multiple ports on a Fibre Channel arbitrated loop or Fibre Channel fabric topology.

➢ See also “Fibre Channel” on page 43 and “Port” on page 135.

System Drive Format

The operating system formats the newly created system drive so that the system drive is accessible through the operating system and can perform I/O operations.

➢ See also “I/O” on page 123 and “System Drives” on page 149.

System Drive State

SCOPE: These terms are used firmware versions earlier than 7.0. The terms changed to online optimal, online critical, online critical and rebuilding, and offline failed for a system drive or logical device. Refer to “Online Optimal” on page 132, “Online Critical” on page 131, “Online Critical and Rebuilding” on page 132, and “Offline Failed” on page 131 for more information.

The current operational state of a system drive: *Online*, *Critical*, or *Offline*.

➢ See also “CAP Strategy for Selecting a RAID Level” on page 92, “Critical” on page 113, “Logical Device States” on page 124, “Offline” on page 130, “Online” on page 131, and “System Drives” on page 149

System Drives

SCOPE: This term has been used by the firmware and configuration tools and utilities through the various firmware versions and releases. The system drive specifications have been modified in firmware version 7.0 to reflect the support for increased number of devices and is supported on the DACFF (FF), FFx, FF2, and FFx2 controllers.

System drives are the logical devices (storage volumes) that are presented to the operating system. During the configuration process, after physical disk drive packs are defined, one or more system drives must be created from the drive packs. System drives have the following properties:

- More than one system drive can be defined on a single drive pack; or a system drive can span 4 (< 7.0) or 16 (> 7.0) packs.
- The minimum size of a system drive is 8 MB. The maximum size is 2 TB.
- Up to 8 (< 7.0) or 32 (> 7.0) system drives can be created.
- Each system drive has a RAID level which is selectable (subject to the number of disk drives in the system drive's pack).
- Each system drive has its own write policy (write-back or write-through).
- Each system drive has its own affinity or LUN mapping.

➢ See also "Drive Packs / Groups" on page 29, "Logical Unit" on page 125, "Physical Disk Drive" on page 135, "Programmable LUN Mapping" on page 52, "RAID Levels" on page 137, "Reset Controllers" on page 54, "SANmapping™" on page 56, "SCSI Cabling" on page 99, "SCSI Termination" on page 100, "Spanning Drive Packs" on page 62, "System Drive Affinity" on page 64, "Write-Back Cache" on page 151, and "Write-Through Cache" on page 151.

Target ID (TID)

SCOPE: The support for Target IDs has evolved through the various firmware versions and releases. As technology has allowed for supporting greater number of devices, addressing has changed to accommodate for the greater number of devices.

Firmware version 3.2 supports 16 SCSI IDs on a drive channel, but reserves two IDs for the controllers.

Firmware version 5.0 supports the fibre loop topology and 16 drives on a drive channel.

Firmware version 7.0 supports 126 drives on a drive channel, but reserves two IDs for the controllers.

The Target ID is the SCSI ID or the fibre ID of a device attached to a controller. In systems with SCSI disk busses, SCSI target ID 7 is reserved for controller 0, and in dual-active controller systems SCSI target ID 6 is reserved for controller 1. Disk drives and environmental devices may use any of the remaining SCSI target IDs; that is, 0-5 and 8-15 on dual-active controller systems or 0-6 and 8-15 on simplex systems.

In systems with fibre drive channels, hard addressing is required. Fibre loop ID 113 (71h) is reserved for controller 0, and in dual-active controller systems fibre loop ID 112 (70h) is reserved for controller 1. Disk drives may use any of the remaining fibre loop IDs; that is, 0-111 (0-6Fh) and 114-125 (72h-7Dh). A maximum of 16 disk drives per fibre drive channel is supported with firmware versions earlier than 7.0. A maximum of 124 disk drives per fibre drive channel is supported with firmware versions 7.0 and greater.

