



ATTO Technology, Inc.

Diamond Storage Array V-Class

Installation and Operation Manual

© 2005 ATTO Technology Inc. All rights reserved. All brand or product names are trademarks of their respective holders. No part of this manual may be reproduced in any form or by any means without the express written permission of ATTO Technology Inc.

2/05

6.4

PRMA-0338-000

Contents

Preface

1.0 Product Overview	1
Features	
Fibre Channel model	
SCSI model	
2.0 Technical Overview	3
ATA disk drives	
ADXT Powers ATA to New Levels	
3.0 Installation Instructions	5
Three major steps are required to set up the Diamond Storage Array	
Step 1: physical setup	
Step 2a: set up the Ethernet connection	
Step 2b: connect to Ethernet	
If the Diamond Storage Array is attached to a DHCP server	
If the Diamond Storage Array is not attached to a DHCP server and you wish to change the defaults	
Step 3: configure drives	
3.1 Components	7
Floor model	
Rack mount	
3.2 Physical Setup	9
Floor Model	
Rack Mount	
General Instructions	
3.2.1 Connecting a Fibre Channel Array	11
Autoconfiguration	
Manual configuration	
3.2.2 Connecting a SCSI Array	13
To connect the SCSI Diamond Storage Array	
3.3 Determining Drive and Sled Designations	15
Numbering conventions	
Physical numbering	
Logical Numbering	
Examples	
Unique serial number for each LUN	

4.0 Accessing the Array	17
Command Line Interface	
ATTO ExpressNAV	
In-band SCSI over Fibre Channel	
RS-232 port	
Ethernet port	
SNMP	
4.1 In-band CLI Over Fibre Channel	19
I/O details	
4.2 Serial Port Access	21
4.3 Ethernet Access: Telnet and SNMP Protocols	23
To connect to the Ethernet port	
To use Telnet	
To use SNMP	
4.4 ATTO ExpressNAV: Browser-based Interface	25
Browser Compatibility	
To optimize ExpressNAV in Internet Explorer	
To open an ExpressNAV session	
To navigate ExpressNAV	
4.4.1 ExpressNAV Pages	27
Status	
Ethernet	
SNMP	
Serial Port	
Fibre Channel	
Storage Management	
RAID	
Clear Data	
Logical Units	
Partitions	
Zoning	
Rebuild	
Configuration	
Advanced	
To use the Advanced Page CLI commands	
Restart	
Help	
4.5 CLI: Interface via ASCII-based Commands	29
4.5.1 Summary of CLI Commands	31

4.5.2 General Use Commands	34
FirmwareRestart	
Help	
PartitionCommit	
SaveConfiguration	
SystemSN	
VerboseMode	
ZoneCommit	
4.5.3 Fibre Channel Configuration Commands	35
FcConnMode	
FcDataRate	
FcFairArb	
FcFrameLength	
FcFullDuplex	
FcHard	
FcHardAddress	
FcPortInfo	
FcPortList	
FcSCSIBusyStatus	
FcWWName	
4.5.4 Serial Port Configuration Commands	37
SerialPortBaudRate	
SerialPortEcho	
SerialPortHandshake	
SerialPortStopBits	
4.5.5 Ethernet Commands	39
EthernetSpeed	
FTPPassword	
IPAddress	
IPDHCP	
IPGateway	
IPSubnetMask	
SNMPTrapAddress	
SNMPTraps	
SNMPUpdates	
TelnetPassword	
TelnetTimeout	
TelnetUsername	
4.5.6 Diagnostic Commands	41
AudibleAlarm	
DiamondModel	
DiamondName	
DriveCopyStatus	
DriveInfo	
FcNodeName	
FcPortList	
FcPortName	
Help	
IdentifyDiamond	
Info	
LUNInfo	
PartitionInfo	

RAID5ClearDataStatus
RAIDRebuildStatus
SerialNumber
SledFaultLED
SMARTData
Temperature
VirtualDriveInfo
ZoneInfo

4.5.7 Drive Configuration Commands43

ATADiskState
AutoRebuild
ClearDiskReservedAreaData
DriveCopy
DriveCopyHalt
DriveCopyResume
DriveCopyStatus
DriveInfo
DriveSledPower
DriveWipe
IdeTransferRate
LUNInfo
LUNState
PartitionCommit
PartitionInfo
PartitionMerge
PartitionSplit
QuickRAID0
QuickRAID1
QuickRAID5
QuickRAID10
RAID5ClearData
RAID5ClearDataStatus
RAIDInterleave
RAIDHaltRebuild
RAIDManualRebuild
RAIDRebuildState
RAIDRebuildStatus
RAIDResumeRebuild
RebuildPriority
ResolveLUNConflicts
RestoreModePages
SledFaultLED
VirtualDriveInfo
ZoneAddDevice
ZoneAddHost
ZoneAddPort
ZoneClearAll
ZoneCommit
ZoneCreate
ZoneDelete
ZoneInfo
ZoneRemoveDevice
ZoneRemoveHost
ZoneRemovePort
ZoneRetrieve
ZoneState

4.5.8 Maintenance Services Commands	48
FcScsiBusyStatus	
FirmwareRestart	
MaxEnclTempAlrm	
MinEnclTempAlrm	
Temperature	
Zmodem	
ZoneRetrieve	
5.0 Configuring Drives	49
JBOD (Just a Bunch of Disks)	
RAID Level 0	
RAID Level 1	
RAID Level 10	
RAID Level 5	
Interleave	
Partitions	
Zones	
Hot Spare sleds	
Enhancing performance	
5.1 JBOD	51
To set up the JBOD configuration	
5.2 RAID Level 0	52
Sled-based versus disk-based	
To set up RAID Level 0 groups	
To remove RAID Level 0 groups from the array	
5.3 RAID Level 1	55
To set up RAID Level 1 groups	
To set up RAID Level 1 with Hot Spare sleds	
To remove RAID groups	
5.4 RAID Level 5	57
Configuring a fully-populated array	
To set up one RAID Level 5 group with one Hot Spare sled	
To set up two RAID Level 5 groups with two Hot Spare sleds	
Configuring a partially-populated array	
To set up one RAID Level 5 group with one Hot Spare sled	
Removing RAID groups	
5.5 RAID Level 10	61
To set up RAID Level 10 groups	
To remove RAID groups	
To set up RAID Level 10 with Hot Spare sleds:	
5.6 Rebuilding RAID Level Configurations	63
To reset LUN status	
To synchronize mirrored drives automatically	
Rebuild priority	
To synchronize mirrored drives manually	

5.7 RAID Interleave	65
To change the RAID Interleave parameter	
5.8 Creating Partitions	67
To create a partition	
To merge partitions	
5.9 Creating Zones	69
Principles of Zoning	
Factors to consider	
Status and Sense Data	
Configuring Zones	
To create a zone	
To remove zones	
To change current zones	
Other operations	
Errors	
6.0 Copying Drives	73
7.0 Updating Firmware	75
Updating firmware via the RS-232 serial port	
Updating firmware via the optional Ethernet card	
8.0 System Monitoring and Reporting	77
8.1 Troubleshooting	81
Windows 2000 special instructions	
Error Messages	
System Fault LED	
Command Line Interface messages	
ERROR. Wrong/Missing Parameters	
ERROR. Invalid Command. Type 'help' for command list	
ERROR. Command Not Processed.	
Audible Alarm	
Specific situations and suggestions	
If a drive fails to respond	
If a power supply fails	
To determine if the problem exists with the Host Interface Card or the connection	
If you can't access the array CLI via Ethernet	
If you do not see the appropriate number of LUNs on the host machine	
8.2 Resetting Defaults	83
Default	
Return to Default settings	
Factory Default	
To reset to Factory Defaults, firmware version 2.5.3 or higher	
9.0 Hardware Maintenance	85

9.1 Hot Swap Operating Instructions	87
Disk Drives	
The following method is the safest way to perform a hot swap of a drive	
Power Supplies	
Blower Assemblies	
To replace a blower assembly	
9.2 Optional Hot Spare Sled	89
To set up RAID Level 1 with Hot Spare sleds	
To set up RAID Level 10 with Hot Spare sleds	
To set up one RAID Level 5 group with one Hot Spare sled	
To set up two RAID Level 5 groups with Hot Spare sleds	
Glossary	i
Fibre Channel technology	
SAN technology	
SCSI protocol	
Appendix A ATA Disk Technology	iii
Appendix B Information command returns	iv
Driveinfo	
LUNInfo	
PartitionInfo	
ZoneInfo	
Appendix C Sample Zoning Command Sequences	x
First time configuration (after download)	
Simple	
Symmetric Model	
Asymmetric Model	
Combined Symmetric/Asymmetric Model	
Appendix D Product Safety	xiii
Safety compliances	
EMC specifications	
Radio and television interference	
Appendix E Specifications	xiv
Environmental and physical	
Rack mount dimensions	
Floor mount dimensions	
Appendix F Part numbers	xv
Appendix G Warranty	xvi
Manufacturer limited warranty	
Contact ATTO Technology, Inc.	

Preface

This guide will take the technology-savvy user through the installation and maintenance of the Diamond Storage Array.

The Diamond Storage Array was designed to meet your need for large amounts of easily accessible storage using proprietary Aggregated Data Transfer Technology (ADXT™) to merge the power of multiple, high performance ATA disk drives with the sustained data transfer rates required by sophisticated computer users.

Your comments help us improve and update our products. Contact us:

ATTO Technology, Inc.
155 CrossPoint Parkway
Amherst, New York 14068
(716) 691-1999 • voice
(716) 691-9353 • fax
<http://www.attotech.com/diamond>

ATTO Technology can also be reached via e-mail at the following addresses:

Sales Support: sls@attotech.com
Technical Support: techsupp@attotech.com

Disclaimer

Although reasonable efforts have been made to assure the accuracy of the information contained herein, this publication could include technical inaccuracies or typographical errors. Manufacturer 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 Manufacturer has been advised of the possibility of such damages). Any questions or comments regarding this document or its contents should be addressed to Manufacturer.

Manufacturer provides this publication as is, without warranty of any kind, either express or implied, including, but not limited to, the implied warranties for merchantability or fitness for a particular purpose.

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

1.0 Product Overview

The Diamond Storage Array offers up to 24 ATA disk drives in a rack mount or floor model configuration.

The Diamond Storage Array is ideally suited for data intensive applications such as web hosting, e-mail servers, on-line transaction processing, digital video editing, medical imaging and digital audio editing. Virtually any high performance computing system with a growing need for storage capacity can use the power of the array.

With the cost effective approach of using ATA disk drives, you can add more storage capacity as your needs grow without the costs of other disk storage technologies. You can also improve the performance and capacity of the array cabinet as technology progresses by simply replacing disk drive sleds and host interface cards.

The array is operating system independent and supports all popular computer hardware platforms and network environments.

Three interface options are available: a 1-gigabit Fibre Channel interface, a 2-gigabit Fibre Channel interface, and an Ultra160 SCSI interface.

The array is a fully populated, turnkey solution with drives pre-installed. It is fully supported by a highly trained customer service and engineering staff.

The Diamond Storage Array uses Aggregated Data Transfer Technology (ADXT™) to merge the performance of multiple ATA drives to achieve sustained, full bandwidth data transfer rates. ADXT provides end users with the power and sophisticated data control needed to take moderately priced ATA disk drives, combine them in a disk storage array, and power them to the performance levels of SCSI or Fibre Channel disk arrays.

Features

- Up to 7.2 Terabytes initial configuration (expandable with future drive technology)
- 24 ATA disk drive capacity
- Aggregated Data Transfer Technology (ADXT™) for high performance/scalability
- Ultra ATA 100 megahertz (Mhz)
- JBOD, RAID Level 0, RAID Level 1, RAID Level 10 and RAID Level 5 configurable
- Partitioning capability
- Zoning capability compatible with third party servers, switches and with deliverables from industry standards organizations.
- Hot spare sleds: replace degraded sleds with spares on the fly using software
- Staggered drive spin-up to reduce peak power demand



- Tagged command queuing to process up to 255 simultaneous data requests
- RS-232 management for local management control; Ethernet option available for setup connection only
- ExpressNAV™ browser-based user interface
- Two power supplies capable of 85-264 V (rated 100-240V AC) operation (340 watts each)
- UL, TUV and CE marked and compliant
- Internal thermal and power management
- Redundant hot swappable power supplies with integrated thermal and power management
- Floor model cabinet or 19" 3U rack mount


Fibre Channel model

- 2 gigabit Fibre Channel Port (single or dual channel)
- SFP-based Fibre Channel interface supports long wave and short wave optical cables
- Built-in hub for daisy-chaining
- Up to 9,500 I/Os per second per port
- Up to 240 MB/sec. sustained Fibre Channel transfer rates per interface

SCSI model

- Ultra 160 SCSI bus.
- Dual stacked VHDCI connectors for daisy-chaining and termination
- SCSI Target ID selection switch
- Support for single-ended and LVD SCSI
- No onboard termination

Exhibit 1.0-1 Back of rack mount model, Diamond Storage Array. Left: 2 gigabit Fibre Channel. Right: SCSI.



2.0 Technical Overview

The Diamond Storage Array uses Aggregated Data Transfer Technology (ADXT) to achieve the high data transfer performance you need. ADXT merges the performance of multiple ATA drives together to achieve sustained, full bandwidth Fibre Channel data transfer rates.

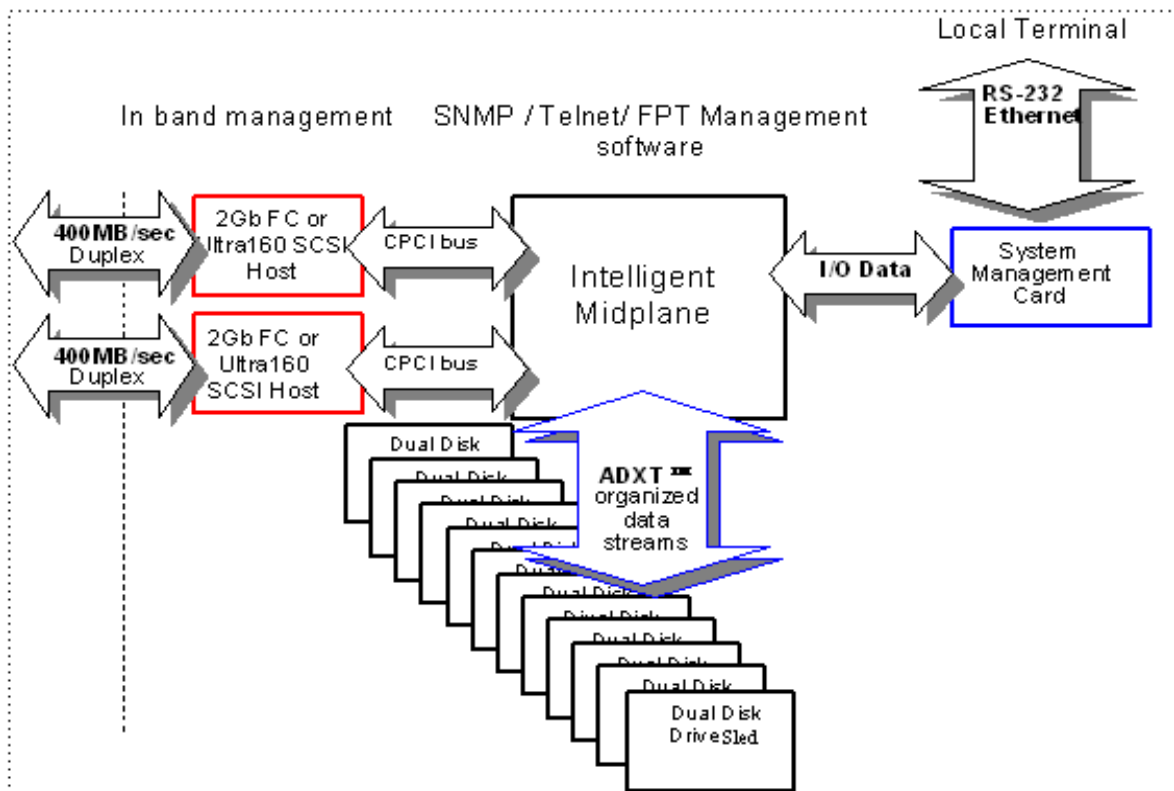
Unlike other storage arrays which use expensive SCSI or Fibre Channel disk drives to achieve performance, the Diamond Series uses lower cost ATA disk drives combined with an intelligent midplane to create a storage array with price and performance characteristics.

The intelligent midplane contains hardware and software which provide the proprietary ADXT, a switched data management and data movement technology. The storage array delivers faster sustained

data transfer rates as well as impressive I/Os per second.

The array is made up of dual SCSI or Fibre Channel host interface cards, the intelligent midplane, a system management card, and 12 independent disk drive sleds containing up to 24 ATA disk drives.

Exhibit 2.0-1 Data pathways and architecture for Fibre Channel operation



ATA disk drives

ATA disk drives were known originally as Integrated Drive Electronics (IDE), a low end disk interface. The original IDE interface was low performance, single threaded (no simultaneous I/O requesting), contained minimal error detection and was unsuitable for

computer applications requiring high performance and high reliability.

As IDE was refined and acquired important capabilities, its name was changed to ATA, Advanced Technology Attachment.

- UltraDMA transfer protocol similar to high performance SCSI disk protocol operating at 66 MB/sec.
- Double-clocking of data transfers, doubling disk data transfer rates
- CRC (Cyclic Redundancy Check) code allowing full error detection and data reliability
- Multi-threaded I/O support
- Overlapped Command Support: allows commands to be simultaneously active on multiple drives on the same ATA bus.
- Command Queuing which allows simultaneous multiple read/write commands to be sent to each drive, reducing command overhead and allowing the drive to service commands in the most efficient manner: similar to the SCSI feature Tagged Command Queuing.
- Faster drive speeds (5400/7200 RPM) with higher media transfer rates
- A communication protocol and interface with a fundamental lower cost structure than SCSI or FC interfaces

ATA disk drives operate at performance and data integrity levels similar to those that were previously available only on SCSI or Fibre Channel disk drives.

ADXT Powers ATA to New Levels

The original notion of RAID was to build high capacity, reliable storage subsystems using large numbers of inexpensive disk drives. Thus its original definition: *Redundant Array of Inexpensive Drives*.

Over time that definition became *Redundant Array of Independent Drives* and the inherent cost advantage in a RAID system was lost.

Intelligent Midplane The heart of the Diamond Series storage array is the intelligent midplane with ADXT to sum or *aggregate* the data rates of individual ATA disk drives to create high data transfer rates. This

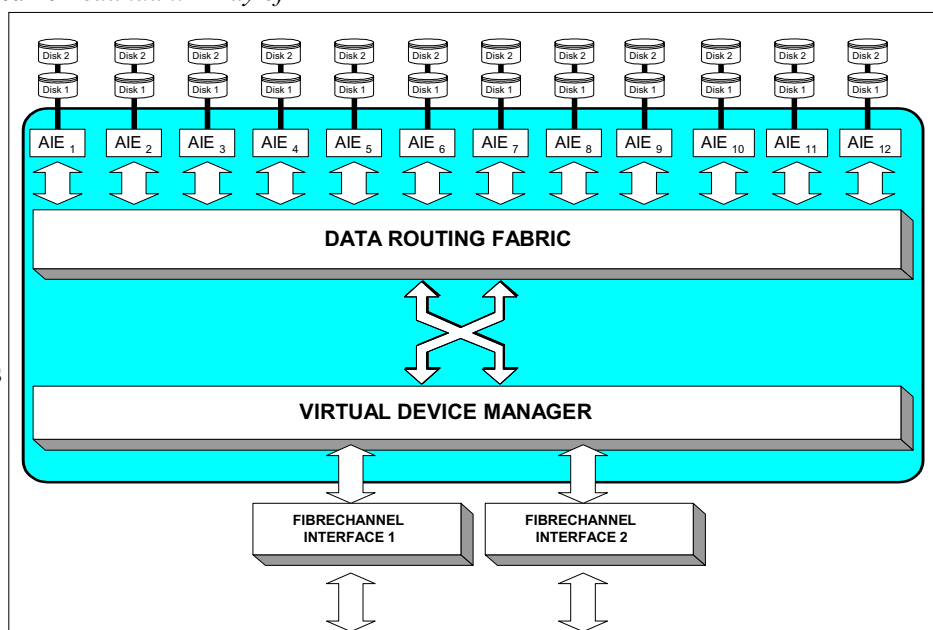
technology enables features such as serverless backup, advanced error protection, metadata storage techniques, virtualization software, thermal management and advanced enclosure services.

The midplane contains a combination of custom Application Specific Integrated Circuits (ASICs), processors and proprietary embedded software divided into three main processing sections which handle the data being read or written to the ATA disk drives from the Fibre Channel or SCSI host interfaces. The Virtual Device Manager (VDM), Data Routing Fabric (DRF) and ATA Interface Engines (AIE) organize data streams for storage or retrieval.

Virtual Drive Manager Data is accessed through Virtual Drives using an implementation of the standard SCSI protocol controlled by the Virtual Device Manager.

Data Routing Fabric incoming or outgoing data is routed between the ATA Interface Engines (AIE) and the Fibre Channel or SCSI interface by the custom Data Routing Fabric ASIC, a high speed, low latency transfer fabric with more than 2 GB/sec. of bandwidth supported by up to 512 Megabytes of memory.

ATA Interface Engine (AIE) The interface to each pair of drives is through a custom ATA Interface Engine ASIC. The AIE implements the typically software-intensive ATA interface completely in silicon. Each AIE contains a dedicated ATA protocol processor to completely automate command and protocol processing.



3.0 Installation Instructions

If you are familiar with the Diamond Storage Array, Fibre Channel, SCSI and RAID configurations, you may set up and configure the array using these instructions. You will find details, illustrations and other guidance for more involved operations and special cases in the rest of this manual.



CAUTION

Before configuring the Diamond Storage Array, ensure that any computer data to be stored on the array is properly backed up and verified. The Manufacturer is not responsible for the loss of any data stored on the array under any circumstances and any special, incidental, or consequential damages that may result thereof.

Three major steps are required to set up the Diamond Storage Array

- 1 Physically set up Diamond Storage Array
- 2 Connect to Ethernet
- 3 Configure the drives

Step 1: physical setup

- 1 Make sure the Diamond Storage Array is mounted properly and has adequate air flow around it.
- 2 Insert the appropriate connector into the interface card in the back of the Diamond Storage Array.
- 3 Connect the Fibre Channel or SCSI cable from your host computer system to the connector. To use ExpressNAV browser-based management interface and configure your Diamond Storage Array, you must connect to the Ethernet port.
- 4 To set up the Ethernet connection: connect a cross-over cable (for a direct connection to a PC) or regular network cable from a network device to the RJ45 Ethernet port on the Ethernet management card on the front of the Diamond Storage Array.

Step 2a: set up the Ethernet connection

The Diamond Storage Array supports service operations over the RS-232 serial port using standard terminal emulation software available with most systems.

- 1 Connect a DB-9 null modem serial cable between the port marked RS-232 on the front of

the Diamond Storage Array and the computer's serial port. The cable must be no longer than three meters.

- 2 Make sure the power switches on the power supplies on the rear of the unit are in the Stand-by position.
- 3 Plug in the power cords to the back of the unit, then into an appropriate power source (100-240 VAC).
- 4 Reboot your host computer system.
- 5 Press the Stand-by power switch for each power supply on the Diamond Storage Array to the ON position.
- 6 Upon successful power up and diagnostics, the unit displays the POST (power up self test) information.

The Diamond is now in Command Line Interface mode. You may modify the setup of the Diamond Storage Array using the CLI (refer to *CLI: Interface via ASCII-based Commands* on page 27), but the easiest method to configure the array is by using ATTO ExpressNAV software, a graphical user management interface accessed through a standard Internet browser. Refer to *ATTO ExpressNAV: Browser-based Interface* on page 23.

Step 2b: connect to Ethernet

If the Diamond Storage Array is attached to a DHCP server

- 1 At the **Ready** prompt after **POST** (refer to Step 6 above), type **set IPDHCP enabled**
- 2 Type **SaveConfiguration Restart**
- 3 At the **Ready** prompt after **POST** (see above), type **get IPAddress**
- 4 Enter this address into your browser.
- 5 The ATTO ExpressNAV screen appears. Log in using the Telnet defaults:
Username: Telnet
Password: Diamond
- 6 Follow the screens to find information about the array or to configure the array from the factory-

default settings. The Diamond Storage Array may be set up in a JBOD, RAID Level 0, 5 or 10 configuration with or without Hot Spare sleds

If the Diamond Storage Array is not attached to a DHCP server and you wish to change the defaults

- 1 At the **Ready** prompt after **POST** (see above), type **set IPAddress [desired IP address]**.
- 2 Type **set IPSubnetMask [desired IP Subnet Mask]**
- 3 Type **set IPGateway [desired IP Gateway]**
- 4 Type **SaveConfiguration Restart** to save the configuration and restart the Diamond Storage Array
- 5 After the powerup and POST complete, type the IP address from step 1 above into your browser.
- 6 The ATTO ExpressNAV screen appears. After logging in (refer to **Step 2b: connect to Ethernet**, Step 5 on page 5), follow the screens to find information about the array or to

configure the array from the factory-default settings. The Diamond Storage Array may be set up in a JBOD, RAID Level 0, 1, 5 or 10 configuration with or without Hot Spare sleds.

Step 3: configure drives

The simplest way to set up configurations is to use the ExpressNAV interface. Refer to *ATTO ExpressNAV: Browser-based Interface* on page 23 for more information on the interface. After logging in, follow the screens to find information about the array or to configure the array from the factory-default settings.

You may the CLI to set up RAID, partitions and zones.



Note

All arrays using RAID level 10 and Hot Spare sled options must be fully populated.

3.1 Components

The Diamond Storage Array has been designed to be easy to use, maintain and upgrade. It features a durable steel outer case and modular components in either a floor or a rack mount model.

Immediately upon receipt, check the shipping carton for damage from mishandling. Contact us at once via the means that is easiest for you (refer to *Warranty* on page xvi) if the carton has been mishandled or displays any signs of damage.

The front of the Diamond Storage Array provides access to the management card and disk drive sleds. The rear of the unit holds the host interface cards, power supplies and blower assemblies.




CAUTION

All modular components must be replaced by qualified personnel only.

Floor model

The management system card is at the top front of the case. At its center is a DB-9 serial RS-232 port, a connection for setup, monitoring and upgrade of the unit from any computer system with an RS-232 interface. The optional 10/100 BaseT Ethernet management services card enables Telnet-based monitoring and



management. It also provides the ability to update the firmware in the array via FTP.

LEDs to the port's right indicate fault, unit ready, host interface cards A and B installation status, and the power status for each power supply.

Below the management system card are individual disk drive sleds which also have LEDs for each drive's status. Each sled contains two hard drives. Up to 24 hard drives may be installed on the 12 sleds. Empty bays should be covered by blank faceplates or empty sleds. Access is provided by loosening two screws and gently pulling on the sled handle.

On the rear of the unit are blowers which support hard drive, cabinet and power supply cooling. The blowers are held in by removable screws. Correct operation is displayed by a LED at the top of each panel.

The power supplies for the array, also in the rear of the unit, are accessible by loosening two screws and pulling on the power supply module handle. The power standby on/off switch is at the top of each module. A yellow LED indicates *caution* and a green LED indicates *on*. The power cord socket is at the bottom of each power supply.

Between the power supplies and blower assemblies are two slots that hold the Host Interface cards. The HIC is the connection point into the array and is available in three options: 1-Gigabit Fibre Channel, 2-Gigabit Fibre Channel, or Ultra 160 SCSI. Host Interface cards have *fault* and *on-line* or *fault* and *activity* LED indicators, depending on the model.