➢ See also “Channel” on page 112, “Configuration Strategies” on page 94, “Controller Specifications” on page 1, “Drive and Controller Addressing” on page 115, “Drive Channel” on page 27, “Dual-Active Controller Configuration” on page 32, “Environmental Device” on page 117, “Fibre Channel” on page 43, “Firmware Versions” on page 95, “Hard Loop ID” on page 120, “Loop Address (Loop ID)” on page 125, “Multiple Target ID (MTID)” on page 48, “Physical Device” on page 133, and “SCSI” on page 145.

Terabyte (TB)

Refers to the capacity of a disk drive or storage device. The value is calculated using two different measurement systems, binary and decimal. The binary value is used by engineers and developers when determining disk drive capacity and how the capacity is utilized. The decimal value is often used in marketing and sales information when referring to disk drive capacity.

The binary value is 2^{40} bytes or 1,099,511,627,776 bytes.

The decimal value is 10^{12} bytes or 1,000,000,000,000 bytes.

➢ See also “Drive Sizing” on page 30, “Gigabyte (GB)” on page 120, “Kilobyte (KB)” on page 123, and “Megabyte (MB)” on page 129.

Topology

The physical layout and connection of devices. Fibre Channel supports three topologies, fabric, arbitrated loop, and point-to-point. Mylex controllers support specific failover topologies available through the Controller Parameters.

➢ See also “Controller Parameters” on page 69, “Fibre Channel” on page 43, “Failover Topology, Inactive Port” on page 37, “Failover Topology, Master/Slave” on page 38, “Failover Topology, Multiple Target ID” on page 39, and “Failover Topology, Multiport” on page 40.

Unconfigured Offline

SCOPE: This term is used by firmware versions 7.0 and greater. The term previously used was offline for physical disk drives. Refer to “Offline” on page 130 for more information.

A physical device state caused by one of the following conditions:

- A new physical device is inserted into a system
- A physical device is powered on in an unconfigured slot
- A physical device has a cleared or no configuration.

A physical device moves from an Unconfigured Offline state to an Online Optimal state when used in a new configuration or RAID device.

➢ See also “Online Optimal” on page 132 and “Physical Device States” on page 134.

Virtual Port

SCOPE: This term applies to Fibre Channel controllers, including the DACSF, DACFL, DACFF (FF), FFx, FF2, and FFx2.

A logical, rather than physical, port that consists of a separate, and possibly unique set of configuration data, and that is referred to and treated as a separate processing entity for initiator port login and Fibre Channel protocol command processing.

➢ See also “Failover Topology, Multiple Target ID” on page 39, “Fibre Channel” on page 43, and “Port” on page 135.

Warm Start

A warm start preserves data in the cache memory until the controller completes the startup process. The data is flushed from cache memory to the disk drives following a successful startup process.

➢ See also “Battery Backup Unit (BBU)” on page 12, “Cache” on page 111, “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, “Standard Data Caching” on page 63, “Write-Back Cache” on page 151, and “Write-Through Cache” on page 151

World Wide Name (WWN)

A 64-bit identifier, with a 60-bit value preceded by a 4-bit Network Address Authority Identifier used to uniquely identify nodes or ports.

➢ See also “Fibre Channel” on page 43, “Node” on page 129, “Port” on page 135, and “World Wide Name (WWN) Assignments” on page 65.

Write-Back Cache

A caching strategy whereby write operations result in a completion signal being sent to the host operating system as soon as the cache receives the data to be written. The target disk drive will receive the data at a more appropriate time in order to increase controller performance. In dual-active controller configurations with write-back caching enabled, the write data is always copied to the cache of the second controller before completion status is issued to the host initiator.

➢ See also “Cache” on page 111, “Cache Coherency” on page 17, “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, “Dual-Active Controller Configuration” on page 32, “Standard Data Caching” on page 63, and “Write-Through Cache” on page 151.

Write-Through Cache

A caching strategy whereby data is written to the disk drive before completion status is returned to the host operating system. This caching strategy is considered more secure, since a power failure is less likely to cause a loss of data; however, write-through cache results in slightly lower performance.

➢ See also “Cache” on page 111, “Cache Flush” on page 112, “Conservative Cache Mode” on page 22, “Dual-Active Controller Configuration” on page 32, “Standard Data Caching” on page 63, and “Write-Back Cache” on page 151.

INDEX

D

S

T

U

V

W