SCSI The SCSI card faceplate has a rotary binary-coded hex switch to set the SCSI ID of the array. The SCSI card also has an in channel, to connect via cable to the unit's communication source, and an out channel, available for daisy-chaining arrays together or to complete termination using an external LVD terminator.

Rack mount


The system management card is at the left front of the case. At its center is a DB-9 serial RS-232 port

which allows a connection for setup, monitoring and upgrade of the unit from any computer system with an RS-232 interface. The optional 10/100 BaseT Ethernet management services card enables Telnet-based monitoring and management. It also provides the ability to update the firmware in the array via FTP. LEDs farthest to the left indicate fault, unit ready, Host Interface cards A and B installation status, and the power status for each power supply.

To the right of the system management card are individual disk drive sleds which also have LED indicators for each drive's status. Each sled contains two hard drives. Up to 24 hard drives may be installed on the 12 sleds. Empty bays should be covered by blank faceplates or filled with empty sleds to promote effective cooling. Access is provided by loosening two screws and pulling on the sled handle.

In the rear of the unit are the blower assemblies which support hard drive, cabinet and power supply cooling. Correct operation is displayed by a LED at the top of each panel. The blowers are held in place by removable screws.

The power supplies for the array are accessible by loosening two screws and pulling on the power supply module handle. The power standby on/off switch is at the top of each module. A yellow LED indicates *caution* and a green LED indicates *on*.



The power cord socket is at the bottom of each power supply.

Between the power supplies and blower assemblies are two slots that hold the Host Interface cards. The HIC is the connection point into the array and is available in 2-Gigabit Fibre Channel, or Ultra 160 SCSI. Host Interface cards have *fault* and *on-line* or *fault* and *activity* LED indicators, depending on the model.

SCSI The SCSI card faceplate has a rotary binary-coded hex switch which allows you to set the SCSI ID of the array. The SCSI card also has an in channel, to connect by a cable to the unit's communication source, and an out channel, available for daisy-chaining arrays together or to complete termination using an external LVD terminator.

3.2 Physical Setup

The Diamond Storage Array is shipped completely assembled with two 120 VAC power cords for use in the United States and Canada.

Immediately upon receipt, check the shipping carton for damage from mishandling. Contact us at once by the means easiest for you (refer to *Warranty* on page xvi) if the carton has been mishandled or displays any signs of damage.

Floor Model

The Diamond Storage Array is heavy (about 92 pounds fully loaded) and requires two people to lift and carry it safely. Place the array on a level surface and make sure there is adequate space in the front and back of the unit for proper cooling and airflow. Continue with the general instructions.

Rack Mount

The array fully loaded is heavy (about 86 pounds). The unit should be handled with care and requires two people to lift, carry and/or install it safely.

The array can be mounted via several different methods in a 19" rack with 3U (5.25") of vertical space required. Air flow should not be restricted in any way.

Installation in a rack may create a differential between the room ambient temperature and the internal ambient temperature in the rack. While the maximum internal operating temperature of the array is 47°C, you should not run the system at the maximum temperature for extended periods. Therefore, ensure that the room ambient temperature is kept below 30°C

for best operation.

Each side of the rack mount array chassis has three pairs of mounting holes. One is located near the front of the rack, one near the unit's center of gravity, and one near the rear of the rack. The holes accommodate 10/32 screws but the screws can protrude no farther than .375 inches into the rack.


Spaced rail pairs in some rackmount cabinets.

You can mount the array using two sets of rail pairs spaced to accommodate the overall length of the unit (approximately 23 inches). Mount using the rack mount brackets on the front and rear of the unit fastened to the rail pairs using 10/32 pan head screws with lock and flat washers.

Rack mount cabinets with stationary shelf or tray system. The shelf or tray must be able to support at least 125 pounds. The shelf or tray must be installed and secured to the rack before installing the array. Secure the front of the array to the rack with 10/32 screws, locks and flat washers.

Sliding shelf or tray type systems should never be used under any circumstances.

Two point open rack system. The rack must be strong enough to support the array properly. Mounting brackets should be moved to the



centermost mounting holes and secured using 10/32 screws.



CAUTION

Do not mount multiple arrays on a two-rail rack or mount the array above the midpoint of a two-rail rack system. Do not mount the array on any kind of rail-type system. The array is too heavy and does not have the proper hole pattern for rails.




Note

Insure the array has adequate air flow.

General Instructions

- 1 Insert the proper connector into the Host Interface Card in the back of the array. (refer to



Connecting a Fibre Channel Array on page 11 for Fibre Channel and Connecting a SCSI Array on page 13 for SCSI).


- 2 Connect the cable (Fibre Channel or SCSI) from your host system to the Host Interface Card connector on the back of the array. The cable you use depends upon your application, the environment and distance.

- 3 Make sure the power switches on the power supplies on the rear of the unit are in the stand-by position. Plug in the power cords to the back of the unit, then into an appropriate power source (100-240 VAC). The power source must be connected to a protective earth ground and comply with local electrical codes. Improper grounding may result in an electrical shock or damage to the unit.
- 4 Press the stand-by power switch for each power supply to the ON position. When the green power LED on the back of the unit is lit, the power supply is fully operational and delivering power to the system. The power LED on the front of the array lights while the firmware executes.

When the power is turned on, the LEDs on the front of the array flash twice. Drives spin up in groups of three every one to two seconds. The individual LEDs blink. After all available drives have spun up, the individual drive LEDs stay lit. When all available drives are operational, the *ready* LED on the top front panel of the management card remains lit.


- 5 Reboot your computer
- 6 Determine the best configuration for your needs (i.e. JBOD, RAID, etc.) and refer to the rest of this manual for configuration information.

Exhibit 3.2-1 Back side of a rack mount array.



3.2.1 Connecting a Fibre Channel Array

The Diamond Storage Array supports up to two Fibre Channel Host Interface Cards (HIC). Physical connections and CLI commands contribute to the Fibre Channel topology.



2 Gb HIC uses
2 SFPs to connect up
to 2 FC cables

The cable you use depends upon your application, the environment and the distance required for your storage area network.

To comply with FCC and CE EMI for the 2-gigabit

Host Interface Card, use fiber optic cables.

Exhibit 3.2-1 Fibre Channel cable options

Cable length	Cable type	Cable size
<10 meters	Unequalized copper	
>10 <30 meters	Equalized copper	
Up to 175 meters	multi mode fiber optic	62.5 micron
Up to 500 meters	multi mode fiber optic	50 micron
Up to 10 kilometers	single mode fiber optic	9 micron

The Diamond Storage Array may have two Fibre Channel Host Interface Cards (HIC). In 2 gigabit Fibre Channel arrays, each HIC is connected by a Fibre Channel cable via a SFP (small form factor pluggable) module into a point-to-point or loop Fibre Channel topology.

Install the SFP according to the manufacturer's instructions.

Each HIC has two ports and an on-board hub. Each port has an SFP module to connect to Fibre Channel. Each HIC is independent of the other, so that one may be connected into a point-to-point topology and the other into a loop. However, if one port in a HIC is connected into a point-to-point topology, the other port cannot be used.

One of the advantages of using loop topology for Fibre Channel connections is that it allows arrays to be daisy-chained together.

Autoconfiguration

The array automatically determines which HICs are installed and if they are in loop or point-to-point topologies.

If you wish to see how the unit has been set up, go to the CLI commands and type Info or go to the **Status** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25). The return displays the data rate and connection mode for each HIC (FC 0 and FC 1).

Manual configuration

You may manually configure the array using CLI commands (refer to *Fibre Channel Configuration Commands* on page 35.) or access the **Fibre Channel** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).

- **FCConnMode** specifies the Fibre Channel topology for both HICs on an array. Options are loop only (loop), point-to-point only (ptp), loop preferred (loop-ptp) or point-to-point preferred (ptp-loop).

Loop connects to either an FC-AL arbitrated loop or a fabric loop port (FL_Port) on a switch.

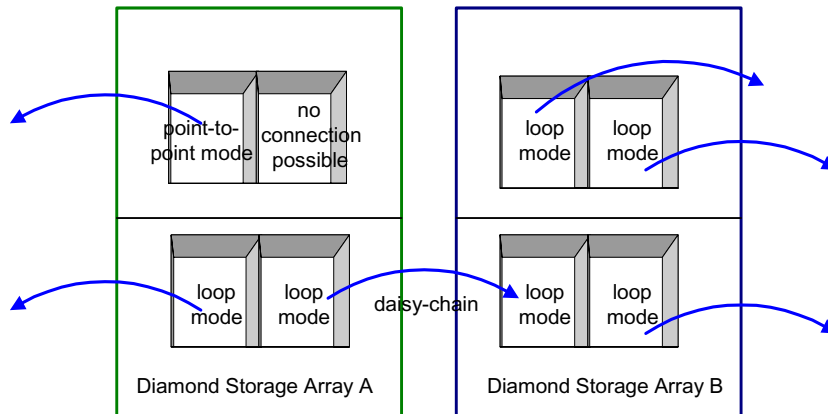
Point-to-point (ptp) connects to a direct fabric connection, such as an F port on a switch.

Loop-ptp allows the array to determine what kind of connection to use, but tries to connect in loop mode first, then point-to-point mode.

Ptp-loop allows the card to determine what kind of connection to use, but tries to connect in point-to-point mode first, then loop mode.

- **FcDataRate** specifies the Fibre Channel data rate at which both HICs on a Diamond operate. Choices are 1 gigabit, 2 gigabit and autodetection.

Exhibit 3.2-2 Possible 2 gigabit Fibre Channel physical connections depending on which Fibre Channel connection mode has been selected.




3.2.2 Connecting a SCSI Array

The SCSI Diamond Storage Array uses a VHDCI connector and SCSI cables to connect to a host. It automatically detects the type of Host Interface Card it is using without any intervention.

To connect the SCSI Diamond Storage Array

- 1 Insert a SCSI VHDCI connector into the Host Interface Card in the back of the array. If the SCSI array is the last device on the bus, you
- 2 The SCSI Host Interface Card has a rotary binary-coded hex switch which allows you to set the SCSI ID of the HIC. Be sure the selected ID is different from all other SCSI devices on the bus.



must attach a VHDCI terminator to one connector of the SCSI Host Interface Card or connect a cable between the second connector and the next device on the SCSI bus.




Note

If slower devices are connected on the same SCSI bus as the Ultra 160 array, the bus communicates at the rate of the slowest device.

Exhibit 3.2-1 SCSI cable options.

	Bus speed, MB/sec. max.	Bus width, bits	Max. bus lengths, meters		Max. device support
			Single-ended	LVD	
SCSI-1	5	8	6	-	8
Fast SCSI	10	8	6	-	8
Fast Wide SCSI	20	16	6	-	16
Wide Ultra SCSI	40	16	3	-	4
Wide Ultra SCSI	40	16	1.5	-	8
Wide Ultra 2 SCSI	80	16	-	12	16
Ultra 3 or Ultra160 SCSI	160	16	-	12	16

Exhibit 3.2-2 SCSI interface cards: left without terminators attached; right with a terminator attached.



3.3 Determining Drive and Sled Designations

The Diamond Storage Array has been designed with 12 sleds, each holding two drives. The easiest way to configure an array is to use all the drives on all the sleds. The firmware uses a numbering system to determine which drives and sleds it is affecting.

All sled slots should be filled contiguously, starting with the first slot next to the management card.

When configuring a array with fewer than 12 drive sleds, you must consider several factors: RAID level, number of physical drives/sleds and the end configuration you are trying to achieve. Review the information about each configuration to determine how each configuration would be affected by using fewer sleds.

Numbering conventions

The Diamond Storage Array with firmware version 3.1 and newer uses a unique numbering convention to orient its drives and sleds to the controlling firmware. Older versions do not use this convention. (refer to *Updating Firmware* on page 75 for information about updating the array firmware.)

Physical refers to the physical drives in the array, the hardware that actually exists in a physical sense.

Logical (or virtual) refers to what the host operating system recognizes as an entity. Two physical drives may be seen as one logical drive by the operating system.

Logical disks do not always map one-to-one with physical disks. In RAID configurations, for example, several physical disk drives (or portions of several physical drives) are grouped into a logical disk, call a RAID Group or a Logical Unit

(LUN). Each RAID group is broken into logical blocks of 512 bytes each, numbered 0 through n (the Logical Block Number or LBN). A 100 GB LUN has approximately 200,000,000 logical blocks.



A RAID group is also referred to as a Virtual Drive.

A LUN is associated with a RAID group or Virtual Drive unless you are using partitions. If you have partitions, a LUN is associated with each partition. A RAID Group or Virtual Drive may then have multiple partitions or LUNs.

Physical numbering

The Diamond RAID Storage Array contains

- Up to 24 physical disk drives
- Two drives mounted on 12 physical drive sleds
- Sleds are numbered 1-12, starting at the top (floor units) or the left (rackmount units).
- Each sled is connected to its own internal ATA bus with two disk drives, numbered 1 and 2.
- Two green LEDs, labeled Drive 1 and Drive 2, indicate activity for the two drives. They remain solidly lighted when there is no activity.



Logical Numbering

Logical numbering depends on the RAID configuration of the storage array. Current firmware includes RAID Level 0, RAID Level 5, RAID Level 1, RAID Level 10, Zoning and Partitioning capability and hot spare sleds. The default configuration is QuickRAID0 0 or JBOD (Just a Bunch of Disks), in a single zone.

Examples

JBOD mode uses 24 LUNs. Each LUN is equivalent to one physical drive. The array can operate with several sleds missing, but the empty sled(s) are treated as offline and cannot be configured.

RAID Level 0 (QuickRAID0 [n]) With a fully populated array, RAID 0 may be configured as 1, 2, 3, 4, 6, or 12 LUNs. As RAID0 1, all 24 physical drives are configured as a single stripe group or LUN. You may also configure two LUNs of 12 drives each, three LUNs of eight drives each, four LUNs of six drives each, six LUNs of four drives each and 12 LUNs of two drives each. (See Exhibit 5.2-3 in *RAID Level 0* on page 52). The command assumes there are 24 drives available to configure the specified number of LUNs.

RAID Level 1 (QuickRAID1) has no options: the array can be configured into six groups, LUNs 1-6, with each LUN containing two physical sleds; each sled in the LUN would be a mirror image of the other sled in the LUN.

RAID Level 10 (QuickRAID10) The QuickRAID10 command first creates six mirrored groups, then stripes them into groups of one, two or three RAID 10 groups. Each group is a LUN. Issuing QuickRAID10 2, the 12 physical sleds are configured as six mirrored pairs, then striped into two LUNs.

RAID Level 5 (QuickRAID5) with a fully populated Array may be configured as 1, 2, 3 or 4 LUNs. As QuickRAID5 1, all 24 physical drives are configured as a single RAID 5 LUN. You may also configure two LUNs of 12 drives each, three LUNs of eight drives each, or four LUNs of six drives each. RAID Level 5 parity reduces the usable capacity of each LUN by the equivalent of one drive sled.

Unique serial number for each LUN

Each LUN in a system has a unique 24-character serial number which is updated when a system configuration changes. It is stored in a Device Association Table on each drive.

When a drive or sled is replaced in QuickRAID configurations that allow for hot swap, a new serial number is computed and is stamped onto all member drives of the RAID group. The CLI command LUNInfo or the **Logical Units** page of the ExpressNAV interface displays the serial number for each drive.

Exhibit 3.3-1 Format of the 24 characters of the unique serial number for each LUN

Character Placement	Valid values	Example of unique serial number for a JBOD configuration: 1231231231231231231XAC0
0-19	Any ASCII char	Drive Serial Number
20	1 for RAID1 A for RAID10 0 for RAID0 5 for RAID5 X for JBOD	RAID Configuration Character
21	A – Z	LUN Configuration Iteration Character: starts at A and increments every time a member drive is replaced in a RAID Group. Reverts back to A any time the Generation Character is updated
22	A – Z	Generation Character: starts at A and increments every time a new configuration is stamped on the system. When it reaches Z, rolls over to A .
23	0	Reserved for future use

4.0 Accessing the Array

Communicate with the Diamond Storage Array through the Fibre Channel link, the RS-232 port or the Ethernet port using Command Line Interface commands or ATTO ExpressNAV, an integrated user management console.

You may configure and tune the Diamond Storage Array for different environments and applications, update the firmware, monitor internal power and temperature status, report on hardware diagnostics and log failures.

Three avenues are available:

- In-band SCSI over Fibre Channel and over SCSI
- RS-232 port
- Telnet and SNMP over Ethernet

The following chapters describe how to access the array and use the Command Line Interface or ATTO ExpressNAV, an integrated user management console.

Command Line Interface

The CLI provides access to the array through ASCII command lines.

An initial display, after powering up the unit or restarting the firmware, contains the information in Exhibit 4.0-1. Once the initial display is complete, with the word **Ready**, you are in the CLI mode.

Type **Help** to display a list of all commands available.

ATTO ExpressNAV

ATTO ExpressNAV is an integrated configuration tool accessible through an Ethernet connection. Platform independent, ExpressNAV contains all the current capabilities of the CLI in a user-friendly GUI console. A menu on each page provides access to information and configuration operations.

Refer to *ATTO ExpressNAV: Browser-based Interface* on page 25 for more information on the program.

In-band SCSI over Fibre Channel

In-band SCSI commands (Write Buffer and Read Buffer) may be issued to the array to manage configuration via two mechanisms:

- In-band CLI over SCSI, where ASCII CLI commands, may be issued via Write Buffer. All CLI commands except Zmodem are supported.
- ID/value, where the application program uses a SCSI CDB (command descriptor block) to select the buffer ID of the configuration parameters to be affected, and the new value of the parameter. Most configuration options are available.

RS-232 port

The array provides remote service operations over the RS-232 serial port using standard terminal emulation software available with most systems.

Set the following serial parameters in your terminal program:

- Bits per second: 115200
- Data Bits: 8
- Parity: None
- Stop Bits: 1
- Flow Control: None.
- Terminal type: ASCII
- Echo: on.

Ethernet port

The 10/100 BaseT Ethernet port provides Telnet- or SNMP-based monitoring and management.

The default IP address is 10.0.0.1; the default subnet mask is 255.255.0.0. To change the defaults, first be configure the array for the network using the RS-232 port to establish the correct IP address. The management port provides TCP/IP-based monitoring and management.

SNMP

SNMP, or Simple Network Management Protocol, is an application layer protocol that allows network devices to exchange management information. Through a combination of standard and custom MIBs (Management Information Base), the array provides status and error

indications to an SNMP server, allowing the array to be managed with other devices in a complex system through a common interface.

Use CLI commands to configure up to six unique Trap addresses. A trap is a way for the array to tell the SNMP server that something has happened.

Exhibit 4.0-1 The POST information displays after boot of the Diamond Storage Array.

```
Diamond Storage Array
(c) 2004 ATTO Technology, Incorporated.

Firmware version 5.40 release date Mar 30 2004, 10:43:06 Build 021G

Power-On Self-Test Completion Status: GOOD
128 Megabytes of RAM Installed.

Interface Port 0 is not installed.
Interface Port 1 is 1.0624 Gb/s Fibre Channel.

Interface 0 World Wide Name = 20 00 00 10 86 10 02 DC
Interface 1 World Wide Name = 20 00 00 10 86 10 02 DC

Diamond Array Serial Number = "MIDP100197"
Diamond Array Name = "      "
System Serial Number = ""
Active Configuration      = ATTO
DiamondClass               = (V) 86
Internal Temperature = 23 C [5 - 47]
ErrorLog Contents: NO ERRORS
For help, type HELP.

Ready.
```

4.1 In-band CLI Over Fibre Channel

In-band Command Line Interface (CLI) configures and manages the Diamond Storage Array using SCSI-based CLI commands over a Fibre Channel port connection.

In-band CLI allows a programmer to configure the Diamond Storage Array while it is moving data. Using a programmer's interface, CLI commands as described previously in this manual may be implemented.

In-band CLI is implemented as part of LUN 0. It uses a different LUN than the array, and reports as a Storage Enclosure Services (SES) device (device type 0x0D).

LUN 0 is visible on all Fibre ports but is actually a single unit. The default value for LUN 0 is 0x00.

LUN 0 must be reserved for each Write Buffer/Read Buffer pair, using the SCSI Reserve command to insure integrity of the in-band CLI session.

- 1 An initiator (host) sends a SCSI Reserve command to LUN 0.
 - If LUN 0 is not reserved by another initiator, LUN 0 is now reserved and available to begin a new CLI session.
 - If the array configuration is reserved by a different CLI session (i.e. serial or Telnet), the in-band session does not allow modifications of the array configuration. If you try, the results buffer of LUN 0 returns:

```
Process X has the configuration reserved.
```

```
ID of this session = Y
Ready.
```

- 2 The initiator issues a SCSI Write Buffer command to LUN 0. A Write Buffer command must be accompanied by an ASCII buffer representing the CLI command string such as

```
set DiamondName Omega1
```

LUN 0 executes the command line and create feedback in the form of ASCII characters into a buffer. This buffer is 8 KB and circular. Retrieve the results by issuing a Read Buffer command before issuing another Write Buffer command.

- 3 A subsequent Write Buffer command executes the new command line and overwrite the previous results in the buffer with new results.
- 4 LUN 0 can be released by issuing a SCSI Release command to the LUN after each Write/Read Buffer pair, or multiple Write/Read Buffer pairs.

Initiator (Host)	Diamond Storage Array
Reserve LUN 0	return: "ok"
Write Buffer LUN 0 bid 'AA' "get Temperature"	executes the CLI command, stores output in buffer
Read Buffer LUN 0 bid 'AA'	return:"Temperature=28C\r\nReady.\r\n0"
Release LUN 0	return: "ok"

I/O details

The buffer sent to the Services LUN during the data out phase of a Write Buffer command must be:

- ASCII data
- maximum 80 bytes length
- terminated with either a carriage return character (0x0D), line feed character (0x0A) or NULL character (0x00)
- Characters following the first carriage return character, line feed character or NULL character are ignored.

The buffer retrieved from the Services LUN during the data-in phase of a Read Buffer command:

- ASCII data
- 8 KBytes (8192 bytes) in length
- terminated with a NULL character (0x00)
- Characters following the NULL character are meaningless.

A CHECK_CONDITION, INVALID_PARAMETER_IN_CDB is returned to an initiator that specifies an incorrect Buffer ID, Mode,

Length or Buffer Offset. The Mode is always Data (0x2), the Buffer ID is always 0 and the Buffer Offset is always 0.

Exhibit 4.1-1 The SCSI command process: reserve the Diamond Storage Array, send the command, release the Diamond Storage Array.

Initiator/Host	Diamond Storage Array
Goal: reserve the Diamond Storage Array for an in-band CLI command	
SCSI cdb: Reserve LUN 0	=>
	<= SCSI success
Goal: retrieve the Diamond Storage Array temperature via in-band CLI	
1. Issue the command:	
SCSI cdb: WriteBuffer LUN 0, bid='AA', "get Temperature\n"	=> places "Temperature=28C\n\r" into the read-data buffer
	<= SCSI success
2. Retrieve the results:	
SCSI cdb: ReadBuffer LUN 0, bid='AA'	=>
	<= Returns "Temperature=28C\n\r" from the read-data buffer
	<= SCSI success
Goal: release the Diamond Storage Array for other in-band users	
SCSI cdb: Release LUN 0	=>
	<= SCSI success

4.2 Serial Port Access

The Diamond Storage Array provides remote service operations over the RS-232 serial port using standard terminal emulation software available with most systems.

The Diamond Storage Array supports service operations over the RS-232 serial port using standard terminal emulation software available with most systems.

- 1 Connect a DB-9 null modem serial cable between the port marked RS-232 on the front of the array and one of the computer's serial ports. A gender changer or DB-9 to DB-25 converter may be needed depending on the cables you are using. The cable must be no longer than three meters.
- 2 Boot the computer used to manage the array.
- 3 Start a terminal emulator program such as Windows HyperTerminal. Set the emulator to use the COM port with cable attached, then the following settings:
 - 115200 baud
 - 8 bit ASCII
 - no parity
 - ASCII terminal type
 - 1 stop bit
 - flow control none
 - echo typed characters locally
- 4 Turn on the array. Upon successful power on and diagnostics, the unit should display the POST (power on self test) information found in Exhibit 4.0-1 on page 18.
- 5 You should now be in the Command Line Interface mode. To see a list of available commands, type help at the **Ready** prompt or refer to this manual's Index.
- 6 Use the CLI to configure the unit as a JBOD, RAID Level 0, RAID Level 1, RAID Level 10 or RAID Level 5 array with partitions, zones and/or hot spare sleds as described later in this manual.

4.3 Ethernet Access: Telnet and SNMP Protocols

The optional 10/100 BaseT Ethernet port provides Telnet- or SNMP-based monitoring and management.

The 10/100 Base T Ethernet management services card provides Telnet-based monitoring and management, including firmware update using FTP.

Remote system monitoring is also available using Simple Network Management Protocol (SNMP). An agent resides in the Diamond Storage Array which takes information from the Array and translates it into a form compatible with SNMP. If certain conditions arise, the agent sends asynchronous notifications (traps) to a client.

To connect to the Ethernet port

- 1 Connect a cross-over cable (for a direct connection to a PC) or regular network cable from a network device to the optional RJ45 Ethernet port on the Ethernet management card on the front of the array.
- 2 If using a direct connection, power on and boot up the host computer.
- 3 Attach a DB-9 null modem serial cable (the cable must be no longer than three meters) from the RS-232 port of the array to a host computer and open a terminal emulator program on the host to set the Ethernet parameters.
- 4 Turn on the array.
- 5 *First time use:* Upon successful power up and diagnostics, set the host computer with the appropriate settings.

The host computer must have appropriate network settings to allow it to communicate with the array. Please see your system administrator for more information.

To use Telnet

- 1 Change the IP address from the default by first accessing the serial connection and changing it using the CLI.

You may change the IP address to a network specific value or, if the local network uses DHCP, you may enable automatic IP addressing (**set IPDHCP enabled**) using the CLI.

- 2 Open a Telnet session on the host computer.

23

- Default IP address: 10.0.0.1
- Port type: telnet
- Terminal type: vt100

- 3 If you make any changes to the network setting on the array, use the **SaveConfiguration Restart** commands.

- 4 Username/password: You are asked for a username and password, up to eight characters each, case insensitive. Only one username/password combination is available per array.

You may change the telnet username and/or password after entering a CLI session using the commands

```
set TelnetUsername [username]
set TelnetPassword [password]
```

Or you may change the telnet username and/or password using the **Configuration** page of the ExpressNAV interface.

RestoreConfiguration default sets the telnet username and password to the default values.

- The username default is **telnet**.and password default is **diamond**.
- 5 In the Command Line Interface mode, see a list of available commands by typing **help** at the **Ready** prompt or refer to this manual's Index.
 - 6 Using the ExpressNAV interface, configure the unit as JBOD, RAID Level 0, RAID Level 1, RAID Level 10 or RAID Level 5 with partitions, zones and/or hot spare sleds as described later in the manual.

To use SNMP

- 1 Enter the CLI through the serial port or Ethernet.
- 2 Change the IP address to a network specific value or, if the local network uses DHCP, you may enable automatic IP addressing.
- 3 Set the number of trap client addresses by typing

```
set SNMPTrapAddress [1-6]
[IPAddress] [Level]
```

- 4 Type **set SNMPUpdates enabled**
- 5 Type **set SNMPTraps enabled**

- 6 Type **SaveConfiguration restart** to reboot the array.
- 7 Install SNMP management software on each client you wish to receive traps (messages).
- 8 Call technical support to get the appropriate MIB file for your array.
- 9 For each client, copy the MIB file to the directory containing the SNMP management software.
- 10 From within the SNMP management software, compile the file attodmnd-mib.mib according to the software's procedures.
- 11 Unload any default MIBs.
- 12 Load the Diamond MIB **ATTODIAMOND**.
- 13 When requested, enter the array's IP address as the **Remote SNMP Agent**.
- 14 The SNMP management software contacts the agent in the array. The screen replies with system information.
- 15 Status is monitored and reported through the SNMP management software.

Traps are generated for the following situations:

- Temperature status changes in any of the sensors located on the array mid-plane.
- A drive or a sled is physically removed from the Array or put into the Array.
- The power supply is turned on or off.
- The fan is physically stopped.

4.4 ATTO ExpressNAV: Browser-based Interface

The easiest way to communicate with the Diamond Storage Array is to use ATTO ExpressNAV, a user-friendly GUI interface accessed through a browser to control the most common capabilities of the array.

Access ATTO ExpressNAV from any browser that supports the latest standards for XHTML 1.0 and CSS1. To take full advantage of the ExpressNAV interface you should have Java script enabled through your browser.

Browser Compatibility

All pages are written in pure XHTML 1.0 and CSS1 to be compatible with the latest versions of Internet Explorer, Netscape, Mozilla (including K-Meleon, Camino, Mozilla Firefox, Epiphany and Galeon), and KHTML (including Konqueror and Safari).

Minimum requirement is for Internet Explorer 5.5 and Netscape 6.2.

To optimize ExpressNAV in Internet Explorer

- 1 Go to the browser toolbar and click on **Tools**
- 2 Click on **Internet Options**
- 3 Click on the **Security** tab
- 4 Click on the **Custom Level** button.
- 5 Click on **Microsoft VM, Java permissions**
- 6 Ensure **Disable Java** is not selected.
- 7 Click on the **Miscellaneous** tab.
- 8 Click on **Metarefresh**

To open an ExpressNAV session


- 1 Obtain the IP address of the array.
- 2 Type the IP address of the array into the browser address window.
- 3 The ExpressNAV interface splash screen is displayed. Click on **Enter**.
- 4 Enter the username and password set previously in *Accessing the Array* on page 17.
 - The default username is **Diamond**
 - The default password is **Password**
- 5 The product faceplate display appears. Click the component you want to manage on the left-hand side menu or go to the Advanced screen to use the CLI.

To navigate ExpressNAV

All pages are accessible by clicking on their titles on the left side of the page. You may also go back one page or go to the **Home** page via the titles on the left side of the page.

Clicking on any of the red option names will bring up a help window. After making changes on a page, click the **Submit** button at the bottom. Clicking this button is equivalent to typing in all the CLI commands and then the command `saveconfiguration norestart`.

Exhibit 4.4-1 Introductory splash screen for ATTO ExpressNAV browser-based configuration tool



Information

Configure choices: red type links to another page to change configuration

ATTO
ATTO Technology, Inc.

ATTO Diamond™

Powered by ATTO

HOME
BACK
Status
Ethernet
Serial Port
Fibre Channel
Storage
RAID
CLEAR DATA
LOGICAL UNITS
PARTITIONS
ZONING
REBUILD
Diamond
Advanced
Restart
Help!

RAID

Type	VDID	#Partitions	State	Capacity	Interleave	
RAID5	59	1	ONLINE	40956MB	16	configure
RAID5	60	1	ONLINE	40956MB	16	configure
RAID5	61	1	DEGRADED	40956MB	16	configure
RAID5	62	1	DEGRADED	40956MB	16	configure

New Configuration

Config Type:

Options:

Num Groups: or Span

Interleave: sled drive

Striping:

Parameter name
Red print links to help text

Submit button
same as typing all CLI commands and *saveconfiguration norestart*

Reset button
return to previous setting without making any changes

options unavailable
because of previous choice

Sidebar:
select the item you wish to view

©2003 ATTO Technology, Inc.

4.4.1 ExpressNAV Pages

Each page in the ATTO ExpressNAV interface provides information and/or configuration parameters based on specific topics. Each page can be reached through the menu on the left hand side of each page.

Status

Contains general information.

- Unit Information such as Vendor ID, Product ID, Firmware Revision, Serial Number
- Environmental Information such as Valid Temperature Range, Midplane Sensor Temperatures
- World Wide Identifiers
- Node Names
- Port Names
- Host Interface Card Status
- Fan Status
- Power Supply Status
- Storage Status
- ATA Disk Errors
- Logical Unit Conflicts

Ethernet

Configures the Ethernet port. Configurable parameters are:

- IP Address
- IP Gateway
- IP Subnet Mask
- Ethernet Speed
- Use DHCP

Refer to *Ethernet Access: Telnet and SNMP Protocols* on page 23 and *Ethernet Commands* on page 39 for details.

SNMP

Remote system monitoring is available using Simple Network Management Protocol (SNMP) including updates, traps and trap recipient IP addresses

Refer to *Ethernet Commands* on page 39 for details on each parameter.

Serial Port

Contains the necessary information to configure the serial port. Configurable options are Baud Rate and Echo.

Refer to *Serial Port Access* on page 21 and *Serial Port Configuration Commands* on page 37 for more information.

Fibre Channel

Contains parameters and information to manage the Fibre Channel port: Data rate, Full duplex mode and Connection mode.

Refer to *Connecting a Fibre Channel Array* on page 11 and *Fibre Channel Configuration Commands* on page 35 for more information.

Storage Management

Shows information about the drives and their status.

Information includes

- Sled Number
- Capacity of each sled
- Number of errors
- Type of configuration (JBOD, RAID5 etc.)
- Virtual ID

You may place sleds on or off line by selecting or de-selecting a check box.

Refer to *Determining Drive and Sled Designations* on page 15 and *Drive Configuration Commands* on page 43 for more information.

Once you open this page, other configuration pages are available on the menu on the left hand side of the page.

RAID

Contains the necessary information and parameters to configure RAID groups. Information provided includes

- Type of configuration (RAID5, JBOD)
- Virtual Disk ID
- Number of Partitions
- State of sleds
- Capacity of sleds
- Interleave values

You may change these parameters:

- Configuration Type (JBOD, RAID5)
- Number of Groups (when applicable)
- Interleave
- Striping method (sled or drive)
- Rebuild priority

Refer to *Configuring Drives* on page 49 and *Drive Configuration Commands* on page 43 for more information.

Clear Data

Allows you to

- view the status of any Clear Data commands in progress
- view the rebuild state of each sled
- initialize a RAID Level 5 Clear Data command
- set a rebuild state for each sled
- change the RAID Interleave parameter

- Enable/disable AutoRebuild

Refer to *Drive Configuration Commands* on page 43 for more information.

Logical Units

Displays information on the logical units which have been configured on the array and allows you to change the status from online or offline or degraded.

Partitions

Allows to you to view the current configuration of the array and to change that configuration, including splitting and merging partitions.

You will lose data in pre-existing RAID groups when you create partitions. Either back up the data to another storage area or only create partitions in data-free RAID configurations.

Do not configure the array into zones until after you have configured partitions.

If a hard disk drive in an existing Virtual Drive is replaced, all partitions that are a part of that Virtual Drive are labeled as degraded. When the Virtual Drive is rebuilt, all partitions are rebuilt.

Refer to *Creating Partitions* on page 67 and *Drive Configuration Commands* on page 43 for more information.

Zoning

Allows you to view the current configuration of the array and to change that configuration.

Refer to *Creating Zones* on page 69 and *Drive Configuration Commands* on page 43 for more information.

Rebuild

Displays the current status of rebuilds on the array and allows you to halt, resume or initiate rebuilds.

Refer to *Rebuilding RAID Level Configurations* on page 63 and *Drive Configuration Commands* on page 43 for more information.

Configuration

Displays information to manage the array.

Configurable options are

- User name
- Password (old password, new password, confirm password)
- Minimum operating temperature
- Maximum operating temperature
- Identify Diamond

- Restore defaults

Refer to *Maintenance Services Commands* on page 48 for details.

Advanced

Allows you to input any CLI command available through the array.

To use the Advanced Page CLI commands

- 1 After the page opens and the **Ready** prompt appears, type in the CLI command
- 2 Click the **Submit** button: this is equivalent to typing the CLI command into a telnet or serial port CLI session.
A text field beneath the box lists the most recent commands issued to the array through this page.
- 3 Type **saveconfiguration norestart**
- 4 Click the **Submit** button. Your changes are implemented.
- 5 To keep the changes through the next power cycle, type **FirmwareRestart** or go to the **Restart** page and click **Restart**.

Restart

Allows you to implement a firmware restart of the array. Access is via the **Restart** link on the left side of the page.



Note

Restarting the firmware may take a few minutes.

- 1 Click the **Restart** button.
A box tells you to wait until the counter gets to 0 and the browser refreshes.
- 2 If the browser does not refresh after the counter gets to 0, click the link to refresh it manually.

Help

Displays help information about the command line interface commands and troubleshooting tips.

Provides links to pages with help text for each of options and one link to the **Troubleshooting Tips and FAQs** page on the ATTO website. Contact information for getting in touch with ATTO technical support is on the right.

When you click a red text box on any page, ExpressNAV asks for your login information, then opens a dialog box with help text.

4.5 CLI: Interface via ASCII-based Commands

The Command Line Interface (CLI) provides access to Diamond Storage Array services through a set of ASCII-based commands.

CLI commands may be entered while in CLI mode or by accessing the **Advanced CLI** configuration page in the ExpressNAV interface.

CLI commands are context sensitive and generally follow a standard format:

```
[Get|Set] Command Parameter  
1|Parameter 2]
```

CLI commands are not case sensitive: you may type all upper or all lower case or a mixture, no matter what the definition either in **Help** or these pages states. Upper and lower case in this manual and the help screen are for clarification only.

Commands generally have three types of operation: get, set and immediate as summarized in Exhibit 4.5-1.

The get form returns the value of a parameter or setting and is an informational command.

Responses to get commands are specified in the Results field for each command, followed by **Ready**.

The set form is an action that changes the value of a parameter or configuration setting. It may require a SaveConfiguration command and a

system restart before changes are implemented. The restart can be accomplished as part of the SaveConfiguration command or by using a separate FirmwareRestart command.

A number of set commands may be issued before the SaveConfiguration command.

Responses to set commands are either an error message or **Ready**. *, which indicates a **SaveConfiguration** command is required.

Set commands which do not require a **SaveConfiguration** command, defined as immediate commands, are immediately implemented.

Responses to Immediate commands are either an error message or **Ready**.



Note

Zone commands do not use the get, set forms. Refer to *Creating Zones* on page 69 for more information on how to use Zone commands.

Symbols, typefaces and abbreviations used to indicate functions and elements of the CLI used in this manual include those found in Exhibit 4.5-2.

Exhibit 4.5-1 Command Line Interface actions and responses

Set commands configure the array and display what you have changed after completing the task. Commands which require a SaveConfiguration command to complete their implementation returns Ready . *. Set commands which do not require a SaveConfiguration command are immediately implemented.
Get commands display information about the configuration of the array. Responses to get commands are specified in the Results field for each command, followed by Ready .
Screen messages, also called returns, may be either terse, with just the current information, or verbose, with labels and the current information. Default is verbose. If you want the terse mode, type set VerboseMode disabled .

Exhibit 4.5-2 Symbols, typefaces and abbreviations used to indicate functions and elements of the Command Line Interface

Symbol	Indicates
[]	Required entry
< >	Optional entry
	pick one of
...	Ellipses, repetition of preceding item
\n	end of line
-	a range (6 – 9 = 6, 7, 8, 9)
Boldface words	must be typed as they appear
<i>Italicized words</i>	Arguments which must be replaced by whatever they represent
Fp	Fibre Channel port number (0 <= fp <= 1)
Fl	Fibre Channel LUN (0 <= fl <= 24), where 0 represents the array unit, and 1-24 represent the disk drives.
device_lun	The LUN of the JBOD or RAID drive (used in zoning)
host_name	In a Fibre Channel environment, the WWPN (World Wide Port Name); in a SCSI environment, SCSI Initiator ID (used in zoning)
port_number	The Diamond port number (0, 1) for the data path (used in zoning)
zone_name	Alphanumeric or '_', character string less than or equal to 16 characters long (used in zoning)

Exhibit 4.5-3 CLI commands returns may be terse (short) or verbose (with parameter names and details of results). Zoning command returns follow these patterns:

return type/mode	return format/content
errors	context sensitive error message\n ERROR message\n
command completion single line output multiple line output	Ready.\n shows the line shows the line count followed by the output lines

4.5.1 Summary of CLI Commands

A summary of the Command Line Interface commands, their defaults, an example of how they might be used, and where you can find the specifics of the command. Commands which have no default values associated with them have a blank entry in that column of the table.

Command	Defaults	Example	Page
AtaDiskState	Online	set AtaDiskState 6 1 offline	43
AudibleAlarm	Disabled	set AudibleAlarm disabled	41
AutoRebuild	Disabled	set AutoRebuild enabled	43
ClearDiskReservedArea		ClearDiskReservedArea 8 2	43
DiamondModel	Diamond	get DiamondModel	41
DiamondName	" "	get DiamondName	41
DriveCopy		DriveCopy 1 1 2 2	43
DriveCopyHalt		DriveCopyHalt 2 2	43
DriveCopyResume		DriveCopyResume 2 2	43
DriveCopyStatus		DriveCopyStatus	41
DriveInfo		DriveInfo 3 2	41, 43
DriveSledPower	On	set DriveSledPower 9 1 off	43
DriveWipe		DriveWipe 2 2	44
EthernetSpeed	Auto	set EthernetSpeed 100	39
FcConnMode	Loop	get FcConnMode	35
FcDataRate	Auto	set FcDataRate 2 gigabit	35
FcFairArb	Enabled	get FcFairArb	35
FcFrameLength	2048	get FcFrameLength	35
FcFullDuplex	Enabled	set FcFullDuplex enabled	35
FcHard	Disabled	get FcHard	35
FcHardAddress	0x03, 0x04	get FcHardAddress 0	35
FcNodeName		get FcNodeName	41
FcPortInfo		get FcPortInfo	35
FcPortList		FcPortList	35, 41
FcPortName		get FcPortName 1	41
FcSCSIBusyStatus	Busy	set FcSCSIStatus qfull	36, 48
FcWWName		get FcWWName 0	36
FirmwareRestart		FirmwareRestart	34, 48
FTPPassword	diamond	set FTPPassword barbw52	40
Help		Help DriveInfo	34, 41
IdentifyDiamond	Disabled	get IdentifyDiamond	41
IdeTransferRate	4	set IdeTransferRate 4	44
Info		Info	41
IPAddress	10.0.0.1	get IPAddress	39

Command	Defaults	Example	Page
IPDHCP	Disabled	set IPDHCP enabled	39
IPGateway	0.0.0.0	set IPGateway 200.10.22.3	39
IPSubnetMask	255.255.255.0	set IPSubnetMask 255.255.255.0	39
LUNInfo		LUNInfo 1	41, 44
LUNState	Online	get LunState 1	44
MaxEnclTempAlrm	47	get MaxEnclTempAlrm	48
MinEnclTempAlrm	5	set MinEnclTempAlrm 10	48
PartitionCommit		PartitionCommit	44
PartitionInfo		PartitionInfo planned	42, 44
PartitionMerge		PartitionMerge 1 all 2 3	44
PartitionSplit		PartitionSplit 1 2 2	44
PowerAudibleAlarm	Enabled	set PowerAudibleAlarm disabled	41
QuickRAID0	sled	set QuickRAID0 6	44
QuickRAID1		set QuickRAID1	44
QuickRAID10		set QuickRAID10 2	45
QuickRAID5		set QuickRAID5 4	45
RAID5ClearData		RAID5ClearData	46
RAID5ClearDataStatus		RAID5ClearDataStatus	45
RAIDHaltRebuild		RAIDHaltRebuild 3	45
RAIDInterleave	128	get RAIDInterleave	45, 65
RAIDManualRebuild		RAIDManualRebuild 2 3	45
RAIDRebuildState		set RAIDRebuildState 2 OK	42
RAIDRebuildStatus		get RAIDRebuildStatus	42, 46
RAIDResumeRebuild		RAIDResumeRebuild 5	46
RestoreConfiguration		RestoreConfiguration default	34
RestoreModePages		RestoreModePages	46
SaveConfiguration		SaveConfiguration restart	34
SerialNumber		get SerialNumber	42
SerialPortBaudRate	115200	set SerialPortBaudRate 9600	37
SerialPortEcho	Disabled	get SerialPortEcho	37
SerialPortHandshake	None	get SerialPortHandshake	37
SerialPortStopBits	1	set SerialPortStopBits 1	37
SledFaultLED		set SledFaultLED 9 on	42, 46
SMARTData		SMARTData 2 1	42
SNMPTrapAddress	10.0.0.1	set snmptrapaddress 1 255.555.555.555 All	39
SNMPTraps	Disabled	get snmptraps	39
SNMPUpdates	Disabled	get snmpupdates	40
SystemSN		set systemsn 5564	34

Command	Defaults	Example	Page
TelnetPassword	diamond	set TelnetPassword 123ABC	40
TelnetTimeout	Disabled	set TelnetTimeout 360	40
TelnetUsername	telnet	set TelnetUsername diamond1	40
Temperature		get Temperature	42, 48
VerboseMode	Enabled	get VerboseMode	34
VirtualDriveInfo		virtualdriveinfo active	42, 46
Zmodem		zmodem receive	48
ZoneAddDevice		zoneadddevice zone1 2	46
ZoneAddHost		zoneaddhost zone1 20:00:00:18:86:00:98:00	46
ZoneAddPort		zoneaddport zone1 0	46
ZoneClearAll		zoneclearall	46
ZoneCommit		zonecommit	46
ZoneCreate		zonecreate zone1	47
ZoneDelete		zonedeletem zone1	47
ZoneInfo		zoneinfo	42
ZoneRemoveDevice		zoneremovedevice zone1 1	47
ZoneRemoveHost		zoneremovehost zone1 20:00:00:18:86:00:98:00	46
ZoneRemovePort		zoneremoveport zone1 0	45
ZoneRetrieve		zoneretrieve	48
ZoneState	Disabled	ZoneState zone1 enabled	47

4.5.2 General Use Commands

The following commands, listed alphabetically, describe or perform general functions.

FirmwareRestart

Causes a warm restart of the Diamond Storage Array.

Immediate command: FirmwareRestart

Help

Displays a list of available commands. Type 'help' followed by a command name to display detailed command-specific information.

Get syntax: Help [Command Name]

PartitionCommit

Commits the current Planned Partition Configuration, making it the persistent, Active configuration.

PartitionCommit must be used to alter any partition settings. Performs a firmware restart.

Set syntax: PartitionCommit

Get syntax: none

SaveConfiguration

If the restart option is selected, the Diamond cycles its power. The norestart option saves changes without restarting. Please note: certain modifications require a SaveConfiguration command and a system restart. If required, the return **Ready**. * displays after the return for the modification. You may make several changes through commands and SaveConfiguration before implementing a restart, but once you have restarted the Diamond, all the

command changes created before the restart and saved are implemented. Changes to zones, however, are unaffected by SaveConfiguration. You must use ZoneCommit.

Restart or no Restart parameter is optional

Set syntax: SaveConfiguration <Restart| NoRestart>

SystemSN

Stores the Diamond Storage Array serial number. The serial number may be one to 16 characters.

Set syntax: set SystemSN [n]

Requires a SaveConfiguration command

Get syntax: get SystemSN

VerboseMode

Specifies the detail of verbal feedback for the CLI.

Disabling this option removes parameter names from 'get' commands and removes descriptions from 'help' commands.

Choices: enabled, disabled

Default: enabled (returns have parameter information)

Set syntax: set VerboseMode [enabled | disabled]

Get syntax: get VerboseMode

ZoneCommit

Commits the current Planned Zone Configuration File, making it the persistent, active configuration.

Set syntax: **ZoneCommit**

4.5.3 Fibre Channel Configuration Commands

The Fibre Channel ports are configured with default settings but may be customized to your specifications using the CLI commands in this section.

FcConnMode

Specifies the Fibre Channel topology for the Diamond Storage Array. Options are loop only (loop), point-to-point only (ptp), loop preferred (loop-ptp) or point-to-point preferred (ptp-loop). Refer to *Connecting a Fibre Channel Array* on page 11 for more information on Fibre Channel topology. Applies to both Host Interface Cards

Default: loop

Set syntax: set FcConnMode [loop | ptp| loop-ptp| ptp-loop]

Requires a SaveConfiguration Restart command

Get syntax: get FcConnMode

FcDataRate

Specifies the Fibre Channel data rate at which the Diamond operates. Applies to both Host Interface Cards

Default: auto

Set syntax: set FcDataRate [1gb | 2gb | auto]

Requires a SaveConfiguration Restart command

Get syntax: get FcDataRate

FcFairArb

Enabling this feature causes the Diamond Storage Array to follow the arbitration fairness rules on the FC-AL.

Applies to both Fibre Channel ports

Default: enabled, enabling arbitration fairness

Set syntax: set FcFairArb [enabled | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get FcFairArb

FcFrameLength

Sets the frame length of a command. If not specified in the set command, current frame length is displayed.

Applies to both Fibre Channel ports

Default: 2048

Set syntax: set FcFrameLength [512 | 1024 | 2048]

Requires a SaveConfiguration Restart command

Get syntax: get FcFrameLength

FcFullDuplex

Enable to allow full duplex Fibre Channel communication between the Diamond Storage Array and

host devices. Disable FcFullDuplex causes half duplex mode. Applies to both Fibre Channel ports

Default: enabled

Set syntax: set FcFullDuplex [enabled | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get FcFullDuplex

FcHard

If hard addresses are enabled, the Diamond Storage Array tries to use its internal hard address as its address on the loop. Under soft addressing, the Diamond Storage Array loop address is assigned during loop initialization. Use FcHardAddress (described below) if you enable hard addressing. Applies to both Fibre Channel ports

Default: disabled

Set syntax: set FcHard [enabled | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get FcHard

FcHardAddress

This hexadecimal value represents the address the Diamond Storage Array tries to use if hard addressing is enabled. When an optional address is not present, the current value is displayed. Each port has individual hard address value

Default: 0x03 for port 0, 0x04 for port 1

Set syntax: set FcHardAddress [fp [[address]]]

Requires a SaveConfiguration Restart command

Get syntax: get FcHardAddress [fp]

FcPortInfo

Retrieves information about the current state of each Fibre Channel port. The status field indicates 'disabled' when a port has been internally disabled.

Immediate command: FcPortInfo

FcPortList

Lists the status of all available Fibre Channel ports.

Immediate command: FcPortList

FcSCSIBusyStatus

Specifies the SCSI status value returned when the Diamond is unable to accept a SCSI command because of a temporary lack of resources.

Default: busy

Actions: set FcSCSIBusyStatus [busy | qfull]

Requires a SaveConfiguration Restart command

Get syntax: get FcSCSIBusyStatus

FcWWName

Sets or view the Word Wide Name (WWPN) of the referenced Fibre Channel port. The WWPN is a unique 8-byte number that identifies the port on a Fibre Channel network. Only the three least significant bytes of the WWPN can be modified. Fabric and loop operations are unpredictable if duplicate WWNs are assigned.

Default: 20 00 0x where x is for 0 for port 0, 1 for port 1

Set syntax: set FcWWName [PortNumber [0 | 1]

Requires a SaveConfiguration Restart command

Get syntax: get FcWWName [PortNumber [0 | 1]

4.5.4 Serial Port Configuration Commands

The serial port configuration may be customized by using the following commands:

SerialPortBaudRate

Sets the baud rate the Diamond Storage Array uses for its terminal interface.

Choices: 2400, 9600, 19200, 38400, 57600, 115200

Default: 115200

Set syntax: set SerialPortBaudRate [2400 | 9600 | 19200 | 38400 | 57600 | 115200]

Get syntax: get SerialPortBaudRate

SerialPortEcho

Controls whether the Diamond Storage Array echoes characters on its RS-232 port. Local ASCII terminal (or terminal emulator) echo settings should be set to disabled while in serialportecho enabled

Default: disabled

Set syntax: set SerialPortEcho [enabled | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get SerialPortEcho

SerialPortHandshake

Describes which handshaking method the Diamond Storage Array uses for its terminal interface (hardware, Xon/Xoff or none).

Choices: hardware, Xon or none

Default: none

Set syntax: set SerialPortHandshake [hard| Xon| none]

Requires a SaveConfiguration Restart command

Get syntax: get SerialPortHandshake

SerialPortStopBits

Configures/reports the number of stop bits per character for the Diamond Storage Array RS -232 serial port. The number of data bits per character is fixed at 8 with no parity.

Choices: 1 or 2

Default: 1 stop bit

Set syntax: set SerialPortStopBits [1 | 2]

Requires a SaveConfiguration Restart command

Get syntax: get SerialPortStopBits

4.5.5 Ethernet Commands

Ethernet configuration commands configure the Ethernet and TCP/IP parameters for the Diamond Storage Array with an optional Ethernet management services card.

EthernetSpeed

Specifies the speed of the Ethernet Network to which the Diamond Storage Array is connected.

Choices: 10 (10 baseT), 100 (100 baseT), auto

Default: auto

Set syntax: set EthernetSpeed [10 | 100| Auto]

Requires a SaveConfiguration Restart command

Get syntax: get EthernetSpeed

FTPPassword

Specifies a password of up to 32 characters for an FTP session.

Default: diamond

Set syntax: set FTPPassword

Requires a SaveConfiguration Restart command

IPAddress

Specifies the IP Address of the Diamond Storage Array on the Ethernet network. If DHCP is enabled, the assigned address of the Diamond is displayed. Setting this value always modifies the internal NVRAM value of the IP address. If IPDHCP is enabled (see below), get command reports current IP address assigned by DHCP server

Default IP Address: 10.0.0.1

Set syntax: set IPAddress xxx.xxx.xxx.xxx

Requires a SaveConfiguration Restart command

Get syntax: get IPAddress

IPDHCP

Selecting DHCP allow the Diamond Storage Array to request an IP address from the network. It requires that the Diamond be attached to a network with at least one

DHCP server. The network must have at least one DHCP server

Default: disabled

Set syntax: set IPDHCP [enabled | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get IPDHCP

IPGateway

Specifies the IP Gateway for the Diamond Storage Array on the Ethernet network. If IPDHCP is enabled (see above), get command reports current IP gateway assigned by DHCP server Must conform to AAA.BBB.CCC.DDD standard network IP addressing.

Default: 0.0.0.0

Set syntax: set IPGateway AAA.BBB.CCC.DDD

Requires a SaveConfiguration Restart command

Get syntax: get IPGateway

IPSubnetMask

Specifies the IP Subnet Mask for the Diamond Storage Array on the Ethernet network. If DHCP is enabled, the assigned subnet mask for the unit is displayed. Setting this value always modifies the internal NVRAM value of the IP Subnet Mask.

Default: 255.255.0.0

Set syntax: set IPSubnetMask AAA.BBB.CCC.DDD

Requires a SaveConfiguration Restart command

Get syntax: get IPSubnetMask

SNMPTrapAddress

Sets up IP trap address and trap level.

Default: 10.0.0.1

Set syntax: set SNMPTrapAddress [Index: 1-6] [Address: XXX.XXX.XXX.XXX] Level: Critical | Warning | All | None]...

Requires a SaveConfiguration Restart command

Get syntax: get SNMPTrapAddress

SNMPUpdates

Enables or disables the SNMP Management Information Base (MIB) database.

Default: disabled

Set syntax: set SNMPTraps [enabled|disabled]

Requires a SaveConfiguration Restart command

Get syntax: get SNMPUpdates

TelnetPassword

Specifies password for telnet session. Only one username/password combination is available per Diamond Storage Array. RestoreConfiguration default sets the telnet username and password to the default values. The password is case insensitive, with 1 to 8 characters.

Default: diamond

Set syntax: set TelnetPassword [password]

Requires a SaveConfiguration Restart command

TelnetTimeout

Specifies the number of minutes of inactivity which elapses before a telnet session automatically times out.

Choices: 1-1440 minutes

Default: disabled

Set syntax: set TelnetTimeout [1-1440 | disabled]

Requires a SaveConfiguration Restart command

Get syntax: get TelnetTimeout

TelnetUsername

Specifies username for telnet session. Only one username/password combination is available per Diamond Storage Array. RestoreConfiguration default sets the telnet username and password to the default values. Username is case insensitive, 1 to 8 characters

Default: telnet

Set syntax: set TelnetUsername [username]

Requires a SaveConfiguration Restart command

The Telnet and SNMP protocols also use CLI commands.

4.5.6 Diagnostic Commands

Diagnostic commands provide information or diagnostic tools for Fibre Channel, SCSI and Serial port configurations, Diamond Storage Array settings and the status of various commands which affect the ATA drives.

AudibleAlarm

Enables or disables the audible alarm in the Diamond Storage Array. When enabled, an alarm sounds when the Fault LED on the front panel blinks.

Choices: enabled, disabled

Default: disabled

Set syntax: set AudibleAlarm [enabled | disabled]

Get syntax: get AudibleAlarm

DiamondModel

Returns specific Diamond Storage Array model information including firmware release and date.

Get syntax: get DiamondModel

DiamondName

Used to identify this Diamond over its Fibre Channel and Ethernet networks. You may customize the name of each Diamond Storage Array enclosure to distinguish it from other units. Maximum eight characters.

Set syntax: set DiamondName [name]

Requires a SaveConfiguration command

Get syntax: get DiamondName

DriveCopyStatus

Displays the status of a DriveCopy, DriveWipe or RAID5ClearData operation

Immediate command: DriveCopyStatus

DriveInfo

Displays information about all disk drives or detailed information about a specific disk drive. Detailed information about an individual drive is obtained by supplying a drive identifier. VD ID is Virtual Disk ID.

Get syntax: DriveInfo [sled ID] [drive ID]

FcNodeName

Returns the Fibre Channel node name stored in NVRAM for this Fibre Channel port: the same as the World Wide Name for port 0.

Get syntax: get FcNodeName

FcPortList

Lists the status of all available Fibre Channel ports.

Immediate command: FcPortList

FcPortName

Returns the Fibre Channel port name stored in NVRAM for this Fibre Channel port.

Get syntax: get FcPortName [port number]

Help

Displays a list of available commands. Type 'help' followed by a command name to display detailed command-specific information.

Get syntax: Help [Command Name]

IdentifyDiamond

Enable this option to identify the current Diamond Storage Array. The fault LED on its front panel blinks. Disable to cancel the 'blink code.'

Set syntax: set IdentifyDiamond [enabled|disabled]

Get syntax: get IdentifyDiamond

Info

Displays version numbers and other key information about the Diamond Storage Array including data rate, connection mode, WorldWideName, Diamond Storage Array name.

Immediate command: Info [all]

LUNInfo

Displays information about all LUNs (logical unit numbers) or detailed information about a specific LUN. Do not specify a LUN to get information about all LUNs. Specify a LUN to get detailed information about that individual LUN.

Get syntax: get LUNInfo [LUN]

PartitionInfo

Displays Partition information for the selected partitions or all partitions.

Get syntax: PartitionInfo [active|planned] [Virtual Drive ID] [Partition ID]

RAID5ClearDataStatus

Displays the status of RAID5 Clear Data processing. S represents the sled number, D represents the drive number.

Choices: applies to RAID5 only

Set syntax: RAID5ClearDataStatus

RAIDRebuildStatus

Displays the RAID1, RAID 5 or RAID 10 Rebuild Status Summary. If no RAID groups are defined, the header information is displayed with no data. The status summary contains the RAID1, RAID 5 or RAID 10. The RAIDRebuildStatus command has no effect on the state of any rebuild in progress.

Member index and the following fields in tabular form:

- Status: OK, DEGRADED, IN PROGRESS, or FAULTED
- Sled Number: Location of the sled
- Current LBA: Logical block currently being rebuilt
- Maximum LBA: Last logical block to be rebuilt
- Status OK: the RAID1 Mirror is in sync and no rebuild activity is occurring. The Current LBA, Maximum LBA and percentage complete values are not displayed.
- Status Degraded: the RAID1 Mirror is out of sync and is waiting to be rebuilt.
- Status In Progress: the RAID1 Mirror is out of sync and a rebuild is occurring on the respective drive.
- Status Faulted: the RAID1 Mirror is out of sync and an error occurred in an attempt to synchronize the drives in the RAID 1 group.

SerialNumber

View the serial board number, a 10-character field. The first four alphanumeric characters are an abbreviation representing the product name. The remaining six digits are the individual unit's number.

Get syntax: get SerialNumber

SledFaultLED

Changes the state of the selected sled LED to the indicated state.

Choices: sled number 1-12 or all, turn on or off

Default: off

Set syntax: set SledFaultLED [all | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12] [on | off]

SMARTData

Displays the current SMART Data (Self-Monitoring Analysis and Reporting Technology) for the specified drive.

Immediate command: SMARTData [Sled#] [Drive#]

Temperature

Returns the current internal temperature of this Diamond Storage Array in degrees Celsius. The value is read only.

Get syntax: get Temperature

VirtualDriveInfo

Displays the named Virtual Drive definitions.

Get syntax: VirtualDriveInfo [active|planned] [Virtual Drive ID]

ZoneInfo

Displays the named zones' definitions. Information about the Active Zone Configuration is the default; if you want information about the Planned Zone Configuration, type ZoneInfo Planned zone_name.

Zone syntax: ZoneInfo [Planned] [zone_name] [all]

4.5.7 Drive Configuration Commands

The Diamond Storage Array ATA drives may be monitored or configured through the CLI using the commands listed below.

ATADiskState

Sets the ATA disk to the specified state.



CAUTION

In a Hot Spare sled configuration, a drive sled should only be taken offline if there is absolutely no activity on that drive. If there is any activity, the rebuild of the Hot Spare sled may be flawed.

Choices: enter sled number (1-12), drive number (1-2) and online or offline

Default: online

Set syntax: set AtaDiskState [sled number] [drive number] [online|offline]

Get syntax: get AtaDiskState [sled number] [drive number]

AutoRebuild

If enabled, initiates an automatic rebuild of a “degraded” RAID group when a sled is replaced by a new sled. If disabled, you must manually rebuild the RAID configuration for the new sled by using RAIDManualRebuild.

Default: disabled

Set syntax: set AutoRebuild [enabled|disabled]

Requires a SaveConfiguration Restart command

Get syntax: get AutoRebuild

ClearDiskReservedAreaData

Clears the data in the disk’s reserved area. Restarting the Diamond Storage Array is required for these settings to take effect. Omitting ReservedAreaIndex clears the entire reserved area.

Choices: enter sled number (1-12), drive number (1-2)

Immediate command: ClearDiskReservedAreaData [sled number] [drive number]

DriveCopy

Copies a drive from the source disk to the destination disk. Parameters are the sled and drive numbers of the source and destination drives. The destination drive must be offline: use the ATADiskState command to determine if the disks are offline. If you choose the same source and destination drive, this command performs a DriveWipe.

Set syntax: DriveCopy [Source Sled] [Source Drive] [Destination Sled] [Destination Drive]

Get syntax: DriveCopyStatus

DriveCopyHalt

Stops a DriveCopy operation in progress.

Set syntax: DriveCopyHalt [Destination Sled] [Destination Drive]

Get syntax: DriveCopyStatus

DriveCopyResume

Resumes a DriveCopy operation that had been stopped.

Set syntax: DriveCopyResume [Destination Sled] [Destination Drive]

Get syntax: DriveCopyStatus

DriveCopyStatus

Displays the status of a DriveCopy, DriveWipe or RAID5ClearData operation

Immediate command: DriveCopyStatus

DriveInfo

Displays information about all disk drives or detailed information about a specific disk drive. Detailed information about an individual drive is obtained by supplying a drive identifier. For examples Refer to *Diagnostic Commands* on page 41.

Get syntax: DriveInfo [sled ID] [drive ID]

DriveSledPower

Gets/sets power to the specified drive sled. Sled must be offline

Default: on

Set syntax: set DriveSledPower [sled number] [on | off]

Get syntax: get DriveSledPower [sled number]

DriveWipe

Initializes a drive: wipes it of all data. Drive must be offline

Set syntax: DriveWipe [Destination Sled] [Dest Drive]
Requires a SaveConfiguration Restart command
Get syntax: DriveCopyStatus

IdeTransferRate

Sets the DMA mode transfer rate for all devices.

Choices: 0, 1, 2, 3, 4
Default: 4

Set syntax: set IdeTransferRate [0 | 1 | 2 | 3 | 4]
Requires a SaveConfiguration Restart command
Get syntax: get IdeTransferRate

LUNInfo

Displays information about all LUNs (logical unit numbers) or detailed information about a specific LUN. Do not specify a LUN to get information about all LUNs. Specify a LUN to get detailed information about that individual LUN. For examples, Refer to *Diagnostic Commands* on page 41.

Get syntax: get LUNInfo [LUN]

LUNState

Sets the LUN to the specified state. May be used to facilitate removal and insertion of sleds and RAID groups during power up/power down of sleds.

Default: online
Set syntax: set LUNState [LUN number] [online|offline]
Get syntax: get LUNState [LUN number]

PartitionCommit

Commits the current Planned Partition Configuration, making it the persistent, Active configuration. PartitionCommit must be used to alter any partition settings. Performs a firmware restart.

Immediate command: PartitionCommit
Information: PartitionInfo

PartitionInfo

Displays Partition information for the selected partitions. Refer to *Diagnostic Commands* on page 41 for examples.

Immediate command: PartitionInfo [active|planned]
[Virtual Drive ID] [Partition ID]

PartitionMerge

Merges the specified partitions into a single larger partition. All partitions to be merged must be stored on contiguous sections of the specified Virtual Drive. If you want this configuration to become the active configuration, follow with a PartitionCommit command.

Set syntax: PartitionMerge [Virtual Drive ID] [Partition ID|all] [<Partition number...>]
Requires a PartitionCommit command
Information: PartitionInfo

PartitionSplit

Create or modify partitions on a Virtual Drive. If you want this configuration to become the active configuration, follow with a PartitionCommit command.

Set syntax: PartitionSplit [Virtual Drive ID] [Partition ID] [Number of partitions]
Requires a PartitionCommit command
Information: PartitionInfo

QuickRAID0

Specifies the RAID Level 0 configuration for the system. The default 0, or no RAID groups, configures the Diamond Storage Array in a JBOD configuration. DRIVE indicates drives on one side of the array are adjacent members of the same stripe group while SLED indicates the two drives on the same sled are adjacent members of a stripe group.

Choices: 0, 1, 2, 3, 4, 6, 12
Set syntax: set QuickRAID0 [0|1|2|3|4|6|12][drive| sled]
Requires a SaveConfiguration Restart command
Information: DriveInfo

QuickRAID1

Specifies the RAID Level 1 configuration for the system. Sets the system to a mirrored array of spanned drives. Causes the Configuration Manager to “stamp” the new configuration onto the drives to take effect at the next system startup. Setting QuickRAID0 0 removes all RAID

QuickRAID5

Specifies the RAID Level 5 configuration for the system. Sets the system to spanned drives with parity information. Causes the Configuration Manager to “stamp” the new configuration onto the drives to take effect at the next system startup.

Setting QuickRAID5 0 removes all RAID configurations and creates JBOD.

Setting QuickRAID5 ALL creates one group that includes all contiguous sleds (minus Hot Spare sleds if applicable).

To complete RAID Level 5 setup, the RAID5ClearData command must be issued after the Diamond Storage Array has been restarted: DO NOT remove power from the array during this operation.

Choices: 0, 1, 2, 3, 4, all

Set syntax: set QuickRAID5 [1 | 2 | 3 | 4 | all] <Number Hot Spare sleds>

Requires a SaveConfiguration Restart command

Get syntax: DriveInfo

QuickRAID10

Specifies the RAID Level 10 configuration for the system; sets the system to a mirrored array of spanned drives and causes the Configuration Manager to “stamp” the new configuration onto the drives to take effect at the next system startup. Setting QuickRAID10 0 removes all RAID configurations and creates a JBOD.

Choices: 0, 1, 2, 3

Set syntax: set QuickRAID10 [1 | 2 | 3] <Number Hot Spare sleds>

Requires a SaveConfiguration Restart command

Get syntax: DriveInfo

RAID5ClearData

To Zero all drives and parity to make newly created RAID5 groups ‘coherent’. The parameter ALL clears data on all RAID5 groups present in the system. Must be used at initial configuration to ensure parity is valid for all drives in the RAID Level 5 group by setting all data and parity to zero. Takes all LUNs offline automatically, then brings them online. The operation takes 3-6 hours, depending on drive capacity. Do not interrupt this

process. This is a destructive operation: all information on these drives is lost. Applies to RAID5 only

Set syntax: RAID5ClearData [ALL | LUN]

Get syntax: DriveCopyStatus or RAID5ClearDataStatus

RAID5ClearDataStatus

Displays the status of RAID5 Clear Data processing. S represents the sled number, D represents the drive number. Applies to RAID5 only

Set syntax: RAID5ClearDataStatus

RAIDInterleave

Specifies the Interleave size (in 512 byte blocks) between members of a RAID group. SPAN indicates that the interleave size between drives in the group is the minimum drive size of all members in the group.

RAIDInterleave options are 16, 32, 64, 128, 256 blocks and span. Span is not available for RAID Level 5.

Choices all but RAID Level 5: 16, 32, 64, 128, 256 or SPAN

Choices for RAID Level 5: 16, 32, 64, 128, 256

Default: 128

Set syntax: set RAIDInterleave [1-256] [SPAN]

Requires a SaveConfiguration Restart command

Get syntax: get RAIDInterleave

RAIDHaltRebuild

Stops a RAID Level 1, 5 or 10 rebuild that is in progress.

Immediate command: RAIDHaltRebuild [Sled Number]

RAIDManualRebuild

Initiates a manual rebuild of a RAID Level 1, 5 or 10 LUN. An error message is returned if the specified LUN is not a RAID Level 1, Level 5 or 10 LUN or if the sled number is not available; no rebuild takes place. Applies to RAID Levels 1, 5 and 10 only

Set syntax: RAIDManualRebuild [LUN] [Sled Number]

RAIDRebuildState

Sets the RAID Level 1, Level 5 or Level 10 rebuild status of the specified sled to OK, degraded or faulted.

Set syntax: set RAIDRebuildState [Sled Number] [Degraded | OK | Faulted]

RAIDRebuildStatus

Displays the RAID1, RAID 5 or RAID 10 Rebuild Status Summary. If no RAID groups are defined, the header information is displayed with no data. The status summary contains the RAID1, RAID 5 or RAID 10. The RAIDRebuildStatus command has no effect on the state of any rebuild in progress.

Member index and the following fields in tabular form: Status, Sled Number, Current LBA, Maximum LBA, Status.

Get syntax: RAIDRebuildStatus

RAIDResumeRebuild

Resumes a RAID Level 1, Level 5 or Level 10 rebuild which had been previously stopped.

Immediate command: RAIDResumeRebuild [Sled Number]

RebuildPriority

Sets the priority of a RAID Level 1, Level 5 or Level 10 rebuild. If you select High priority, rebuild I/O requests are implemented before system I/O requests. If you select Low priority, rebuild I/O requests executes only when there are no pending I/O requests. If you select Same priority, rebuild I/O and system I/O receive equal consideration.

Set syntax: set RebuildPriority [high|low|same]
Requires a SaveConfiguration Restart command
Get syntax: get RebuildPriority

ResolveLUNConflicts

Re-numbers any conflicting Logical Unit numbers that exist in the Diamond. Conflicts may occur when a unit is taken from one Array and inserted into another Array.

Set syntax: ResolveLUNConflicts

RestoreModePages

Restores all mode pages to the factory set default. Restarting the Diamond Storage Array is required for settings to take effect.

Set syntax: RestoreModePages

SledFaultLED

Changes the state of the selected sled LED to the indicated state.

Choices: enter sled number 1-12 or all, on or off
Default: off
Set syntax: set SledFaultLED [all|sled number] [on| off]

VirtualDriveInfo

Displays the named Virtual Drive definitions. For examples, refer to *Diagnostic Commands* on page 41.

Immediate command: VirtualDriveInfo [active|planned] [Virtual Drive ID]

ZoneAddDevice

Adds one or more LUNs (devices) to an existing zone. LUNs not added to zones are not available. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Zone syntax: ZoneAddDevice [zone_name] [device_LUN...]

ZoneAddHost

Adds one or more hosts to an existing zone. The host is the WWPN of the HBA attached to the system. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

ZoneAddPort

Adds one or more ports [port_name] to an existing zone. The port refers to the specific Host Interface Card on the Diamond Storage Array, either 0 or 1. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneAddPort [zone_name] [0|1]

ZoneClearAll

Removes all entries from the Planned Zone Configuration. Removes any active zones if followed immediately by a ZoneCommit command.

Set syntax: ZoneClearAll

ZoneCommit

Commits the current Planned Zone Configuration, making it the persistent, Active configuration. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command.

Set syntax: ZoneCommit

ZoneCreate

Creates a new named zone. Names may be up to 16 characters. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneCreate zone_name

ZoneDelete

Deletes one or more named zones. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneDelete [zone_name...]

ZoneInfo

Displays the named zones' definitions. Information about the Active Zone Configuration is the default; if you want information about the Planned Zone Configuration, type ZoneInfo Planned zone_name.

Get syntax: ZoneInfo

ZoneRemoveDevice

Removes one or more LUNs (devices) from an existing zone. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneRemoveDevice [zone_name] [device_LUN...]

ZoneRemoveHost

Removes one or more hosts from an existing zone. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneAddHost [zone_name] [host_name...]

ZoneRemovePort

Removes one or more ports [port_name] from an existing zone. The port refers to the specific HIC, either 0 or 1. To complete this procedure, the ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Set syntax: ZoneAddPort [zone_name] [0|1]

ZoneRetrieve

Retrieves the Active Zone Configuration into the Planned Zone Configuration to allow modifications of the current configuration.

Get syntax: ZoneRetrieve

ZoneState

Changes the specific state of a zone. The ZoneState command must be entered as enabled to activate the zone before using the ZoneCommit command (which makes the Planned Zone Configuration the Active configuration).

Default: disabled

Set syntax: ZoneState zone_name [enabled|disabled]

Get syntax: ZoneInfo

4.5.8 Maintenance Services Commands

Maintenance commands allow updating and maintenance of the Diamond Storage Array.

FcScsiBusyStatus

You may set the Diamond to report busy or queue full when it is unable to accept a command.

Default: Busy

Set syntax: set FcScsiBusyStatus [busy|qfull]

FirmwareRestart

Causes a warm restart of the Diamond Storage Array

Immediate command: FirmwareRestart

MaxEnclTempAlrm

Sets/displays the maximum enclosure temperature alarm of the Diamond Storage Array in degrees Celsius. Valid entries are between 5 and 52 degrees and above the current minimum enclosure temperature alarm

Default: 47

Set syntax: set MaxEnclTempAlrm [5-52]

Requires a SaveConfiguration command

Get syntax: get MaxEnclTempAlrm

MinEnclTempAlrm

Sets/displays the minimum enclosure temperature alarm of the Diamond Storage Array in degrees Celsius. Valid entries are between 5 and 47 degrees and below the current maximum enclosure temperature alarm

Default: 5

Set syntax: set MinEnclTempAlrm [5-47]

Requires a SaveConfiguration command

Get syntax: get MinEnclTempAlrm

Temperature

Returns the current internal temperature of this Diamond Storage Array in degrees Celsius. The value is read only.

Get syntax: get Temperature

Zmodem

Use the Zmodem protocol to transfer a file to or from the Diamond Storage Array. The filename is required if the 'send' option is specified. Available only through the RS-232 interface



WARNING

After a firmware image is downloaded to the Diamond Storage Array, the image is placed into flash memory. During this time (about 90 seconds), DO NOT remove power to the Diamond Storage Array or the flash may become corrupted. Power should not be removed until the READY prompt appears.

Immediate command: Zmodem [Send filename|Receive]

ZoneRetrieve

Retrieves the Active Zone Configuration into the Planned Zone Configuration to allow modifications of the current configuration.

Zone syntax: ZoneRetrieve

5.0 Configuring Drives

The Diamond Storage Array can be configured as JBOD, RAID Level 0, RAID Level 1, RAID Level 10 or RAID Level 5 with zones, partitions and/or hot spare sleds. The default is JBOD with a single zone which includes all LUNs (devices), all ports and all hosts. RAID is a storage configuration which uses multiple disk drives to increase capacity, performance and/or reliability.

You may configure your Diamond Storage Array in several different ways depending on your needs although the Diamond makes some choices for you. The following elements must be considered when you are configuring your Diamond.

- RAID level
- Interleave
- Hot Spares option
- Number of partitions
- Number of zones

Using the ExpressNAV browser-based interface is the easiest way to set up your Diamond. You may also use the Command Line Interface commands.



CAUTION

Changing these parameters causes all previous drive data on the Diamond Storage Array to be erased. Make sure you back up all information before setting up a different configuration.

JBOD (Just a Bunch of Disks)

JBOD (Just a Bunch of Disks) configuration, the default for the Diamond Storage Array, allows many individual disk drives to be available for normal storage operations. A JBOD configuration allows you to access each disk drive in the array independently. Any action you can do to a normal disk drive can be performed on any disk in the JBOD.

RAID Level 0

RAID Level 0 (striping) is based on the fact that increased disk performance can be achieved by simultaneously accessing data across multiple disk drives in an array. This arrangement increases data transfer rates while reducing average access time by overlapping drive seeks. RAID Level 0 groups provide data that is striped across several drives. RAID Level 0 is pure striping, without redundancy, meaning there is no data protection. If one disk fails, all data within that stripe set is lost.

RAID Level 0 is used by applications requiring high performance for non-critical data.

The **QuickRAID0** command, accessed through the Command Line Interface, allows a simple, fast, out-of-the-box setup of the array into evenly-sized RAID Level 0 stripe groups.

RAID Level 1

RAID Level 1 (mirroring) ensures the security of data by writing the exact same data simultaneously to two or more different drives. This application is for users with critical data which cannot be lost or corrupted due to the failure of a single drive.

With RAID Level 1, the host sees what it believes to be a single physical disk of a specific size: it does not know or care about the mirrored pair. The RAID controller manages where data is written and read, allowing one disk to fail without the host knowing it has failed. The array sends notification of the failure over the serial or Ethernet port and the fault LED is illuminated. Service personnel can then replace the failed drive and initiate a rebuild.

RAID Level 1 is used in applications containing mission critical data. The QuickRAID1 command, accessed through the CLI, allows a simple, fast, out-of-the-box setup of the array into RAID Level 1 mirrored groups.

RAID Level 10

RAID Level 10 (mirroring with striping) increases data transfer rates while ensuring security by writing the exact same data simultaneously to two or more different drives. RAID Level 10 is used in applications requiring high performance and redundancy, combining the attributes of RAID Levels 1 and 0. The QuickRAID10 command, accessed through the CLI, allows a simple, out-of-the-box setup of RAID Level 10 groups.

RAID Level 5

RAID Level 5 increases reliability while using fewer disks than mirroring by employing parity redundancy. Distributed parity on multiple drives provides the redundancy to rebuild a failed drive from the remaining good drives. Parity data is added to the

transmitted data at one end of the transaction, then the parity data is checked at the other end to make sure the transmission has not had any errors.

In the array, transmitted data with the added parity data is striped across disk drives. A hardware XOR engine computes parity, thus alleviating software processing during reads and writes.

The array operates in degraded mode if a drive fails.

Interleave

The interleave size sets the amount of data to be written to each drive in a RAID group. This is a tunable parameter which takes a single stream of data and breaks it up to use multiple disks per I/O interval.

The CLI command **RAIDInterleave** allows you to change the size of the sector interleave between RAID groups. The value will depend upon the normal expected file transfer size. If the normal file transfer size is large, the interleave value should be large, and vice versa.

The value entered for the **RAIDInterleave** command refers to blocks of data: one block is equivalent to 512 bytes of data.

Valid entries are 16, 32, 64, 128, 256 and SPAN.

SPAN, not available in RAID Level 5, indicates that interleave size between the drives in the group will be the minimum drive size of all members in the group.

Partitions

With the introduction of larger and larger GB-sized drives, the array may have up to 7.2 TB total data capacity. Partitioning can increase storage efficiency by providing more LUNs without using lower capacity RAID groups.

Partitioning allows the creation of multiple logical volumes. Long LBA (64 bit addressing) allows you to take full advantage of the increasing storage capacity made possible through the new high capacity disk drives. Applications and host operating systems which do not support Long LBAs are able to access larger array capacities which otherwise would not have been possible.

Using the CLI or by accessing the **Advanced CLI** configuration page in the ExpressNAV interface, you are able to divide an individual Partition into a set of equally-sized subpartitions which can then be presented to hosts as separate LUNs.

Zones

Zoning is a collection of related Diamond capabilities supporting flexible array configuration management configurable via CLI commands in the CLI or the **Advanced CLI** configuration page of the ExpressNAV interface. Zoning supports security by granting or denying access between initiators and devices as defined by an administrator.

A zone is a collection of devices which can access each other. The devices in a zone usually include one or more initiators, one or more devices, and one or more paths between the initiators and the devices.

Hot Spare sleds

In most configurations, if a member of a virtual device becomes degraded, you must swap out the faulted sled as defined in *Hot Swap Operating Instructions on page 87*. If you have not enabled *AutoRebuild*, you must also start a manual rebuild.

For four configurations, however, Hot Spare sleds may be designated as replacements for faulted sleds without intervention by you or a host.

Each configuration requires a certain number of Hot Spare sleds. These sleds, once designated as Hot Spares, are not available for other use.

The following configurations support optional Hot Spare sleds:

RAID Level 1: 2 Hot Spare sleds

RAID Level 10: 1 group, 2 Hot Spare sleds

RAID Level 5: 1 group, 1 Hot Spare sled

RAID Level 5: 2 groups, 2 Hot Spare sleds

Enhancing performance

SpeedWrite, enabled by the CLI command **SpeedWrite**, improves the performance of WRITE commands

5.1 JBOD

The Diamond Storage Array is set up in a JBOD (Just a Bunch of Disks) configuration as default. JBOD configuration allows for many individual disk drives to be available for normal storage operations.



CAUTION

Selecting JBOD configuration causes all previous drive data on the Diamond Storage Array to be erased. Make sure all of your information is backed up before setting up the array in a JBOD configuration.

A JBOD configuration allows you to access each of the possible 24 disk drives in the Diamond Storage Array independently. In this configuration, any action you can do to a normal disk drive can be performed on any disk in the JBOD.

To set up the JBOD configuration

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in the Command Line Interface mode.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).

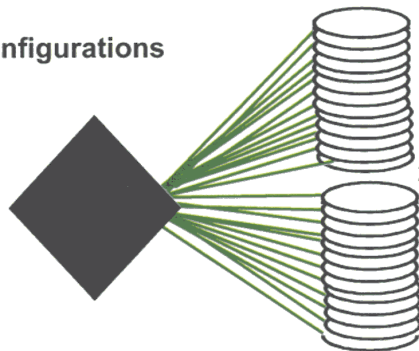


Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

- 3 Type **set QuickRAID0 0**. The command configures the array in the JBOD configuration.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer that is connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

JBOD configurations



24 separate physical drives, each addressed individually

5.2 RAID Level 0

The Diamond Storage Array can be set up into RAID Level 0 (striping) groups to allow it to read and store data at a higher transfer rate. QuickRAID0, a CLI command, allows you to set up the system as if it were a single drive instead of separate drives.



CAUTION

Selecting RAID configuration causes all previous drive data on the Diamond Storage Array to be erased. Make sure all of your information is backed up before setting up RAID groups. You may copy drives first. Refer to Drive Configuration Commands on page 43.

RAID Level 0 groups provide data that is striped across several drives. The QuickRAID0 command, accessed through the Command Line Interface, sets up the Diamond Storage Array into evenly-sized RAID Level 0 stripe groups. Each stripe group is a Virtual Drive named with its own LUN (logical unit number). With a fully populated array, RAID 0 may be configured as 1, 2, 3, 4, 6, or 12 LUNs. As RAID0 1, all 24 drives are configured as a single stripe group. You may also configure two LUNs of 12 drives each, three LUNs of eight drives each, four LUNs of six drives each, six LUNs of four drives each and 12 LUNs of two drives each. (See Exhibit 5.2-4). The command assumes there are 24 drives available to configure the number of LUNs.

Sled-based versus disk-based

RAID Level 0 can be configured one of two ways, sled-based or disk-based. The default is sled-based.

Sled-based Use sled-based if an external RAID controller controls the array to ensure that both drives on a sled are members of the same RAID group (LUN). Removing one sled does not affect other LUNs.

Drive-based RAID0 designates each drive on a sled as a either partner 1 or partner 2. Stripe groups are made by combining all partners designated as 1 together, and all partners designated as 2 together. Removing one sled affects more than one LUN because each partner belongs to a different LUN.

To set up RAID Level 0 groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing*

the Array on page 17). You should now be in CLI.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

- 3 QuickRAID0 options are 0, 1, 2, 3, 4, 6 and 12 RAID groups, sled-based or drive-based. The number indicates the number of RAID groups the array is divided into. *Sled* or *drive* indicates the way you want the array striped. The QuickRAID0 command divides the total number of drives in the same array equally by the number called out in the command. Type

```
set QuickRAID0 [0|1|2|3|4|6|12]
[DRIVE|SLED]
```

- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

If sled(s) or drive(s) are physically missing from the array, the entire stripe group (LUN) containing the drive(s) is unavailable. To determine which drives would be unavailable in various configurations, see the exhibits below.

To remove RAID Level 0 groups from the array

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.
- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).

3 Type set **QuickRAID0 0**. This configures the array in JBOD mode.

4 Continue with steps 4 through 8 from the previous section.

Exhibit 5.2-1 A graphical representation of RAID Level 0 configuration.

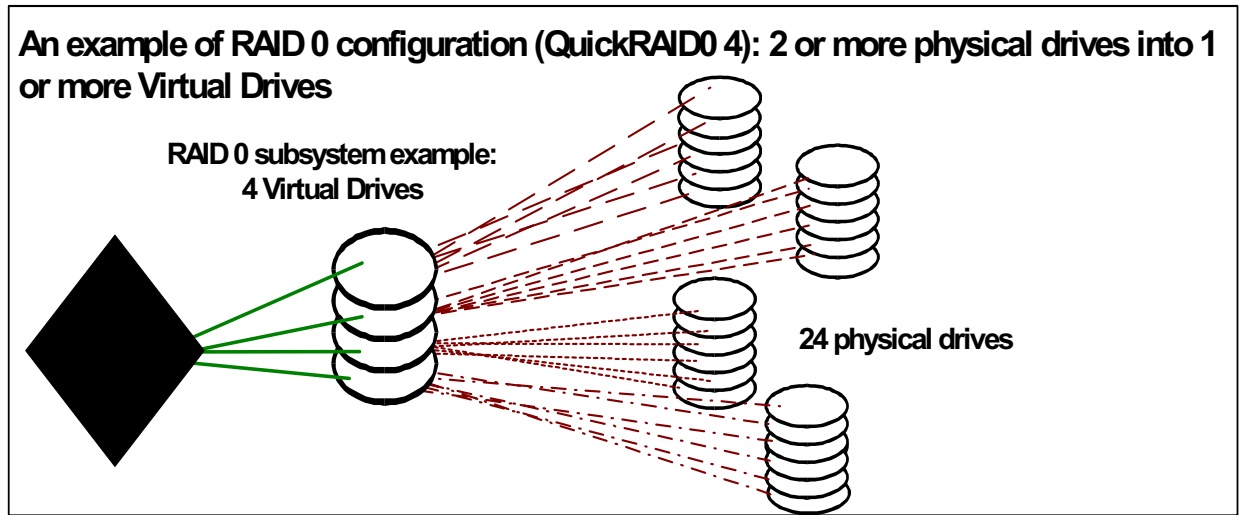


Exhibit 5.2-2 Sled-based QuickRAID0 stripe groups with LUN designations in a fully populated Diamond Storage Array. If the Array were set up as a QuickRAID0 6 SLED and sled 6 were to be withdrawn from the array, LUN 3 (grayed boxes) would be unavailable.

Physical Sled	1		2		3		4		5		6		7		8		9		10		11		12	
Physical Drive	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
QuickRAID0 parameters	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
	2	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	2
	3	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	3	3	3	3	3	3	3	3
	4	1	1	1	1	1	1	2	2	2	2	2	2	3	3	3	3	3	4	4	4	4	4	4
	6	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4	5	5	5	5	6	6	6
	12	1	1	2	2	3	3	4	4	5	5	6	6	7	7	8	8	9	9	10	10	11	11	12

LUNs: Virtual Drive numbers

Exhibit 5.2-3 Drive-based QuickRAID0 stripe groups with LUN designations in a fully populated Array. If the Array were set up as QuickRAID0 6 DRIVE and sled 6 were to be withdrawn from the array, LUNs 2 and 5 (grayed boxes) would be unavailable.

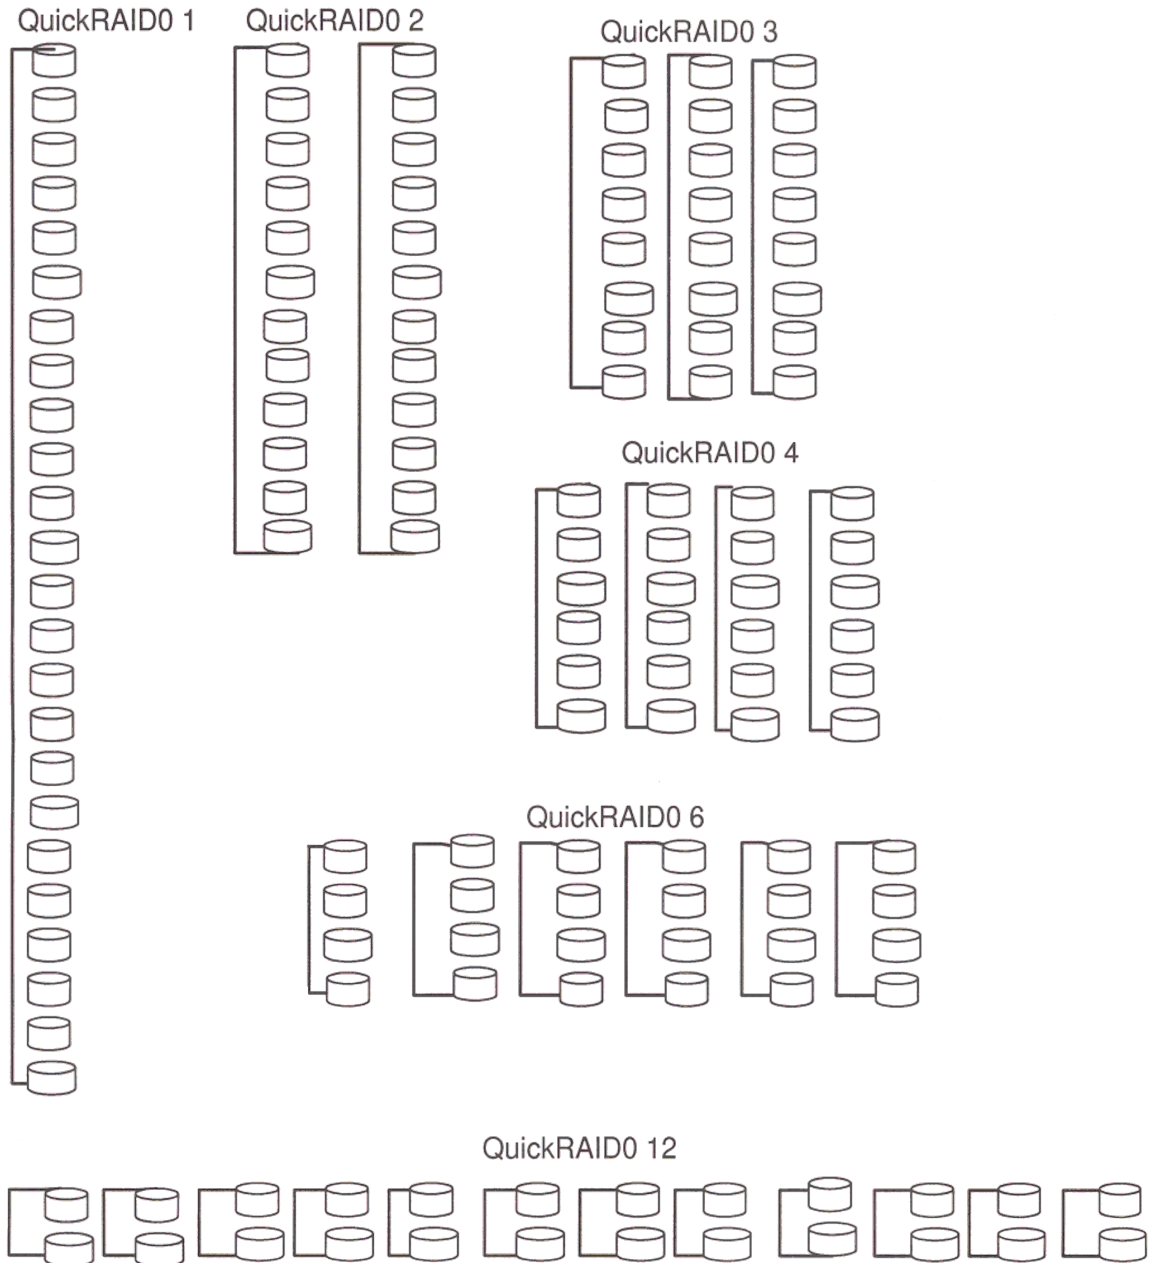
Physical Sled	1		2		3		4		5		6		7		8		9		10		11		12			
Physical Drive	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2		
QuickRAID0 parameters	JBOD	0	1	13	2	14	3	15	4	16	5	17	6	18	7	19	8	20	9	21	10	22	11	23	12	24
	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	
	3	1	2	1	2	1	2	1	2	1	3	1	3	1	3	1	3	2	3	2	3	2	3	2	3	
	4	1	3	1	3	1	3	1	3	1	3	1	3	2	4	2	4	2	4	2	4	2	4	2	4	
	6	1	4	1	4	1	4	1	4	2	5	2	5	2	5	2	5	3	6	3	6	3	6	3	6	
	12	1	7	1	7	2	8	2	8	3	9	3	9	4	10	4	10	5	11	5	11	6	12	6	12	

LUNs: Virtual Drive numbers

Exhibit 5.2-4 A fully populated array may be configured in several different ways in RAID Level 0.

RAID Level 0 striping groups

Each column represents one seen by the operating system as one virtual drive instead of 1 or more physical drives



5.3 RAID Level 1

The Diamond Storage Array can be set up into RAID Level 1 (mirrored) groups, with or without hot spare sleds, to provide greater reliability by simultaneously writing data to two sleds. Each sled partnered through QuickRAID1, a CLI command, has the same data as its partner.



CAUTION

Selecting RAID configuration causes all previous drive data on the Diamond Storage Array to be erased. Make sure all of your information is backed up before setting up RAID groups.

The configuration of RAID Level 1 performs the same operations on two partnered sleds at the same time, providing an automatic backup of data. The operating system sees the two sleds as one Virtual Drive with its own LUN (Logical Unit Number).

The **QuickRAID1** command allows the Diamond Storage Array to be set into mirrored drives. The command first spans two drives on a sled together, then partners two sleds to be a RAID Level 1 group designated by a LUN (logical unit number).

When you initially set up RAID groups using the **QuickRAID** command, groups are synchronized automatically because there is no pre-existing data on the drives. However, drives may display as “degraded, and you need to set all LUNs to OK status. Refer to *Rebuilding RAID Level Configurations* on page 63.



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

To set up RAID Level 1 groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in Command Line Interface.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 QuickRAID1 has no options: the command sets up each sled and its mirror image. Type `set QuickRAID1`.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline. In a fully populated array, your operating system displays you have six drives.



Note

In a less than fully populated array, if both partners of a LUN are missing, the LUN does not exist. If only one partner is missing, the LUN does exist, but it is degraded (unprotected). See Exhibit 5.3-1 to determine which LUNs would be affected.

To set up RAID Level 1 with Hot Spare sleds

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV browser interface **Advanced CLI** configuration page.
 - 2 Type `set QuickRAID 1 2`
 - 3 Type **SaveConfiguration Restart**
- The Diamond is configured into one RAID Level 1 group with two Hot Spare sleds.

To remove RAID groups

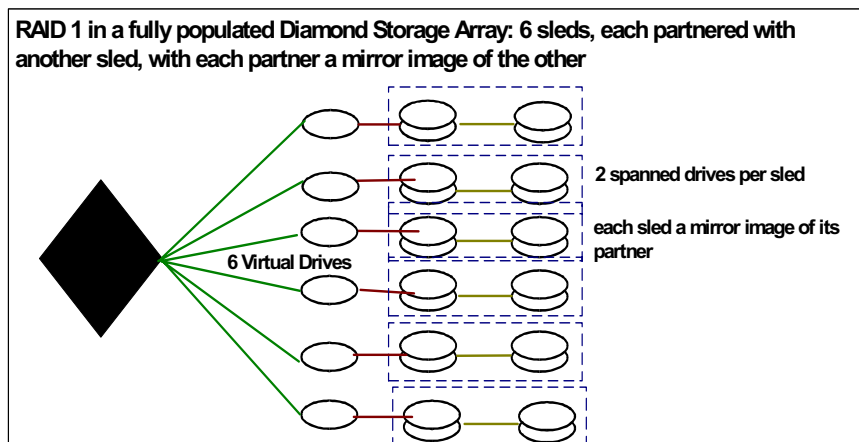
- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **set QuickRAID0 0** to configure the array in JBOD mode.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline

Exhibit 5.3-1 Mirrored stripe groups with LUN designations in a fully populated Diamond Storage Array. If sled 12 were removed, the drives marked LUN 6 would be available but degraded (unprotected by mirroring). If both sleds 11 and 12 were missing, LUN 6 would be unavailable.

Sled	1	2	3	4	5	6	7	8	9	10	11	12
Drive	1 2	1 2	1 2	1 2	1 2	1 2	1 2	1 2	1 2	1 2	1 2	1 2
LUN	1 1	1 1	2 2	2 2	3 3	3 3	4 4	4 4	5 5	5 5	6 6	6 6
Mirror	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2

Exhibit 5.3-2 Configuration of sleds in RAID Level 1: five LUNs with two Hot Spare sleds



5.4 RAID Level 5

RAID Level 5 increases reliability while using fewer disks than mirroring by employing parity redundancy. Distributed parity on multiple drives provides the redundancy to rebuild a failed drive from the remaining good drives.



CAUTION

Selecting RAID configuration causes all previous drive data on the Diamond Storage Array to be erased. Make sure all information is backed up before configuring RAID groups.

In RAID Level 5, parity data is added to the transmitted data at one end of the transaction, then the parity data is checked at the other end to make sure the transmission has not had any errors.

In the Diamond Storage Array, transmitted data with the added parity data is striped across disk drives. A hardware XOR engine computes parity, thus alleviating software processing during reads and writes.

The array uses parity declustering, a special case of RAID Level 5. Parity information is spread across each LUN, not concentrated on one drive or sled.

When you initially set up RAID groups using the QuickRAID command, groups are synchronized automatically because there is no pre-existing data on the drives. However, drives may display as “degraded, and you need to set all LUNs to OK status. Refer to *Rebuilding RAID Level Configurations* on page 63.

Configuring a fully-populated array

You may set up a fully-populated Diamond (12 sleds) into one, two, three or four RAID Level 5 groups, with or without hot spare sleds, using the QuickRAID 5 command.

To set up RAID Level 5 groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in the CLI.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

- 3 Decide how many RAID Level 5 groups you want (1, 2, 3 or 4).
- 4 Type **set QuickRAID5 [0|1|2|3|4]**
0 returns the array to JBOD
- 5 Type **SaveConfiguration Restart** to save the RAID Level 5 configuration.
- 6 A **Ready** prompt displays. You must zero all drives and parity to make all RAID Level 5 drives coherent. Type **RAID5ClearData all**



CAUTION

This is a destructive operation: all information stored on these drives is lost.

DO NOT interrupt power until the RAID5ClearData operation has completed (three to six hours).

- 7 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 8 Type **SaveConfiguration**.
- 9 Type **FirmwareRestart** to reboot the array.
- 10 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 11 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

To set up one RAID Level 5 group with one Hot Spare sled

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the Command Line Interface or enter ATTO ExpressNAV browser interface **Advanced CLI Configuration** page.

- 2 Type **set QuickRAID5 1 1**
- 3 Type **SaveConfiguration Restart**
The array is configured into one RAID Level 5 group with one Hot Spare sled.

the CLI or enter ATTO ExpressNAV browser interface **Advanced CLI Configuration** page.

- 2 Type **set QuickRAID5 2 2**
- 3 Type **SaveConfiguration Restart**
The array is configured into two RAID Level 5 groups with two Hot Spare sleds. Refer to Exhibit 5.4-1

To set up two RAID Level 5 groups with two Hot Spare sleds

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use

Exhibit 5.4-1 LUNs are set up using the drives and sleds shown here. Parity information is spread throughout each LUN.

Sled	1	2	3	4	5	6	7	8	9	10	11	12				
Drive	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
QuickRAID5 1	LUN 1															
QuickRAID5 2	LUN 1								LUN 2							
QuickRAID5 3	LUN 1				LUN 2				LUN 3							
QuickRAID5 4	LUN 1				LUN 2				LUN 3				LUN 4			

Exhibit 5.4-2 Drives are striped, and parity information is interspersed among the sleds.

Physical Sled	1	2	3	4	5	6	7	8	9	10	11	12				
Physical Drive	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
Step 1	Virtual Drive		1	2	3	4	5	6	7	8	9	10	11	12		
Step 2	Striping		→		→				→				→			
Step 3	Created LUN		LUN 1			LUN 2			LUN 3			LUN 4				

Exhibit 5.4-3 Configuration of sleds in RAID Level 5 in a fully populated array (12 sleds).

Sled	1	2	3	4	5	6	7	8	9	10	11	12		
Drive	1	2	1	2	1	2	1	2	1	2	1	2	1	2
QuickRAID5 1 1	LUN 1											Hot Spare		
QuickRAID5 2 2	LUN 1					LUN 2					Hot Spare	Hot Spare		

Configuring a partially-populated array

The simplest way to attain RAID Level 5 in a partially-populated array (an array with three or more sleds but less than 12 sleds) is to create one RAID Level 5 group, with or without hot spare

sleds, encompassing all the available sleds by using the CLI command QuickRAID5 ALL.

To set up a partially populated arrays, you must have at least three sleds filling contiguous slots as shown in Exhibit 1.28-4, beginning with the slot closest to the management card.

Exhibit 1.28-4 Examples of RAID Level 5 configurations in partially-populated array with at least six sleds. Hot Spare sleds are in the highest slot numbers.

Sled	1	2	3	4	5	6	7	8	9	10	11	12
QuickRAID5 4 3 contiguous sleds in slots 1-3 and slots 4-6 with no Hot Spare sled	LUN 1			LUN 2								
QuickRAID5 3 4 contiguous sleds in slots 1-4 and slots 5-8 with no Hot Spare sled	LUN 1				LUN 2							
QuickRAID5 2 6 contiguous sleds in slots 1-6 with no Hot Spare sled	LUN 1											
QuickRAID5 2 1 5 contiguous sleds in slots 1-5 with up to 2 Hot Spare sleds	LUN 1											Hot Spare
QuickRAID5 ALL 3-12 contiguous sleds with no Hot Spare sleds	LUN 1											
QuickRAID5 ALL 1 3-11 contiguous sleds with 1 Hot Spare sled	LUN 1											Hot Spare

To set up RAID Level 5 groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in the CLI.
- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

- 3 Decide the RAID Level 5 configuration you want based on the number of sleds you are using and Exhibit 1.28-4 above.
- 4 Type **set QuickRAID5 [2 | 3 | 4 | ALL]**

- 5 Type **SaveConfiguration Restart** to save the RAID Level 5 configuration.
- 6 A **Ready** prompt displays. You must zero all drives and parity to make all RAID Level 5 drives coherent. Type **RAID5ClearData all**



CAUTION

This is a destructive operation: all information stored on these drives is lost.



CAUTION

DO NOT interrupt power until the RAID5ClearData operation has completed (three to six hours).

- 7 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 8 Type **SaveConfiguration**.
- 9 Type **FirmwareRestart** to reboot the array.
- 10 Reboot the host computer connected via Fibre Channel or SCSI to the array.

- 11 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

To set up one RAID Level 5 group with one Hot Spare sled

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the Command Line Interface or enter ATTO ExpressNAV browser interface **Advanced** CLI Configuration page.
- 2 Type **set QuickRAID5 ALL 1**
- 3 Type **SaveConfiguration Restart**
The array is configured into one RAID Level 5 group with one Hot Spare sled.

Removing RAID groups

To remove RAID groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet

management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.

- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **set QuickRAID5 0**. This configures the array in JBOD mode.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

5.5 RAID Level 10

The Diamond Storage Array can be set up into RAID Level 10 (striped and mirrored) groups, with or without hot spare sleds, to provide greater reliability by simultaneously writing data to two sleds. Each sled partnered through QuickRAID10, a CLI command, has the same data as its partner:



CAUTION

Selecting RAID configuration causes all previous drive data on the Diamond Storage Array to be erased. Make sure all of your information is backed up before configuring RAID groups.

The configuration of RAID Level 10 stripes information across several mirrored drives, performing the same operations on two partnered sleds at the same time thus providing an automatic backup of data.


The **QuickRAID10** command, accessed through the Command Line Interface or ExpressNAV interface, first creates six mirrored groups, then stripes them into groups of one, two or three RAID 10 groups (see Exhibit 5-6).

When you initially set up RAID groups using the QuickRAID command, groups are synchronized automatically because there is no pre-existing data on the drives. However, drives may display as “degraded, and you need to set all LUNs to OK status. Refer to *Rebuilding RAID Level Configurations* on page 63

To set up RAID Level 10 groups

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.

An example of RAID 10 (**QuickRAID10 2**) in a fully populated Diamond Storage Array: 6 sleds, each partnered with another sled, with each partner a mirror image of the other, configured into two stripe groups or Virtual Drives



- 1 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

- 2 Decide how many RAID Level 10 groups you want (0, 1, 2 or 3).
- 3 Type **set QuickRAID10 [0 | 1 | 2| 3]**
0 returns the array to JBOD.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

To remove RAID groups


- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.
- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **set QuickRAID10 0** to configure the array in JBOD mode.

- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer.

- 8 The array is now configured. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

To set up RAID Level 10 with Hot Spare sleds:

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV browser interface Advanced CLI configuration page.
- 2 Type **set QuickRAID10 1 2**
- 3 Type **SaveConfiguration Restart**



- 4 The array is now configured into one RAID Level 10 group with two Hot Spare sleds. To verify the configuration, type **DriveInfo**. If any sleds are missing, the sled slots are reported as offline.

Exhibit 5-6 QuickRaid10 first spans drives across sleds, then partners sleds into mirrored groups, then enables striping across the mirrored groups.

Stripe	QuickRAID10 1 = LUN 1											
	Mirrored group 1		Mirrored group 2		Mirrored group 3		Mirrored group 4		Mirrored group 5		Mirrored group 6	
	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2
Sled	1	2	3	4	5	6	7	8	9	10	11	12
Drive	1	2	1	2	1	2	1	2	1	2	1	2

Stripe	QuickRAID10 2 = LUN 1						QuickRAID10 2 = LUN 2					
	Mirrored group 1		Mirrored group 2		Mirrored group 3		Mirrored group 4		Mirrored group 5		Mirrored group 6	
	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2
Sled	1	2	3	4	5	6	7	8	9	10	11	12
Drive	1	2	1	2	1	2	1	2	1	2	1	2

Stripe	QuickRAID10 3 = LUN 1				QuickRAID10 3 = LUN 2				QuickRAID10 3 = LUN 3			
	Mirrored group 1		Mirrored group 2		Mirrored group 3		Mirrored group 4		Mirrored group 5		Mirrored group 6	
	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2	Partner 1	Partner 2
Sled	1	2	3	4	5	6	7	8	9	10	11	12
Drive	1	2	1	2	1	2	1	2	1	2	1	2

All configurations are not available if the Diamond Storage Array has less than 24 physical drives in 12 physical sleds. For example, QuickRAID 10 1 only works with a

fully populated array. QuickRAID10 3 works if sleds 9-12 are removed: LUNs 1 and 2 are available but LUN 3 is not available.

5.6 Rebuilding RAID Level Configurations

If a sled must be removed and a new sled inserted into the Diamond Storage Array while it is configured in a RAID Level 1, 5 or 10, you must rebuild the RAID Level using CLI commands or the ExpressNAV interface.



WARNING

Selecting RAID parameters causes all previous drive data on the Diamond Storage Array to be erased. Make sure all of your information is backed up before setting up RAID groups.

When you initially set up RAID groups using the QuickRAID command, groups are synchronized automatically because there is no pre-existing data on the drives. However, drives may display as “degraded, and you need to set all LUNs to OK status.

The simplest method to check RAID group status is to access the **RAID** page of the ExpressNAV interface. Refer to *ATTO ExpressNAV: Browser-based Interface* on page 25.



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

To reset LUN status

- 1 Display the status of the array by typing **RAIDRebuildStatus**.
- 2 Set the sleds which are listed as degraded to a rebuild state of OK by entering
set RAIDRebuildState [sled number] OK

To synchronize mirrored drives automatically

If mirrored drives are removed for more than 15 to 30 seconds and then re-inserted or replaced, the replaced drives are labeled “degraded” when you check the array’s status by typing **RAIDRebuildStatus**. If you enable AutoRebuild, the array rebuilds the degraded drives automatically when a new drive is inserted.



Note

Drive rebuilding reduces performance. You may want to leave AutoRebuild at the default disabled and manually rebuild during off-peak hours, or use the RebuildPriority command described below.

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.
- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **AutoRebuild enabled**.
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 To verify the configuration, type **RAIDRebuildStatus**
- 9 If a rebuild is necessary, the array will automatically rebuild drives.

Rebuild priority

Drive rebuilding reduces performance. Use the RebuildPriority command through the CLI or on the **RAID** page of the ExpressNAV interface to customize when your rebuilds will occur.

Choices are High, Low and Same.

- If you select High priority, rebuild I/O requests are implemented before system I/O requests.
- If you select Low priority, rebuild I/O requests execute only when there are no pending I/O requests.
- If you select Same priority, rebuild I/O and system I/O receive equal consideration.

You must use a **SaveConfiguration Restart** command to implement the rebuild priority command.

To synchronize mirrored drives manually

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.
- 2 Continue with the CLI or access the RAID page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).

3 If **AutoRebuild** has not been disabled, type
`set AutoRebuild disabled saveconfiguration restart.`

4 After the array reboots and completes its diagnostics, enter the CLI and type **RAIDManualRebuild L S** where L is the LUN and S is the sled to be rebuilt.

This procedure may take a few hours, depending on the size of the LUN.

5 To check the rebuild status, type **RAIDRebuildStatus**

5.7 RAID Interleave

The interleave size sets the amount of data to be written to each drive in a RAID Level group. This is a tunable parameter which takes a single stream of data and breaks it up to use multiple disks per I/O interval.



WARNING

Changing this parameter causes all previous drive data on the Diamond Storage Array to be erased. Make sure you back up all information before setting up different interleave sizes.

The default sector interleave set by the QuickRAID command is 128 blocks (64k). The CLI command, **RAIDInterleave**, allows you to change the size of the sector interleave between RAID groups. The value depends upon the normal expected file transfer size. If the normal file transfer size is large, the interleave value should be large, and vice versa.

The value entered for the RAIDInterleave command refers to blocks of data: one block is equivalent to 512 bytes of data.

To change the RAID Interleave parameter

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in Command Line Interface.
- 2 Continue with the CLI or access the **RAID** page of the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).



Note

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

RAIDInterleave options are 16, 32, 64, 128, 256 blocks and span. Span is not available for RAID Level 5.

- 3 In all RAID levels except RAID Level 5, type

```
set RAIDInterleave [16|32|64|128|256|span]
```

For RAID Level 5 type

```
set RAIDInterleave [16|32|64|128|256]
```
- 4 Information displays on the screen while the array updates NVRAM ending with a **Ready***.
- 5 Type **SaveConfiguration**.
- 6 Type **FirmwareRestart** to reboot the array.
- 7 Reboot the host computer connected via Fibre Channel or SCSI to the array.
- 8 The array is now configured. To verify the configuration, type **get RAIDInterleave**.

5.8 Creating Partitions

With the introduction of larger and larger GB-sized drives, the Diamond Storage Array may have up to 6 TB total data capacity. Partitioning can increase storage efficiency by providing more LUNs without using lower capacity RAID groups.

Partitioning allows the creation of multiple logical volumes.

Using the Command Line Interface, you may divide an individual Partition into a set of equally-sized subpartitions which can then be presented to hosts as separate LUNs.



CAUTION

Before configuring the Diamond Storage Array, ensure that any computer data stored on the array is properly backed up and verified. The manufacturer is not responsible for the loss of any data stored on the Diamond Storage Array under any circumstances and any special, incidental or consequential damages that may result thereof.

If your Diamond Storage Array has been flashed with version 5.2 firmware (which allows partitions), and you flash the Array with a pre-5.2 version of the firmware, the configuration reverts to defaults. When you flash the Diamond to v5.1, all drives are offline and require a restamp to be accessed.

You will lose data in pre-existing RAID groups when you create partitions. Either back up the data to another storage area or only create partitions in data-free RAID configurations.

The array is set up in a JBOD (Just a Bunch of Disks) configuration as default and is available for normal storage operations immediately.

The array may be set up in a JBOD, RAID Level 0, RAID Level 1, RAID Level 10 or RAID Level 5 before partitions can be created.

Partitions allow better data management. For example, when using a RAID 5 configuration, two drives' worth of capacity are required for parity data for each physical LUN. Instead of creating four physical RAID Level 5 LUNs,

requiring eight drives of capacity for parity, you can create a single physical RAID 5 group with only two drives' worth dedicated to parity. You can then use partitioning to divide this single RAID 5 group into any number (up to 16) equally-sized, addressable LUNs.

If you do not want equally-sized partitions, you can merge partitions to create different capacity configurations. For example, if you create a Virtual Drive with 1TB capacity, then partition it into eight partitions of 128 GB each, you can merge several partitions into a larger partition. Partitions 2, 3, 4 and 5 could become a single LUN of 512 GB leaving partitions 1, 6, 7 and 8 at 128 GB. Partitions 3, 4 and 5 would no longer exist.

If you do not create partitions, the array reports a logical partition spanning the entire Virtual Drive by default. Each JBOD or RAID group is a Virtual Drive. A LUN is usually associated with a RAID group or Virtual Drive, but if you are using partitions, a LUN is associated with each partition. A RAID Group or Virtual Drive may then have multiple partitions or LUNs.

NOTE You must reboot the operating system to scan the array after any changes to the configuration.



Note

Do not configure the array into zones until after you have configured partitions.

If a hard disk drive in an existing Virtual Drive is replaced, all partitions that are a part of that Virtual Drive are labeled as degraded. When the Virtual Drive is rebuilt, all partitions are rebuilt.

To create a partition

- 1 The array must be configured to JBOD or the appropriate QuickRAID configuration before applying the Partition configuration. Zoning

may only be applied after Partition configuration.

- Each RAID group is a Virtual Drive. The array assigns an ID to each Virtual Drive. (refer to *RAID Level 1* on page 55, *RAID Level 10* on page 61 and *RAID Level 5* on page 57 about how to create Virtual Drives). Type `VirtualDriveInfo [active]` to determine the Virtual Drive ID and LUNs of any Virtual Drives already set up on your array.
- Type **VirtualDriveInfo [planned]** to determine if any partitioning has been planned but not completed. If a merge has been planned and you want the merge, go to step 5.
- You may create up to 16 partitions on any single Virtual Drive, with no more than 127 partitions across the entire array. Type

```
PartitionSplit [Virtual Drive ID]
[Partition ID] [number of Partitions]
```

- Type `PartitionCommit` to create the partitions. The array completes the configuration and reboots. The planned configuration you have entered becomes persistent and active.

- The array completes the configuration and reboots.

To merge partitions



CAUTION

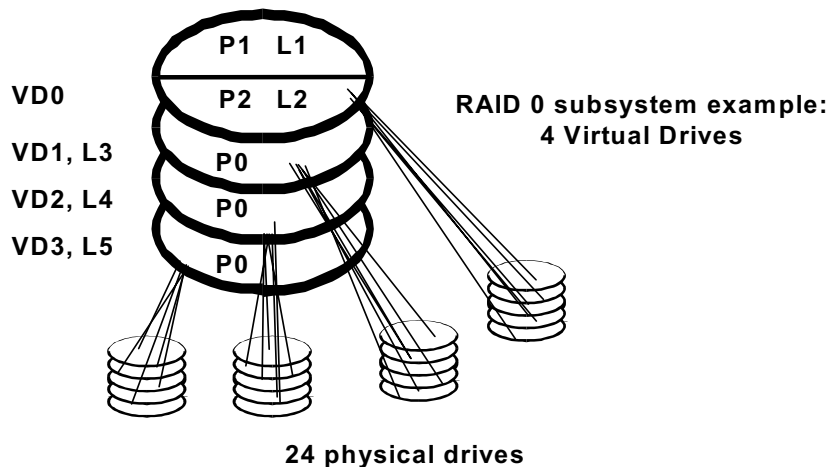
You may lose the ability to access data when you merge partitions. Either back up the data to another storage area or only merge partitions which are data-free.

- Type `PartitionInfo [active]` to determine the current partitions and their IDs.
- Type `PartitionInfo [planned]` to determine if partitioning had been planned but not completed. If a merge has been planned and you want that merge, go to step 4.
- Type `PartitionMerge [Virtual Drive ID] [Partition ID|all] [<Partition Number...>]` to create a planned partition configuration.
- Type `PartitionCommit` to merge the partitions and make the planned partition configuration the active configuration.
- The array completes the configuration and reboots.

Exhibit 5.8-1 An example of Virtual Drives set up in drive-based QuickRAID0 stripe groups.

An example of RAID 0 configuration (QuickRAID0 4): 2 or more physical drives into 1 or more Virtual Drives.

Virtual Drive 0 has been configured into two partitions: LUNs 1 and 2. Each Virtual Drive 1, 2 and 3 is configured by default as one partition and designated LUNs 3, 4, and 5



5.9 Creating Zones

Zoning is a collection of related Diamond Storage Array capabilities supporting flexible Diamond configuration management configurable via CLI commands in the Command Line Interface mode or in the Advanced CLI Configuration page in the ExpressNAV interface. Zoning supports security by granting or denying access between initiators and devices as defined by an administrator.

A zone is a collection of devices which can access each other. The devices in a zone usually include one or more initiators, one or more devices, and one or more paths between the initiators and the devices.

To set up zones, use the Command Line Interface (refer to *Accessing the Array* on page 17) or the **Advanced** CLI page in the ExpressNAV interface. Zone CLI commands only take effect after you enter the ZoneCommit command.

The individual elements are referred to as device_lun, host_name, port_number and zone_name as defined in Exhibit 5.9-1.

Exhibit 5.9-1 Definitions of zone configuration entries.

device_lun	The LUN of the JBOD or RAID drive
host_name	In a Fibre Channel environment, the WWPN; in a SCSI environment, SCSI Initiator ID
port_number	The Diamond port number (0, 1) for the data path
zone_name	Alphanumeric or '_', character string less than or equal to 16 characters long

Principles of Zoning

Zoning provides a validation filter for each SCSI command.

Each zone entry includes a named zone, a host portion, a port portion and a device portion. The components of a valid path from a host to a device satisfy the following conjunction:

<host_name> AND <port_list> AND device_list>

The zone is named to identify it from other zones. It appears as zone_name in this manual.

The host portion defines the valid access path from a host through a port to a device (LUN),

representing the Initiator ID in a SCSI environment or the World Wide Port Name in a Fibre Channel environment. It appears as host_name in this manual.

The device portion defines the LUN(s) participating in the zone. It appears as device_lun in this manual.

The port portion defines the Fibre Channel or SCSI port in the Diamond Storage Array. It appears as port_n in this manual.

The process: Each command received by the array is parsed to determine its host/HBA identifier, its port number and the target LUN, forming the zone nexus. This zone nexus is looked up in the defined zones table. If the zone nexus is present, the operation continues; if it is not found, the command is rejected with the appropriate status and sense data.

Two zone configurations accessed through the CLI regulate zoning:

The Planned configuration is a work-in-process configuration used to build or edit the desired configuration. The Planned configuration does not control I/O access until it is transformed into the Active Zone Configuration via successful completion of the ZoneCommit command.

Other than as a site for zone configuration editing, the Planned configuration has no impact on the Active configuration or the array. Changes to the Planned configuration may be made without considering synchronization with other configuration commands.

Use the ZoneClearAll command to clear the Planned configuration.

If, while working in the Planned configuration, you decide you want to negate that configuration

and edit the Active configuration, use the ZoneRetrieve command. The information from the Active configuration is copied into the Planned configuration.

To determine what is in the Planned configuration, type ZoneInfo Planned.

Active configuration is persistent and establishes the Diamond zoning configuration after power-up.

The Planned configuration becomes the Active configuration after successful execution of the ZoneCommit command. The Active configuration is replicated as the Planned configuration after the successful ZoneCommit, after power-up and after ZoneRetrieve. This replication eases incremental modifications to the current zone configuration: you only need to enter changes. Modifications to the Planned Configuration made before ZoneCommit are not persistent and are lost in the case of power-up.

To determine what is in the Active Configuration, type ZoneInfo.

Factors to consider

Several factors must be considered when configuring an array with zones:

- You must be careful when changing array zoning configurations. Internal validation logic cannot detect misconfigurations.
- The array zoning may be driven by external applications which handle considerations such as aliasing of parameters.
- Stopping or pausing I/O operations during zoning changes is the responsibility of the host computer, external to the array.
- You must refer to a device by a consistent LUN across zones in accordance with Fibre Channel specifications.
- JBOD/RAID configuration changes require planning and preparation independent of whether any zones are enabled. Such changes affect data integrity, and any write to an incorrect LUN may result in data corruption.
- Degraded operation and RAID rebuilding occur at a lower level than the Zoning features. SCSI command operation continues to operate, and you can modify the Zoning configuration via the ZoneCommit command at any time.
- Zone validation of switch/fabric routing is not supported. The array operates within a SAN environment including host systems, host bus adapters, switches and other devices. The Operating System you use may limit zoning flexibility in your SAN.
- You may create up to 32 zones. Each zone may have up to two ports, up to 24 devices and up to 32 hosts.

Status and Sense Data

Commands sent to a device may be rejected with sense key, code, qualifier as follows:

LOGICAL UNIT NOT CONFIGURED: 68 00 00	No such LUN exists; the initiator does not have access to this LUN
LOGICAL UNIT NOT SUPPORTED: 25 00 00	The Logical unit is not in an accessible zone
LOGICAL UNIT NOT READY: 04 03 00	The Logical unit is in an accessible zone, but is not available; it may be offline or busy

If the logical unit inventory changes for any reason, including completion of initialization, removal of a logical unit, or creation of a logical unit, the device server generates a Unit Attention command for all initiators, telling them a *ZoneCommit* procedure has been successful. The device server sets the additional sense code to

REPORTED LUNS DATA HAS CHANGED: 3F 0E 00 (LUN has been added to or removed from the zone)

Configuring Zones

Be careful when changing Diamond Storage Array zoning configurations. Internal validation logic cannot detect misconfigurations

An unrestricted zone configuration, exactly mimicking the LUN configuration, is created internally after the *first* power-up or restart after installation of the array. No special operating modes are required and Zoning can be easily installed with no impact on previous configurations. The unrestricted zone can be considered an all/all/all zone: all hosts, all ports and all devices.

Examples of initial configurations are available in *Sample Zoning Command Sequences* on page x.

To create a zone

- 1 The array must be configured to JBOD, the appropriate QuickRAID and/or Partition configuration before applying the Zoning configuration.



Note

Interpretation of the Zoning command is a single forward pass, so any entities referenced in any command must have been previously defined (ZoneCreate, followed by ZoneAdd, etc.) or you receive an error message.

- Type ZoneInfo Planned to determine the status of a Planned Zone Configuration.
- Type ZoneInfo to determine the status of an Active Zone Configuration.
- If you want to start fresh and create zones without reference to the definitions in the current Active Zone Configuration, type ZoneClearAll to remove all entries from the Planned zone definition table.

- 2 Create a new named zone. Type

```
ZoneCreate zone_name
```

- 3 Add the devices, hosts and ports you want to include in the zone.

```
ZoneAddDevice zone_name device_lun
```

```
ZoneAddHost zone_name host_name
```

```
ZoneAddPorts zone_name port_number
```

- 4 Enable the zone, type

```
ZoneState zone_name enabled
```

- 5 Type ZoneCommit to make this Planned Zone Configuration the Active Zone Configuration.

Refer to *Sample Zoning Command Sequences* on page x for sample configurations.

- 6 If the Diamond pauses operation by using a queue while a ZoneCommit command is executing:
 - a. The array completes any in-process I/O requests received before the ZoneCommit command was issued
 - b. The array performs the zoning changes
 - c. The array resumes I/O operations.

To remove zones

- 1 Remove the named zone. Type

```
ZoneRemove zone_name
```

- 2 Type ZoneCommit to commit this Planned Zone Configuration to become the Active Zone Configuration.

To change current zones



CAUTION

Be careful when changing Diamond Storage Array zoning configurations. Internal validation logic cannot detect misconfigurations

- 1 Type ZoneInfo Planned to determine the status of a Planned Zone Configuration.
- 2 Type ZoneInfo to determine the status of an Active Zone Configuration.
 - a. If necessary, type ZoneRetrieve to overwrite the Planned configuration with the current Active configuration.
- 3 Create or delete named zones as outlined above.
- 4 Add or delete devices, hosts and ports as outlined above.
- 5 Type ZoneCommit to make this Planned Zone Configuration the Active Zone Configuration.

Other operations

- To disable a zone, type

```
ZoneState zone_name disabled
```

- To clear the Planned Zone Configuration of all entries, type

```
ZoneClear All
```

- Using RestoreConfiguration default or RestoreConfiguration factory default does not affect the zoning configuration. To restore the array to factory default, type

```
ZoneClearAll
```

```
ZoneCommit
```

Errors

Zone definition tables: The Zone commands manage entries in the Zone definition tables which manage the overall zoning process. Definition tables are indexed by unique keys (zone_name). All definition tables are repositories for their respective data and participate in establishing the configuration by executing the ZoneCommit command.

The integrity of these tables is essential to the data integrity of the array. If the Zone definitions are faulty problems can occur. Although Zone command processing provides a level of command and configuration validation, you must be very careful managing configurations with any zoning system.

Validation of the command line is performed before the command is deemed acceptable to be executed.

These descriptions are generalizations.

- Configuration inaccuracies occurring before application or operating system data is written to the drives may have no impact on the array, but results after an operating system or application have written to the drives are unpredictable.
- Incorrect Zone entries can include mis-specification of resources to a zone.
- Verify each command line is properly formed (number of parameters, proper spelling of keywords).
- For commands defining entities, the name being defined must not already be defined.
- If you want to undefine an entity, the name being undefined must already be defined. (A warning displays if the name to be undefined doesn't exist.)
- World wide port names are validated according to basic format rules. Content verification of WWPN occurs at runtime.
- LUN must be in the range defined by the JBOD/RAID<n> configuration.
- Errors detected in the CLI command line are described in Exhibit 5.9-1.
- Errors detected while writing the Active Zone Configuration result in an error message and no change to the zoning configuration. The Active configuration continues to match the persistent configuration.
- Errors detected while reading the Active Zone Configuration result in an error message and the zoning configuration remains unchanged.

6.0 Copying Drives

Copying drives using the DriveCopy CLI command may be necessary on drives in the JBOD configuration. RAID Level 1 and RAID Level 10 configurations already provide mirroring of drives.

DriveCopy can create a backup of a drive onto another drive sled. If one drive fails, DriveCopy may be used to copy the data off the remaining drive on the sled to another drive on another sled. Once completed, the sled containing the failed drive may be replaced.



CAUTION

Make sure the destination backup drive sled does not contain any important data because it will be completely overwritten by the DriveCopy command.

- 1 Use the DriveCopy CLI command (refer to *Drive Configuration Commands* on page 43) to copy a drive from the source disk to the destination disk. Parameters are the sled and drive numbers of the source and destination drives.
- 2 The destination drive must be offline: use the ATADiskState command to determine if the destination disk is offline. If you choose the same source and destination drive, this command performs a DriveWipe.

DriveCopyHalt stops a DriveCopy operation in progress.

DriveCopyResume resumes a DriveCopy operation that had been stopped.

DriveCopyStatus displays the status of a DriveCopy or DriveWipe operation

JBOD configurations: You may want to copy drives for backup. However, since several configurations may be present on a Diamond Storage Array, you must be confident you are accessing and overwriting the appropriate drives and volumes.

RAID Level 1 and RAID Level 10 configurations should not need to be backed up by copying drives since these settings already provide mirrored copies of drives. However, you may copy a drive to another drive within a RAID Level 1 or RAID Level 10 system, but the destination drive cannot contain data you want to save.

You might want to keep an entire group (LUN) free of data to use as spare drives within the array system.

RAID Level 0 and RAID Level 5 systems

DriveCopy may be used to coordinate the generation of a backup of an entire RAID Level 0 or RAID Level 5 LUN.

7.0 Updating Firmware

Engineers, technicians and/or system administrators/integrators may update the firmware of the Diamond Storage Array using the Command Line Interface (CLI) (refer to *Accessing the Array on page 17*) via the RS-232 serial port or the optional Ethernet management services card.

Updating firmware via the RS-232 serial port

To update the firmware via a connection to the RS-232 serial port, you need

- a host computer with a terminal emulation program such as HyperTerminal in Windows
 - binary information file, "...".ima, available from technical support or on our website, www.attotech.com
 - a null modem serial cable with a DB-9 connector
- 1 Connect to Diamond Storage Array services via the RS-232 port (refer to *Accessing the Array on page 17*). You should now be in the Command Line Interface mode.
 - 2 Copy the latest array image file, "...".ima, onto the host computer.
 - 3 Type **ZModem Receive** at the **Ready** prompt. The terminal program on the host should be in Zmodem only mode, with no other parameters. The array displays information that it is preparing to receive a file from your terminal program.
 - 4 On the terminal program, choose

Transfer Send File.

- 5 In the **Send File** box, enter the current Diamond flash, "...".ima, filename or click the **Browse** button to find it.
- 6 Click **Send File**
- 7 The array acknowledges receiving the file and displays a message not to interrupt power for 90 seconds.



CAUTION

Do not interrupt the flash process. If the process is interrupted, the Diamond Storage Array becomes inoperable and must be returned to the factory for repair.

Do not turn off the Diamond Storage Array until the display returns the Ready prompt.

If upgrading the firmware from versions older than 2.5.3, follow the procedures outlined in *Resetting Defaults on page 83*.

- 8 When the flash procedure is complete, cycle power on the array.

Updating firmware via the optional Ethernet card

To update the firmware via the optional Ethernet management services card, you need

- the optional Ethernet management services card installed in your array
- a host computer with a network card or a network-connected device such as a hub
- binary information file, "...".ima, available from technical support or on our website, www.attotech.com
- a crossover network cable for a direct connect or standard network cable if attached to a network device

- a valid IP address
- 1 Connect a cross-over cable (for a direct connection to a PC) or regular network cable from a network device to the optional RJ45 Ethernet port on the Ethernet management card on the front of the array.
 - 2 Power on and boot up the host computer. You may also attach a DB-9 null modem serial cable from the RS-232 port of the array to a host computer and open a terminal emulator program on the host to set the Ethernet parameters.
 - 3 Turn on the array.

- 4 Copy the latest array image file, "...".ima, onto the host computer and note its directory such as

```
c:\diamond\flash\"...".ima
```

- 5 *First time use:* Upon successful power up and diagnostics, set the host computer with the appropriate settings such as IP Address.

The host computer must have appropriate network settings to allow it to communicate with the array. Please see your system administrator for more information.

- 6 Change directories to the place where you copied the "...".ima file such as

```
cd c:\diamond\flash
```

- 7 Open a FTP session using a user-defined IP address. The IP address must be a valid address for your network.
- 8 At the FTP login prompt, type **sysadmin** as the userID.

- 9 Press **return** at the password prompt.

- 10 Type **put "...".ima**.

The array should acknowledge receiving the file and display a message not to interrupt power for 90 seconds.



CAUTION

Do not interrupt the flash process. If the process is interrupted, the Diamond Storage Array becomes inoperable and must be returned to the factory for repair.

Do not turn off the Diamond Storage Array until the display returns the Ready prompt.

If upgrading the firmware from versions older than 2.5.3, follow the procedures outlined in *Resetting Defaults* on page 83.

- 11 When the flash procedure is complete, cycle power on the array.

8.0 System Monitoring and Reporting

The Diamond Storage Array provides a number of visual, audible and computer system-generated indicators to identify the operational status of the array. System status and error information is readily available.

RS-232 Monitoring Port and CLI Use a host computer with an RS-232 port and terminal emulation software to connect a null modem serial cable to the RS-232 port on the array management card and control the array via the CLI management software built into the array. The CLI can be used to configure the unit, modify key parameters and read back key system information. (refer to *Accessing the Array* on page 17)

Ethernet Monitoring Port and CLI If the optional Ethernet management system card has been installed, use the RJ45 Ethernet port and Telnet to access the CLI software on the array. Use the same CLI commands to configure, modify or read key systems information (refer to *Accessing the Array* on page 17) or the ExpressNAV browser-based interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25.) If you restart the array, the Telnet session ends and the session cannot be re-established until the array completes the POST.

Power On Self Test (POST) Each time the array is powered up, it performs a series of internal tests called POST. The POST sequence takes from one to two minutes to complete.

LEDs blink at various points in the test sequence and, if the RS-232 system management port is connected, a list of tests and test results scrolls across the screen. When the POST is nearly complete, all the LEDs on the array flashes simultaneously twice. If the system is fully operational, the ready LED lights. If the POST fails, the System Fault LED lights.

Ready LED The Ready LED (RDY) indicates the operational status of the array. At power up, a green Ready (RDY) LED is disabled until successful completion of the POST.

If a host computer is connected to either a Fibre Channel or a SCSI port, the Ready LED blinks, then stays lit if the connection is good. If the host is rebooted, the Ready LED blinks and stays lighted when the connection is reestablished.

Audible Alarm The audible alarm beeps twice at power up and beeps repeatedly when the System Fault (FLT) light on the system management card is activated. The alarm turns off when the fault condition is cleared or you have disabled the alarm by issuing a PowerAudibleAlarm or AudibleAlarm command via CLI as described in *Diagnostic Commands* on page 41. Disabling the alarm for a specific error does not silence the alarm for other errors.

Thermal Monitoring The array provides advance warning of temperature problems through visual, audible and software warning mechanisms and through an automatic system which protects the disk drives under abnormal conditions.

Temperature sensors provide data to the software to trip a temperature warning alarm and, at excessively high or low temperatures, flushes cache memory to prevent data loss and disables disk drive activity to protect the drives. The Diamond reports the temperature and the state of the warning (Not Present, OK, Warning, Critical) through SES, SNMP, CLI and the **Status** page of the ExpressNAV interface. The temperature warning alarm reports OK during normal operating conditions.

If an abnormal operating condition, such as blower failure, occurs and the array internal midplane temperature reaches a critical point, the temperature alarm reports Warning and activates the audible alarm and fault LED. If the internal midplane temperature reaches a higher point, the temperature warning alarm reports Critical, the array is taken off line, and all disk drive activity is

disabled. When the ambient temperature decreases to within standard operating range, the drives are powered back on and the host is allowed to access data. Typical Diamond Storage

Array operation over a range of external ambient temperatures. Alarms are audible, visual, CLI and SES cues as described above.

Exhibit 8.0-1 Typical Diamond Storage Array operation over a range of external ambient temperatures. Alarms are audible, visual, CLI and SES cues as described above.

Temperature	25°C	25°C	32°C	32°C	35°C	35°C	40°C	40°C
Condition	blowers OK	1 blower fails	blowers OK	1 blower fails	blowers OK	1 blower fails	blowers OK	1 blower fails
Read/Writes	normal	normal	normal	normal	normal	normal	normal	Array offline
Alarms	OK	OK	OK	OK	OK	WARNING	OK	CRITICAL

Power Supply Monitoring The array monitors the operation of its power supplies and blower assemblies every 30 seconds. If the status changes, the system reports it visually with an LED and a message in the CLI.

If a power supply or blower fails, the management system sends a CLI message and turns off the corresponding power supply LED.

The blower assembly directly adjacent a power supply must be functioning properly for the power supply to work properly. If a blower assembly fails, the power supply shuts down and the management system sends a CLI message. The corresponding power supply LED on the system management card also turns off.

If a good replacement power supply or blower assembly is inserted into the array, the management system sends a CLI message and the corresponding system management card LED lights. It may take up to 30 seconds for the system to note these changes.

System Fault LED and Error Codes If a serious hardware or software error occurs in the array, the System Fault LED displays a series of flashes or blink codes. Error information is reported via the CLI if it is operational.

At power up, the green array Ready (RDY) LED turns on after successful completion of the POST (Power On Self Test) indicating the array is available for normal operation.

During operation, the Ready LED stays lit even if the amber Fault (FLT) LED lights.

If a serious error occurs in the array, the amber Fault LED flashes repeatedly in a blink code pattern:

- an initial series of blinks indicating the system problem

Number of blinks	Problem area
1, 2 or 3	processor or memory
4	Fibre Channel interface
5	SCSI interface
7	Fibre Channel connection
8	general internal processing
9	SCSI Enclosure Services

- a two second pause
- another series of blinks providing more detailed information for technical personnel.
- a four-second pause
- the blink code sequence repeats from step 1 until the error is cleared.

In general, any fault requires notification of Diamond Storage Array technical personnel for resolution or for further debug instructions. When you report an error code, please provide both the first and second blink code values.

During a fault condition, more detailed information about the fault may be available via the CLI or the ExpressNAV interface over the RS-232 interface port or the optional Ethernet port.

These error messages should be reported to technical personnel to assist in debugging the problem.

The blink codes are also saved internally by the array to NVRAM (Non Volatile Random Access Memory) and are displayed at power up if the power to the array is recycled.

Disk Drive Activity and Disk Fault LEDs Each dual disk drive sled assembly contains two green activity LEDs and an amber Disk Fault LED. Once the system has successfully powered up and passed POST, the green activity LEDs are full on and the Disk Fault LED off.

The Drive 1 and Drive 2 activity LEDs stays full on when the system is operational and no disk drive activity is present. As the disk drives are accessed, the green LEDs flashes. If the disk drives are heavily accessed the green activity LEDs appears to flash at a high rate or may even appear to be completely turned off.

The amber Disk Fault LED is off under normal operation. If either of the disk drives on a dual disk drive assembly reports a disk error of any kind, the amber Disk Fault LED lights. The Disk Fault LED can be activated by minor issues such as a disk drive writing to a bad sector (which is

usually corrected by the disk drive the next time it writes) or major issues such as a head crash or complete drive failure.

When the Disk Fault LED is turned on, the system issues a detailed message via the CLI. These messages are not written permanently to the error log file but should be recorded to help assess the disk problem.

If the disk drive error is a non fatal error and the drive is still functional, the array continues to read and write data to the disk drive but the Disk Fault LED remains on. If you repeat a drive command or action and it completes successfully, the Disk Fault LED may have been set by an anomaly in the disk drive. You can clear the Disk Fault LED by either power cycling the array or issuing the SledFaultLED command in CLI as per *Diagnostic Commands* on page 41

If you repeat a disk command or action and the Disk Fault LED remains on, the disk error may be serious. You should write down the error message issued by the CLI and contact technical support via the means easiest for you for assistance (refer to *Warranty* on page xvi). If you choose to replace the suspected faulty dual disk drive sled assembly, follow the appropriate procedures.

8.1 Troubleshooting

The Diamond Storage Array provides a number of visual, audible and computer system-generated indicators to identify the operational status of the array. If your situation is not defined here or elsewhere in the manual, if these solutions do not help, or if you have any questions or concerns about any aspect of operating the array, contact technical support.

Windows 2000 special instructions

When using Windows 2000, the **New Hardware** screen pops up and asks for a driver when the Diamond Storage Array is first booted up. While a driver is not necessary for operation, you should install our dummy driver to eliminate the **New Hardware** screen's appearance. Download **AttoDM2k.zip** from our website, www.attotech.com, unzip it, and install the driver **AttoDM2k.inf** according to the instructions in **AttoDM2k.pdf**.

Error Messages

System Fault LED

If a serious hardware or software error occurs in the array, the System Fault LED displays a series of flashes or blink codes in a pattern.

- an initial series of blinks indicating the system problem
- | Number of blinks | Problem area |
|------------------|-----------------------------|
| 1, 2 or 3 | processor or memory |
| 4 | Fibre Channel interface |
| 5 | SCSI interface |
| 7 | Fibre Channel connection |
| 8 | general internal processing |
| 9 | SCSI Enclosure Services |
- a two second pause
 - another series of blinks providing more detailed information for technical personnel.
 - a four-second pause
 - the blink code sequence repeats from step 1 until the error is cleared.

In general, any fault requires notification of Diamond Storage Array technical personnel for resolution or for further debug instructions. When you report an error code, provide the first and second blink code values.

During a fault condition, more detailed information about the fault may be available via the CLI or the ExpressNAV interface over the RS-232 interface port or the optional Ethernet port.

The blink codes are also saved internally by the array to NVRAM (Non Volatile Random Access Memory) and are displayed at power up if the power to the array is recycled.

Command Line Interface messages

ERROR. Wrong/Missing Parameters

Check **Help** for the correct input and retype command

ERROR. Invalid Command. Type 'help' for command list

Check **Help** to find a list of all commands which are available. Contact technical support via the means easiest for you. Refer to *Warranty* on page xvi for additional information.

ERROR. Command Not Processed.

The array did not accept the command you requested. Check **Help** for a list of commands or check this manual for the function you wish to access. If you cannot accomplish what you want to do with the commands listed, contact array technical support via the means easiest for you (refer to *Warranty* on page xvi) for more information.

Audible Alarm

The array audible alarm warns of potential problems or faults. It beeps repeatedly when the System Fault (FLT) light on the system management card is activated. The alarm turns off when the fault condition is cleared or the alarm can be disabled by issuing a

PowerAudibleAlarm or **AudibleAlarm** command via CLI as described in *Diagnostic Commands* on page 41. Disabling the alarm for a specific error does not silence the alarm for other error conditions.

Specific situations and suggestions

For all problems, first check the pages of the ExpressNAV interface appropriate for the problem especially the **Status** and **Storage**

Management or use the appropriate CLI commands

If a drive fails to respond

- Determine which drive has failed by observing the amber fault LED on the drive sled or connect to the CLI and type `driveinfo` for a list of all drives and their status.
- For specific information on a particular drive in CLI, type `driveinfo [sled number] [drive number]`. Record all errors.
 - Determine if the drive is configured in JBOD, RAID Level 0 or RAID Level 1.
 - Follow the instructions in the appropriate chapters of this manual for removing, replacing and reconfiguring the drive (refer to *Hot Swap Operating Instructions on page 87*)
- You may copy drives by using CLI commands. (refer to *Copying Drives on page 73* and *Drive Configuration Commands on page 43*).

If a power supply fails

- Verify the power cord is correctly plugged and there is power at the power receptacle.
- If there is power, the cord is secure and the blower and power supply LEDs are off, replace the blower unit (refer to *Hot Swap Operating Instructions on page 87*).
- If there is power, the cord is secure and the power supply LED is off, but the blower LED is on, replace the power supply (refer to *Hot Swap Operating Instructions on page 87*)
- Command Overlap: The array contains special software to take advantage of the Command Overlap feature offered in some high performance disk drives. Because all 24 disk drives in a array may be seeking at the same time under Command Overlap, only arrays with two operational power supplies and blower assemblies supports Command Overlap: the array automatically disables the Command Overlap feature if a power supply fails and notifies the system administrator through the audible alarm and CLI.

If you cannot read or write to the array: You may have lost connection to the host via the Host Interface Card. The Host Interface Card LED on the back of the array should be lit green. The amber LED should go out and the green LED should light when the connection is complete.

Also check the host bus adapter (HBA) in the host machine for proper functioning and drivers.

- Verify that the connector and both ends of the cable are completely seated.
- Try connecting directly to the host, bypassing any hubs or switches.

To determine if the problem exists with the Host Interface Card or the connection

- Swap the cable from one HIC to the other HIC.
- If the LED is now green, swap the cable back to the original HIC. If the LED is amber, the HIC is defective.
- Swap the defective HIC (refer to *Hot Swap Operating Instructions on page 87*).

If you can't access the array CLI via Ethernet

- Verify there are three or fewer concurrent sessions using Telnet or FTP. You may be the fourth session or someone in another session has entered a command which requires a **SaveConfiguration** command.
- Verify you are using a crossover cable for direct connection, or a network cable for a network connection.
- Verify the array is set to **IPDHCP enabled** if you expect an IP address will be assigned by the network and DHCP is an option available on the network.
- Verify the IP address is compatible with the host machine.
- Verify the system is functioning and accessible via in-band inquires such as Disk Management.
- Try setting your terminal emulator with a different baud rate, starting at 2400, then 9600, 19200, 38400, 57600, and 115200.
- If you perform a hardware restart, the Ethernet connection is dropped. You must re-establish the Ethernet connection.
- If you enter a command that requires a **SaveConfiguration** command in either the serial interface window or the Ethernet connection window, you cannot not be able to access the inactive window until the **SaveConfiguration** command is complete in the active window.
- If you still fail to communicate with the array, swap out the management card and try to connect using default settings.

If you do not see the appropriate number of LUNs on the host machine

- Ensure any configuration changes are appropriate
- Type **FirmwareRestart** in CLI
- Re-start the host computer
- Verify all drives associated with the missing LUN(s) are inserted properly and powered up.

8.2 Resetting Defaults

Resetting the Diamond Storage Array to defaults does not alter the RAID configuration, zoning configuration, IP configuration or Telnet information. However, resetting the array to factory defaults is a last-ditch effort to recover from corrupt configurations or complete failure. All data is lost, but the zoning configuration remains.

Default

If you need to return to the default settings of the Diamond Storage Array but do not want to lose data or Ethernet settings, use the CLI command `RestoreConfiguration default` in CLI mode or in the *Advanced CLI* configuration page of the ExpressNAV interface.



Note

Resetting the array to factory defaults is a last-ditch effort to recover from corrupt configurations Using **RestoreConfiguration default** or **RestoreConfiguration factory default** does not affect the zoning configuration.

To restore the array to factory default, type

ZoneClearAll

ZoneCommit

Because the ExpressNAV pages take you through this process easily, the following instructions are based on the CLI commands. Use these instructions as a guide in ExpressNAV.

Return to Default settings

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in Command Line Interface.
- 2 Continue with the CLI or access the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **RestoreConfiguration default** to reset the system configuration (See Exhibit 8.2-1 for a list of configurations which change).
- 4 Type **FirmwareRestart** or cycle power of the array.

- 5 Reboot the host PC after the array completes its power on cycle.

Factory Default



Note

Resetting the array to factory defaults is a last-ditch effort to recover from corrupt configurations Using **RestoreConfiguration default** or **RestoreConfiguration factory default** does not affect the zoning configuration.

To restore the array to factory default, type

ZoneClearAll

ZoneCommit



CAUTION

Data will be lost if you follow these procedures. Make sure you have no other choice before resetting the array to factory defaults.

To reset to Factory Defaults, firmware version 2.5.3 or higher

- 1 Connect to Diamond Storage Array services via the RS-232 port or the optional Ethernet management services card (refer to *Accessing the Array* on page 17). You should now be in CLI.
- 2 Continue with the CLI or access the ExpressNAV interface (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 3 Type **RestoreConfiguration factorydefault** to reset the system configuration. (See Exhibit 8.2-1 for a list of configurations which change).
- 4 Type **FirmwareRestart** or cycle power of the array.
- 5 Reboot the host PC after the array completes its power on cycle.

Exhibit 8.2-1 Configurations which change during a **RestoreConfiguration** command

Command	Default	Reset during default	Reset during factorydefault
AudibleAlarm	Disabled	✓	✓
AutoRebuild	Disabled		✓
DiamondName	“ ”		✓
EthernetSpeed	Auto		✓
FcConnMode	Loop		✓
FcDataRate	Auto	✓	✓
FcFairArb	Enabled	✓	✓
FcFrameLength	2048	✓	✓
FcFullDuplex	Enabled	✓	✓
FcHard	Disabled	✓	✓
FcHardAddress	0x03	✓	✓
IdentifyDiamond	Disabled	✓	✓
IdeTransferRate	4	✓	✓
IPAddress	10.0.0.1		✓
IPDHCP	Disabled		✓
IPGateway	0.0.0.0		✓
IPSubnetMask	255.255.255.0		✓
MaxEnclTempAlrm	47	✓	✓
MinEnclTempAlrm	5	✓	✓
PowerAudibleAlarm	Enabled	✓	✓
QuickRAID0	0		✓
QuickRAID1			✓
QuickRAID10			✓
QuickRAID5			✓
RAIDInterleave	128		✓
SerialPortBaudRate	115200	✓	✓
SerialPortEcho	Disabled	✓	✓
SerialPortHandshake	None	✓	✓
SerialPortStopBits	1	✓	✓
SNMPSendTrap	Disabled	✓	✓
SNMPTrapAddress	10.0.0.1	✓	✓
SNMPTraps	4	✓	✓
SNMPUpdates	Disabled	✓	✓
TelnetPassword	diamond		✓
TelnetTimeout	Disabled		✓
TelnetUsername	telnet		✓
VerboseMode	Enabled	✓	✓

9.0 Hardware Maintenance

The disk drive sleds, blower assemblies, power supplies, host interface cards, and system management card may be replaced with identical or upgraded parts.




CAUTION

Do not leave empty openings on the front or rear of the Diamond Storage Array under any circumstances. Empty openings affect airflow and may cause the unit to overheat and shut down.



WARNING

The only way to completely de-energize the unit is



to turn off both power supplies and unplug both power cords from the back of the unit. Turning the power switch to the Stand-by position on one power supply does not completely turn off power

to the array; it is not an AC on-off switch. Power may still be in the unit through the other power supply.

All modular components must be replaced by qualified personnel only. Use a static wriststrap when handling any of the cards inside the Diamond Storage Array. Components are electrostatic sensitive. Use proper grounding methods when working with or around the Diamond Storage Array. Always store spare components in proper ESD containers when not in use.

- The power supply and blower assembly may be replaced while the unit is running. (refer to *Hot Swap Operating Instructions on page 87*)
- Host interface cards and management cards may only be replaced when the array is **off**. Backup the unit fully before replacing these components.
- You may remove a disk drive sled while the array is powered on. Refer to the instructions in *Hot Swap Operating Instructions on page 87* for details.

Management card: To remove a management card, power down both power supplies, loosen the screws holding the card in place, pull out the assembly and replace it with another. Securely tighten all screws after replacing the component.



Exhibit 9-1: The management card may be accessed via a serial port DB-9 connector or an optional Ethernet connection.

Disk drive sled: To remove a disk drive sled (Exhibit 9-2), loosen the screws on either side of the assembly, then pull on the assembly's handle and carefully slide it out of its bay.

FC or SCSI Host Interface Card To remove a Fibre Channel or SCSI Host Interface Card from the back of the array (Exhibit 9-3), power down both power supplies and remove any cable attached to the port. Loosen the retaining screws and pull the Host Interface Card out of the unit. To replace the card, push it back into the unit and tighten the retaining screws.

Power supply: To remove the power supply (Exhibit 9-4), press the Stand-by power switch to the off position, remove the power cord, and, using a No. 1 Phillips screwdriver, loosen the screws holding the assembly in place. Pull out the assembly and replace it with another. Securely tighten all screws after replacing the component.

Blower assembly: To remove a blower assembly (Exhibit 9-4), using a No. 1 Phillips screwdriver, loosen the screws holding the assembly in place. Pull out the assembly and replace it with another. Securely tighten all screws after replacing the component.




Exhibit 9-2: Above, disk drive sled partially pulled out of the Diamond Storage Array. Bottom left, top of disk drive sled. Bottom right, underside of disk drive sled showing individual drives




Exhibit 9-3: The Fibre Channel or SCSI Host Interface Card may be replaced by shutting power down, removing any cable attached to the port, removing the SFP according to manufacturer's instructions, loosening the screws at the top and bottom of the card, then carefully pulling out the unit.



CAUTION

Do not leave empty openings on the front or rear of the array under any circumstances. Empty openings affect airflow and may cause the unit to overheat and shut down.



Exhibit 9-4: A power supply pulled out from a rackmount Diamond Storage Array: do not leave an empty opening while the Diamond Storage Array is operating. Access the blower assembly and the power supply by loosening the screws on either side of the component, then pulling out the part. If a blower or power supply stops working, keep it in place until another component is installed. (Note: host interface configuration pictured is not supported)

9.1 Hot Swap Operating Instructions

To maintain array up time, individual disk drive sled assemblies, power supplies and blower assemblies can be replaced with the unit fully operational. Special instructions need to be followed to perform these operations.

Disk Drives

CAUTION

Individual disk drive sled assemblies may be replaced while the array is operating with no other intervention only if there is absolutely no activity on that drive. Failure to ensure no activity may destroy any data on that drive and possibly stop the entire Diamond Storage Array operation.

Follow the instructions below to replace drives.

Removing a drive sled that is part of a RAID Level 0 group results in the loss of all data in the groups associated with that sled. Follow your backup procedures when removing sleds.

Do not leave empty openings on the front or rear of the array under any circumstances. Empty openings may cause the unit to overheat.

WARNING


All modular components must be replaced by qualified personnel only. Components are electrostatic sensitive. Use a static wriststrap when handling any of the cards inside the array. Use proper grounding methods. Always store spare components in proper ESD containers.

Follow your backup procedures before removing a sled. You may copy drives using the appropriate CLI commands (refer to *Drive Configuration Commands* on page 43).

The following method is the safest way to perform a hot swap of a drive

The general approach: identify the disk drive sled to be replaced, take it offline using the appropriate CLI commands, turn off its power using CLI commands, remove and replace with a new disk drive sled, power up the new sled, and place back on line.

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and enter the CLI.
- 2 The fault LED should blink on the disk drive sled which requires replacement. If it is not,



type `set SledFaultLED [n] on`. The LED of the drive you want [n] lights.

- 3 Determine the Drive Sled Number. The disk drive sled closest to the management card is always the number 1 disk drive sled
- 4 Take the disk drive sled offline by entering the following CLI commands at the control computer **Ready** prompt:

```
Set AtaDiskState (SledNum, 1, OFFLINE)
Set AtaDiskState (SledNum, 2, OFFLINE)
Set DriveSledPower (SledNum, OFF)
```

CAUTION

In a Hot Spare configuration, a drive sled should only be taken offline if there is absolutely no activity on that drive. If there is any activity, the rebuild of the Hot Spare sled may be flawed.

- 5 Wait 30 seconds for the disk drive sled to spin down and complete any remaining I/O activity.
- 6 Unscrew the two screws on the disk drive sled with the appropriate tool.
- 7 Carefully pull the disk drive sled out of chassis using its handle.
- 8 Mark or tag the disk drive sled with the array serial number, the date removed, and its slot number.
- 9 Place the disk drive sled in an appropriate ESD container or bag.
- 10 Install the replacement disk drive sled into the array chassis using proper ESD control steps. The disk drive sled assembly is keyed and can only be inserted one way.
- 11 Tighten the two screws on the disk drive sled with the appropriate tool.

- 12 If you do not have AutoRebuild enabled, using the disk drive sled number identified in step 2, on the control computer **Ready** prompt type:

```
ClearDiskReservedArea (SledNum, 1)
ClearDiskReservedArea (SledNum, 2)
Set AtaDiskState (SledNum, 1, ONLINE)
Set AtaDiskState (SledNum, 2, ONLINE)
ResolveLUNConflicts
```

- 13 The new disk drive sled is available for system use 10 to 15 seconds after the disk drives spin up and communication is reestablished.

Power Supplies



CAUTION

Do not leave empty openings on the front or rear of the array under any circumstances. Empty openings may cause the unit to overheat.



WARNING

Hazardous voltage and stored energy hazard when removing power supplies.

In a system with at least one operational power supply, the other power supply can be successfully removed and replaced without powering the system down and with no loss of array functionality. The green activity LED on the front of the system management card identifies the operational status of each power supply (Green means the power supply is operating correctly).



Note

System command overlap is discontinued across some drives when only one power supply is operational.

Labels on the rear of the array point to the A and B power supplies.

- 1 Identify the power supply to be swapped.
- 2 Turn off the power supply on/off switch.
- 3 Disconnect AC line cord.
- 4 Unscrew the two screws on the power supply with the appropriate tool.
- 5 Pull the power supply out of the chassis using the power supply module handle.
- 6 Install a new power supply in the chassis. The power supply is keyed and can only be inserted one way.
- 7 Tighten the two screws on the power supply with the appropriate tool.
- 8 Connect AC line cord

- 9 Turn on the power supply switch.
- 10 Verify correct operation by observing that the green light on the rear of the power supply is lit and the appropriate power supply light on the system management card on the front of the unit is lit.



Note

*It takes up to 30 seconds for the system to recognize the insertion or removal of a power supply and change the LED on the system management board. The CLI issues messages about the change (refer to *System Monitoring and Reporting* on page 77).*

Blower Assemblies



CAUTION

Do not leave empty openings on the front or rear of the array under any circumstances. Empty openings affect airflow and may cause the unit to overheat and shut down.

The array contains two blower assemblies. The blowers are critical to proper array cooling operation. However, the array can operate with only one functional blower within certain ambient temperatures. The blowers are electronically connected to the power supplies and a power supply will not run without its corresponding blower: if the blower adjacent to power supply A is removed, the 'A' power supply shuts down, turning off the corresponding LED on the system management card.

To replace a blower assembly

- 1 Unscrew the two screws on the blower assembly with the appropriate tool.
- 2 Pull the blower assembly out of chassis.
- 3 Install a new blower assembly in the chassis> The blower assembly is keyed and can only be inserted one way.
- 4 Tighten the two screws on the blower assembly with the appropriate tool.
- 5 Verify correct operation by observing that the green light on the rear of the power supply is lit, and the appropriate power supply green light on the system management card on the front of the unit is also on.



Note

It takes up to 30 seconds for the system to recognize the insertion or removal of a power supply and change the LED on the system

9.2 Optional Hot Spare Sled

To maintain array up time with minimal risk of data loss, individual sleds which fail may be replaced with a spare sled in some configurations.

In most configurations, if a member of a virtual device becomes degraded, you must swap out the faulted sled as defined in *Hot Swap Operating Instructions* on page 87. If you have not enabled *AutoRebuild*, you must also start a manual rebuild.

For four configurations, however, Hot Spare sleds may be designated as replacements for faulted sleds without intervention by you or a host.

Each configuration requires a certain number of Hot Spare sleds. These sleds, once designated as Hot Spares, are not available for other use.

The following configurations support optional Hot Spare sleds:

RAID Level 1: 2 Hot Spare sleds

RAID Level 10: 1 group, 2 Hot Spare sleds

RAID Level 5: 1 group, 1 Hot Spare sled

RAID Level 5: 2 groups, 2 Hot Spare sleds

If a sled becomes degraded and a Hot Spare sled has been designated:

- the Diamond replaces the degraded sled with the Hot Spare sled, simulating a hot swap of a sled with *AutoRebuild enabled*, without intervention
- any sled with a Faulted drive is not used. Faulted sleds maintain their faulted status until they are removed and re-inserted or if the system is restarted.
- Hot Spares are handled as sleds, not as individual drives
- the Hot Spare sled replaces the Faulted sled in the Virtual Device
- a rebuild automatically starts after the hot spare is switched into the Virtual Device, even if *AutoRebuild* is disabled.
- the *DriveInfo* command lists the number of Hot Spare sleds currently in the system

- you can replace the faulted drive or sled later. However, faulted is a non-persistent state. After a power cycle, the faulted sled displays on the DriveInfo screen as having no type, i.e, the **type** field is blank. The sled cannot be accessed but must be removed and replaced. If it is replaced, it becomes part of any open Virtual Disk in this order:

1. if a Virtual Device is missing a sled or you are hot swapping a sled, the new sled becomes part of the existing Virtual Device.

- 2 if a Hot Spare sled is missing, the new sled becomes a Hot Spare sled.

To set up RAID Level 1 with Hot Spare sleds

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV **Advanced CLI Configuration** page (refer to *ATTO ExpressNAV: Browser-based Interface* on page 25).
- 2 Type `set QuickRAID1 2`
- 3 Type `SaveConfiguration Restart`
The Diamond is configured into one RAID Level 1 group with two Hot Spare sleds.

To set up RAID Level 10 with Hot Spare sleds

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV browser interface **Advanced CLI Configuration** page.
- 2 Type `set QuickRAID10 1 2`
- 3 Type `SaveConfiguration Restart`
The Diamond is configured into one RAID Level 10 group with two Hot Spare sleds.

To set up one RAID Level 5 group with one Hot Spare sled

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV browser interface **Advanced CLI Configuration** page.
- 2 Type `set QuickRAID5 1 1`
- 3 Type `SaveConfiguration Restart`
The Diamond is configured into one RAID Level 5 group with one Hot Spare sled.

To set up two RAID Level 5 groups with Hot Spare sleds

The Diamond is configured into two RAID Level 5 groups with two Hot Spare sleds.

- 1 Connect to Diamond Storage Array services (refer to *Accessing the Array* on page 17) and use the CLI or enter ATTO ExpressNAV **Advanced** CLI Configuration page.
- 2 Type **set QuickRAID5 2 2**
- 3 Type **SaveConfiguration Restart**

Exhibit 9.2-1 Configuration of sleds in RAID Level 1: five LUNs with two Hot Spare sleds

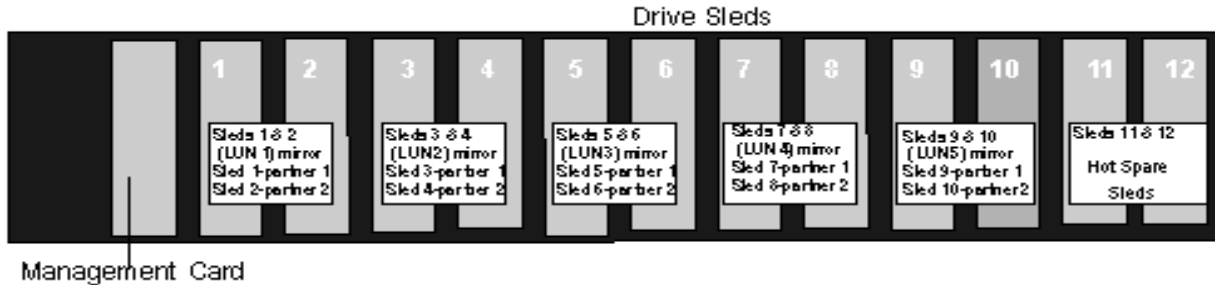


Exhibit 9.2-2 Configuration of sleds in RAID Level 10: five mirrored pairs of sleds in one Virtual Drive and two Hot Spare non-mirrored sleds

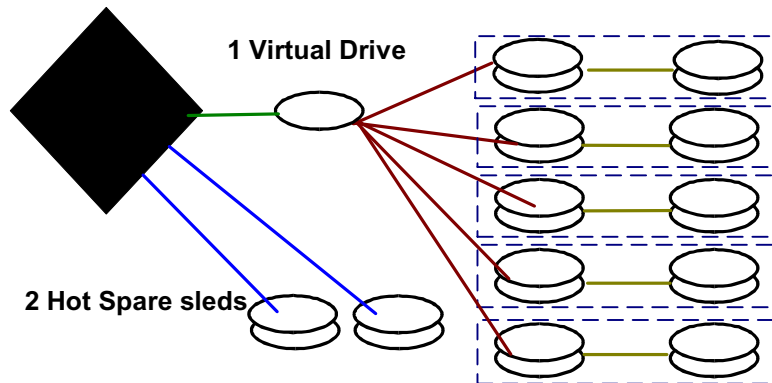


Exhibit 9.2-3 Configuration of sleds in RAID Level 5

Sled	1		2		3		4		5		6		7		8		9		10		11		12	
	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2	1	2
QuickRAID5 1	LUN 1																				Hot Spare			
QuickRAID5 2	LUN 1										LUN 2										Hot Spare		Hot Spare	

Glossary

The following terms are only a few examples of the language used for Fibre Channel, SCSI and storage systems. For more information, visit the websites of these professional organizations: the Fibre Channel Industry Association (www.fibrechannel.com); the Fibre Channel Consortium (www.iol.unh.edu/consortiums/index.html click on Fibre Channel); the Storage Area Networking Industry Association (www.snia.org), or SCSI Trade Association (www.scsita.org)

Fibre Channel technology

Fibre Channel is a serial communication interface designed to transfer large amounts of data among a variety of hardware systems over long distances. It is becoming a key technology for applications that require shared, high-bandwidth access to storage.

Fibre Channel provides a logical point-to-point serial channel for the transfer of data between a buffer at a source device and a buffer at a destination device. It moves buffer contents from one port to another, without regard to the format or meaning of the data, allowing different upper level protocols to run over the Fibre Channel connection.

The Fibre Channel architecture is a hierarchical set of protocol layers. Defined within these layers are rules for signal interfaces, serial encoding and decoding, error control, frame format and communications protocols. Fibre Channel provides a range of implementation possibilities and purposely isolates the transmission medium from the control protocol so that each implementation may use the technology best suited to the application.

Fibre Channel arbitrated loop (FC-AL) technology allows 127 systems and devices to be distributed over a range of up to 10 kilometers running at speeds of 1

gigabit/sec. Multiple FC-AL loops can be connected via a fabric switch topology to extend the range.

SAN technology

A Storage Area Network (SAN) is a shared storage architecture. SANs connect computers and storage devices for online data access. Each connected system can directly access any attached storage device. A SAN maintains greater fault tolerance and load balancing by supporting server clustering and fail-over (the ability for one server to take over for another in the event of a failure).

SCSI protocol

SCSI (pronounced scuzzy) stands for Small Computer System Interface, a processor-independent standard for system-level interface between a computer and SCSI devices such as hard disks, CD-ROM and printers, and to other computers and local area networks.

The protocol has many standards specifying data transfer rates, data width and speeds which allow SCSI devices to exchange data and control signals with a computer's CPU. Each device connected to a SCSI bus must have a different SCSI ID.

Term	Definition
arbitrate	process of selecting one L_Port from a collection of ports which ask for use of the arbitrated loop at the same time
arbitrated loop	a loop topology in which two or more ports are interconnected, but only two ports can communicate at one time
ATA	Advanced Technology Attachment: A disk drive interface standard, formerly called IDE; the ATA specification deals with the power and data signal interface of an integrated disk controller and drive
auto negotiation	hardware senses and automatically responds depending on configuration
CLI	Command Line Interface: used by applications and advanced users to perform diagnostic or configuration tasks; access is available through a computer terminal or terminal emulator via connections to the RS-232 port or the optional Ethernet management services card via the local area network using Telnet
fabric	A Fibre Channel switch or two or more Fibre Channel switches interconnected to physically transmit data between any two N_Ports on a switch or switches

Term	Definition
F_port	A port in the Fibre Channel fabric where a N_port may attach
FC-AL	Fibre Channel Arbitrated Loop: a Fibre Channel network in which up to 126 systems and devices are connected in a loop topology, with each transmitter connecting to the receiver of the device to its logical right; multiple FC-AL loops can be connected via a fabric switch to extend the network
FL-port	A port in the Fibre Channel fabric where a NL_port may attach in an arbitrated loop
firmware	Software stored in read-only memory (ROM) or programmable ROM; easier to change than hardware but more difficult to change than software stored on disk; firmware is responsible for the behavior of the system when it is first switched on
host	a processor which communicates with devices over an interface
hot swapping	Components are removed and replaced while the unit is running, with power to either the component or a device connected to the unit
initiator device	A component which originates a command
JBOD	Just a Bunch of Disks: a storage subsystem of multiple independent disk drives
L_port	a port in an arbitrated loop, either a NL_port or a FL_port
LED	Light-emitting diode, a type of diode that emits light when current passes through it; visible LEDs are used as indicator lights on the Diamond Storage Array
LUN	Logical Unit Number: a SCSI or Fibre Channel identifier of a device. A LUN is associated with a RAID group or Virtual Drive unless you are using partitions. If you have partitions, a LUN is associated with each partition.
N_port	a port attached to a node used with point to point or fabric configurations
NL_port	a port attached to a node in Fibre Channel arbitrated loop or fabric loop
originator	an initiating device; a component which originates a command
partition	a continuous range of addresses in a storage subsystem that appear as a logical unit; used to subdivide and manage a storage subsystem
port	an access point in a device: see N_port, NL_port, etc.
receiver	the ultimate destination of data transmission; a terminal device
RAID	originally Redundant Array of Inexpensive Disks, now Redundant Array of Independent Drives: a storage system spanning multiple disk drives. The following standard RAID specifications will be used here: RAID 0: non-redundant striped array; drives are accessed alternately, as if stacked one on top of the other; no data protection RAID 1: mirrored arrays: information written to one disk is also written to another simultaneously; provides complete data protection RAID 10: striping of mirrored groups RAID 5: striping across multiple disks offering redundancy by including a parity block in every stripe; the parity block is used to reconstruct the data if a single drive in the subsystem fails
subpartition	a convenient term used to describe the allocation nature of a partition; a partition may be split into a number of subpartitions (each a partition in its own right) and a collection of contiguous subpartitions may be merged into a partition
Virtual Drive	the presentation of a collection of disk drives as a single logical disk, also known as a RAID group.
Virtual Drive ID	a unique identifier associated with each Virtual Drive and displayed via CLI

Appendix A ATA Disk Technology

ATA is the dominant disk drive technology. It offers all the cost advantages of a mass produced, consumer-driven technology which is rapidly being driven forward. Today ATA also offers all the performance and reliability features needed to create high performance ATA-based disk storage arrays.

We view ATA as the perfect disk drive technology for “in-the-box” disk storage array solutions.

Aggregated Data Interface Technology (ADXT™) with its built-in Virtual Device Manager (VDM), Data Routing Fabric (DRF) and ATA Interface Engines (AIE) smoothly and efficiently organizes Fibre Channel data streams for storage or retrieval on ATA disk drives. ADXT provides end users with the power and sophisticated data control needed to take moderately priced ATA disk drives combine them in a disk storage array and power them to the performance levels of SCSI or Fibre Channel disk arrays

Unlike other storage arrays which use expensive SCSI or Fibre Channel disk drives to achieve performance, the Diamond Series uses lower cost ATA disk drives combined with an intelligent midplane to create a storage array with exceptional price and performance characteristics.

ATA disk drives are used in the enormous PC marketplace and an estimated 85% of all disk drives sold today contain an ATA interface. Shipments of ATA drives are expected to grow rapidly over the next five years with increasing consumer and business demand. In general, ATA disk drives are the “driver” of all disk storage technology as the costs, volumes and demands of the PC marketplace are the primary focus of the world’s largest disk drive manufacturers.

ATA disk drives today offer capacities comparable to SCSI and FC disk drives – although at a tremendously lower cost. Advances in platter and head technology will push drive capacity along a very steep curve in the next five years. Exhibit A-1 summarizes ATA and other disk interface standards. It clearly indicates that ATA drive technology offers characteristics that ideally suit it for “in the box” usage in high capacity, high performance disk arrays.

Feature	IDE	ATA	SCSI	Fibre Channel
Connectivity • Max Devices • Cable Length	2 46 cm	2 46 cm	16 25 m	126...millions 10 km
Performance • Bandwidth • Bandwidth/drive • Multi-threaded I/O	16 MB/sec. 8 MB/sec. No	100 MB/sec. 33 MB/sec. Yes	160 MB/sec. 10 MB/sec. Yes	200 MB/sec. 1.6 MB/sec. Yes
Max Initiators	1	1	<16	125...millions
Topology	Bussed	Bussed	Bussed	Loop/Fabric
Error Detection	None	Yes	Yes	Yes
Protocol “weight”	very light	light	medium to heavy	heavy
Cost • Interface • Drives	n/a n/a	lowest lowest	medium to high medium to high	highest medium to high
Manageability	none	low to medium	medium	medium+

Appendix B Information command returns

Driveinfo

Sled	Capacity	Errors	TYPE	VDID	Capacity	Errors	TYPE	VDID
1	76292MB	0	JBOD	1	76292MB	0	JBOD	13
2	76292MB	0	JBOD	2	76292MB	0	JBOD	14
3	76292MB	0	JBOD	3	76292MB	0	JBOD	15
4	76292MB	0	JBOD	4	76292MB	0	JBOD	16
5	76292MB	0	JBOD	5	76292MB	0	JBOD	17
6	76292MB	0	JBOD	6	76292MB	0	JBOD	18
7	76292MB	0	JBOD	7	76292MB	0	JBOD	19
8	76292MB	0	JBOD	8	76292MB	0	JBOD	20
9	76292MB	0	JBOD	9	76292MB	0	JBOD	21
10	76292MB	0	JBOD	10	76292MB	0	JBOD	22
11	76292MB	0	JBOD	11	76292MB	0	JBOD	23
12	73307MB	0	JBOD	12	73307MB	0	JBOD	24

Driveinfo 7 2

JBOD VDID 19
 Sled = 7, Drive = 2
 Drive Model: Maxtor 98196H8
 S/N: V80D15GC F/W Rev: ZAH814Y0 ATA Rev: 007e
 Capacity: 76292 MB (09502790 blocks)
 PIO: 03 MDMA: 07 UDMA: 07 QD: 00 CMDSet: 7c69 4309

Operations Blocks
 Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

Driveinfo RAID0 6

Drive 1				Drive 2				
Sled	Capacity	Errors	TYPE	VDID	Capacity	Errors	TYPE	VDID
1	76292MB	0	RAID0	26	76292MB	0	RAID0	26
2	76292MB	0	RAID0	26	76292MB	0	RAID0	26
3	76292MB	0	RAID0	27	76292MB	0	RAID0	27
4	76292MB	0	RAID0	27	76292MB	0	RAID0	27
5	76292MB	0	RAID0	28	76292MB	0	RAID0	28
6	76292MB	0	RAID0	28	76292MB	0	RAID0	28
7	76292MB	0	RAID0	29	76292MB	0	RAID0	29
8	76292MB	0	RAID0	29	76292MB	0	RAID0	29
9	76292MB	0	RAID0	30	76292MB	0	RAID0	30
10	76292MB	0	RAID0	30	76292MB	0	RAID0	30
11	76292MB	0	RAID0	31	76292MB	0	RAID0	31
12	73307MB	0	RAID0	31	73307MB	0	RAID0	31

Driveinfo 3 1

RAID0 VDID 27 State: ONLINE
 RAID0 Member Index 0 of total members 4 disks
 RAID0 Interleave Size 128 Blocks
 Sled = 3, Drive = 1
 Drive Model: Maxtor 98196H8
 S/N: V80DTE5C F/W Rev: ZAH814Y0 ATA Rev: 007e
 Capacity: 76292 MB (09502790 blocks)
 PIO: 03 MDMA: 07 UDMA: 07 QD: 00 CMDSet: 7c69 4309
 Operations Blocks

Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

DrivelInfo RAID1

Drive 1				Drive 2				
Sled	Capacity	Errors	TYPE	VDID	Capacity	Errors	TYPE	VDID
1	76292MB	0	RAID1	50	76292MB	0	RAID1	50
2	76292MB	0	RAID1	50	76292MB	0	RAID1	50
3	76292MB	0	RAID1	51	76292MB	0	RAID1	51
4	76292MB	0	RAID1	51	76292MB	0	RAID1	51
5	76292MB	0	RAID1	52	76292MB	0	RAID1	52
6	76292MB	0	RAID1	52	76292MB	0	RAID1	52
7	76292MB	0	RAID1	53	76292MB	0	RAID1	53
8	76292MB	0	RAID1	53	76292MB	0	RAID1	53
9	76292MB	0	RAID1	54	76292MB	0	RAID1	54
10	76292MB	0	RAID1	54	76292MB	0	RAID1	54
11	76292MB	0	RAID1	55	76292MB	0	RAID1	55
12	73307MB	0	RAID1	55	73307MB	0	RAID1	55

DrivelInfo 5 1

RAID1 VDID 52 State: ONLINE
 RAID1 Member Index 0 of total members 4 disks
 RAID1 Interleave Size 256 Blocks
 Sled = 5, Drive = 1
 Drive Model: Maxtor 98196H8
 S/N: V80EE4TC F/W Rev: ZAH814Y0 ATA Rev: 007e
 Capacity: 76292 MB (09502790 blocks)
 PIO: 03 MDMA: 07 UDMA: 07 QD: 00 CMDSet: 7c69 4309
 Operations Blocks
 Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

DrivelInfo RAID1 with optional Hot Spare sleds activated

Drive 1				Drive 2				
Sled	Capacity	Errors	TYPE	VDID	Capacity	Errors	TYPE	VDID
1	76292MB	0	RAID1	50	76292MB	0	RAID1	50
2	76292MB	0	RAID1	50	76292MB	0	RAID1	50
3	76292MB	0	RAID1	51	76292MB	0	RAID1	51
4	76292MB	0	RAID1	51	76292MB	0	RAID1	51
5	76292MB	0	RAID1	52	76292MB	0	RAID1	52
6	76292MB	0	RAID1	52	76292MB	0	RAID1	52
7	76292MB	0	RAID1	53	76292MB	0	RAID1	53
8	76292MB	0	RAID1	53	76292MB	0	RAID1	53
9	76292MB	0	RAID1	54	76292MB	0	RAID1	54
10	76292MB	0	RAID1	54	76292MB	0	RAID1	54
11	76292MB	0	RAID1	55	76292MB	0	RAID1	55
12	Drive is HOT SPARE				Drive is HOT SPARE			

DrivelInfo RAID10 3

Drive 1				Drive 2				
Sled	Capacity	Errors	TYPE	VDID	Capacity	Errors	TYPE	VDID
1	76292MB	0	RAID10	56	76292MB	0	RAID10	56
2	76292MB	0	RAID10	56	76292MB	0	RAID10	56
3	76292MB	0	RAID10	56	76292MB	0	RAID10	56

```

4 76292MB 0 RAID10 56 76292MB 0 RAID10 56
5 76292MB 0 RAID10 57 76292MB 0 RAID10 57
6 76292MB 0 RAID10 57 76292MB 0 RAID10 57
7 76292MB 0 RAID10 57 76292MB 0 RAID10 57
8 76292MB 0 RAID10 57 76292MB 0 RAID10 57
9 76292MB 0 RAID10 58 76292MB 0 RAID10 58
10 76292MB 0 RAID10 58 76292MB 0 RAID10 58
11 76292MB 0 RAID10 58 76292MB 0 RAID10 58
12 73307MB 0 RAID10 58 73307MB 0 RAID10 58

```

DriveInfo 7 2

```

RAID10 VDID 57 State: ONLINE
RAID10 Member Index 5 of total members 8 disks
RAID10 Interleave Size 128 Blocks
Sled = 7, Drive = 2
Drive Model: Maxtor 98196H8
S/N: V80D15GC F/W Rev: ZAH814Y0 ATA Rev: 007e
Capacity: 76292 MB (09502790 blocks)
PIO: 03 MDMA: 07 UDMA: 07 QD: 00 CMDSet: 7c69 4309
Operations Blocks
Read 00000000 00000000
Write 00000000 00000000
Errors 00000000

```

LUNInfo

LunInfo for a JBOD setup

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	JBOD	ONLINE	76292MB	0	V80EE4YCXA10000
2	JBOD	ONLINE	76292MB	0	V80EE4HCXA10000
3	JBOD	ONLINE	76292MB	0	V80DTESCXA10000
4	JBOD	ONLINE	76292MB	0	V80DTCWCXA10000
5	JBOD	ONLINE	76292MB	0	V80EE4TCXA10000
6	JBOD	ONLINE	76292MB	0	V80CG4RCXA10000
7	JBOD	ONLINE	76292MB	0	V80D44LCXA10000
8	JBOD	ONLINE	76292MB	0	V80DTCXCXA10000
9	JBOD	ONLINE	76292MB	0	V80DTELCXA10000
10	JBOD	ONLINE	76292MB	0	V80ED9GCXA10000
11	JBOD	ONLINE	76292MB	0	V80CFVHCXA10000
12	JBOD	ONLINE	73307MB	0	YSDYSG9HXAI0000
13	JBOD	ONLINE	76292MB	0	V80EE4GCXA10000
14	JBOD	ONLINE	76292MB	0	V80EE5LCXA10000
15	JBOD	ONLINE	76292MB	0	V80DTDTCXA10000
16	JBOD	ONLINE	76292MB	0	V80DTDLCXA10000
17	JBOD	ONLINE	76292MB	0	V80EE4PCXA10000
18	JBOD	ONLINE	76292MB	0	V80CLN4CXA10000
19	JBOD	ONLINE	76292MB	0	V80D15GCXA10000
20	JBOD	ONLINE	76292MB	0	V80DTDRCXA10000
21	JBOD	ONLINE	76292MB	0	V80DTDYCXA10000
22	JBOD	ONLINE	76292MB	0	V80EE5ECXA10000
23	JBOD	ONLINE	76292MB	0	V80CG2GCXA10000
24	JBOD	ONLINE	73307MB	0	YSDYSG9MXAI0000

LunInfo 13

```

ATADisk
SerialNumber V80EE4GC XAI0000
Sled = 1, Drive = 2

```

Drive Model: Maxtor 98196H8
 S/N: V80EE4GCXAI0000 F/W Rev: ZAH814Y0 ATA Rev: 007e
 Capacity: 76292 MB (09502790 blocks)
 PIO: 03 MDMA: 07 UDMA: 07 QD: 00 CMDSet: 7c69 4309

Operations Blocks
 Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

LUNinfo for RAID0 6

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	RAID0	ONLINE	305171MB	0	V80EE4YC0AJ0000
2	RAID0	ONLINE	305171MB	0	V80DTESC0AJ0000
3	RAID0	ONLINE	305171MB	0	V80EE4TC0AJ0000
4	RAID0	ONLINE	305171MB	0	V80D44LC0AJ0000
5	RAID0	ONLINE	305171MB	0	V80DTELC0AJ0000
6	RAID0	ONLINE	293231MB	0	V80CFVHC0AJ0000

LunInfo 4

RAID0
 SerialNumber V80D44LC0AJ0000
 Virtual Drive ID 29, Capacity: 305171 MB (0x00025409e00 blocks)
 Operations Blocks
 Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

LunInfo for a RAID1 configuration

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	RAID1	ONLINE	152585MB	0	V80EE4YC1AK0000
2	RAID1	ONLINE	152585MB	0	V80DTESC1AK0000
3	RAID1	ONLINE	152585MB	0	V80EE4TC1AK0000
4	RAID1	ONLINE	152585MB	0	V80D44LC1AK0000
5	RAID1	ONLINE	152585MB	0	V80DTELC1AK0000
6	RAID1	ONLINE	146615MB	0	V80CFVHC1AK0000

LunInfo 4

RAID1
 SerialNumber V80D44LC1AK0000
 Virtual Drive ID 53, Capacity: 152585 MB (0x00012a04f1e blocks)
 Operations Blocks
 Read 00000000 00000000
 Write 00000000 00000000
 Errors 00000000

LunInfo for a RAID10 3

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	RAID10	ONLINE	305171MB	0	V80EE4YCAAL0000
2	RAID10	ONLINE	305171MB	0	V80EE4TCAAL0000
3	RAID10	ONLINE	293231MB	0	V80DTELCAL0000

LunInfo 2

RAID10

SerialNumber V80EE4TCAAL0000

Virtual Drive ID 57, Capacity: 305171 MB (0x00025409e00 blocks)

Operations Blocks

Read 00000000 00000000

Write 00000000 00000000

Errors 00000000

LunInfo for a RAID5 1 configuration

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	RAID5	ONLINE	1612749MB	0	V80EE4YC5AM0000

LunInfo 1

RAID5

SerialNumber V80EE4YC5AM0000

Virtual Drive ID 59, Capacity: 1612749 MB (0x000c4de6f00 blocks)

Operations Blocks

Read 00000000 00000000

Write 00000000 00000000

Errors 00000000

LunInfo for a RAID5 4 configuration

LUN	TYPE	State	Capacity	Errors	SerialNumber
0	Processor	ONLINE	0MB	0	
1	RAID5	ONLINE	305167MB	0	V80EE4YC5AN0000
2	RAID5	ONLINE	305167MB	0	V80DTCWC5AN0000
3	RAID5	ONLINE	305167MB	0	V80D44LC5AN0000
4	RAID5	ONLINE	293227MB	0	V80ED9GC5AN0000

LunInfo 4

RAID5

SerialNumber V80ED9GC5AN0000

Virtual Drive ID 62, Capacity: 293227 MB (0x00023cb5a00 blocks)

Operations Blocks

Read 00000000 00000000

Write 00000000 00000000

Errors 00000000

Part ID Name VDID LUN Capacity StartVDLBA

Part ID	Name	VDID	LUN	Capacity	StartVDLBA	
0	0000	59	1	7629	1MB	0x0000000000
1	0001	59	5	7629	1MB	0x0009501f00
2	0002	59	6	7629	1MB	0x0012a03e00
3	0003	59	7	7629	2MB	0x001bf05d00
0	0000	60	2	30516	7MB	0x0000000000
0	0000	61	3	30516	7MB	0x0000000000
0	0000	62	4	29322	7MB	0x0000000000

PartitionInfo

PartitionInfo active 59 1

Part ID	Name	VDID	LUN	Capacity	StartVDLBA
0	0000	59	1	76291MB	0x0000000000

TYPE	VDID	#Parts	State	Capacity	Serial Number
RAID5	59	4	ONLINE	305165MB	V80EE4YC5AN
RAID5	60	1	ONLINE	305167MB	V80DTCWC5AN
RAID5	61	1	ONLINE	305167MB	V80D44LC5AN
RAID5	62	1	ONLINE	293227MB	V80ED9GC5AN

VirtualDriveInfo active 59

TYPE	VDID	#Parts	State	Capacity	Serial Number
RAID5	59	4	ONLINE	305165MB	V80EE4YC5AN

Verbose return

ZoneInfo

ZoneInfo abcd
 Zone: abcd disabled
 Host: Remaining
 IID: none
 Port: 0
 Device: 123

ZoneInfo 1234

Zone 1234 disabled
 Host: 20:00:00:10:86:00:21:12
 IID: none
 Port: 0 1
 Device: 3 4 5 6

ZoneInfo 1

Zone: zone1 enabled
 Host: 20:00:00:10:00:00:00
 IID: none
 Port: 1
 Device: 12345

Appendix C Sample Zoning Command Sequences

First time configuration (after download)

QuickRAID or JBOD under Zoning “first time” configuration, an all/all/all configuration in the selected QuickRAID configuration. All zone commands except *ZoneCommit* affect only the planned configuration. *ZoneCommit* sets the Planned Configuration into the Active Configuration.

ZoneClearAll	Clears any previously-defined zones
ZoneCommit	Sets the new zoning configuration.

Simple

2 hosts, Diamond in JBOD configuration (24 LUNs). Hosts can use different operating systems. All zone commands except *ZoneCommit* affect only the planned configuration. *ZoneCommit* sets the Planned Configuration into the Active Configuration.

ZoneClearAll	Clears any previously defined zones from the planned configuration.
ZoneCreate zone_1	Creates a zone named zone_1
ZoneAddHost zone_1 20:00:00:10:86:00:01:34	Adds host WWPN to zone_1
ZoneAddDevice zone_1 1 4	Adds LUN 1 and 4 to zone_1
ZoneAddPort zone_1 0	LUN 1 and LUN 4 will be available through HIC 0 for only the host in zone_1
ZoneCreate zone_2	Creates a zone named zone_2
ZoneAddHost zone_2 20:00:00:60:00:00:12:34	Adds host WWPN to zone_2
ZoneAddDevice zone_2 7	Adds LUN 7 to zone_2
ZoneAddPort zone_2 1	LUN 7 will be available through HIC 1 for only the host in zone_2
ZoneState zone_1 enabled	Enables zone_1.
ZoneState zone_2 enabled	Enables zone_2.
ZoneCommit	Sets the Planned Zoning Configuration into the Active configuration. No other LUNs will be available to any host.

Symmetric Model

1 host, 1 switch, Diamond in JBOD configuration (24 LUNs). All host ports travel through a switch to both Diamond ports. Zoning is not required, since there is no reason to restrict host/HBA/Port access. All zone commands except *ZoneCommit* affect only the planned configuration. *ZoneCommit* sets the Planned Configuration into the Active Configuration.

ZoneClear all	Clears any previously-defined zones from the planned configuration
ZoneCommit	Copies what is in the Planned Zoning Configuration to the Active Configuration: in this case, clears any previously-defined zones.

Asymmetric Model

3 hosts, 1 metadata controller (also a host), RAID 5 with 3 LUNs. All zone commands except *ZoneCommit* affect only the planned configuration. *ZoneCommit* sets the Planned Configuration into the Active Configuration.

ZoneClearAll	Clears any previously defined zones
ZoneCreate zone_1	Creates a zone named zone_1
ZoneAddHost zone_1 20:00:00:10:86:00:01:34	Adds host 1 WWPN to zone_1
ZoneAddHost zone_1 20:00:00:10:10:10:86:01	Adds metadata controller host WWPN to zone_1
ZoneAddDevice zone_1 1 2	Adds LUN 1 and LUN 2 to zone_1
ZoneAddPort zone_1 0	LUN 1 and LUN 2 will be available through HIC 0 for only the hosts in zone_1
ZoneCreate zone_2	Creates a zone named zone_2
ZoneAddHost zone_2 20:00:00:60:32:01:02:12	Adds host 2 WWPN to zone_2
ZoneAddHost zone_2 20:00:00:10:10:10:86:01	Adds metadata controller host WWPN to zone_2
ZoneAddDevice zone_2 1 2	Adds LUN 1 and LUN 2 to zone_2
ZoneAddPort zone_2 0	LUN 1 and LUN 2 will be available through HIC 0 for only the hosts in zone_2
ZoneCreate zone_3	Creates a zone named zone_3
ZoneAddHost zone_3 20:00:00:08:11:01:22:22	Adds host WWPN to zone_3
ZoneAddHost zone_3 20:00:00:10:10:10:86:01	Adds metadata controller host WWPN to zone_3
ZoneAddDevice zone_3 3	Adds LUN 3 to zone_3
ZoneAddPort zone_3 1	LUN 3 will be available through HIC 1 for only the hosts in zone_3
ZoneState zone_1 enabled	Enables Zone_1.
ZoneState zone_2 enabled	Enables Zone_2.
ZoneState zone_3 enabled	Enables Zone_3.
ZoneCommit	Sets the Planned Zoning Configuration into the Active configuration. No other LUNs will be available to any host.

Combined Symmetric/Asymmetric Model

Many hosts, RAID Level 1 with 6 LUNs. All zone commands except *ZoneCommit* affect only the planned configuration. *ZoneCommit* sets the Planned Configuration into the Active Configuration.

ZoneClearAll	Clears any previously defined zones
ZoneCreate zone_1	Creates a zone named zone_1
ZoneAddHost zone_1 20:00:00:10:86:00:01:34	Adds host 1 WWPN to zone_1
ZoneAddHost zone_1 20:00:00:10:10:10:86:01	Adds host 2 WWPN to zone_1
ZoneAddDevice zone_1 1 2 ZoneAddPort zone_1 0	Adds LUN 1 and LUN 2 to zone_1 LUN 1 and LUN 2 will be available through HIC 0 for only the hosts in zone_1
ZoneCreate zone_2 ZoneAddHost remaining	Creates a zone named zone_2 Allows all other hosts (not host 1 or host 2) access to zone_2
ZoneAddDevice zone_2 3 4 5 6 ZoneAddPort zone_2 1	Adds LUNs 3, 4, 5 and 6 to zone_2 LUNs 3, 4, 5 and 6 will be available through HIC 1 for all hosts in zone 2 (excludes host 1 and host 2)
ZoneState zone_1 enabled	Enables Zone_1
ZoneState zone_2 enabled	Enables Zone_2
ZoneCommit	Sets the Planned Zoning Configuration into the Active configuration. No other LUNs will be available to any host.

Appendix D Product Safety

Safe handling of the Diamond Storage Array helps protect its components as well as the people working with them.


- The array is heavy (92 pounds for the floor model, 86 pounds for the rack mount system). Two people are needed to move it safely.
- Mechanical, shock and energy hazards are present through the system if one or more of the modules is removed. There are no operator-serviceable components inside the unit except those indicated in **Hardware Maintenance** on page 85 and **Hot Swap Operating Instructions** on page 87 of this manual.
- This unit is designed to only have inside communication lines connected to it (no lines outside the building proper).
- The array must be provided with adequate air flow. To reduce the risk of damage, the ambient temperature of the inlet air should not exceed 40°C.
- Service personnel must follow proper grounding methods when working with or around the system because of the array's electrostatic sensitive components. Use a static wriststrap when handling any components from the array.



CAUTION

All modular components must be replaced by qualified personnel only.

- This equipment must be connected to a protective earth ground in accordance with the instructions provided in this guide. Improper grounding may result in an electrical shock.



- When the power supply green LED is on, the power source is fully operational, delivering power to the system. The only way to completely disconnect from the system is to

remove the AC power cord from the back of the unit from both power supply inputs. The switch on the power supply is only a standby switch and does not disconnect the AC from the system.

Safety compliances

- UL 1950 3rd Edition
- TUV EN 60950, 1992+A1+A2+A3+A4+A11
- CSA 22.2 No. 950-95
- LEDs and fibre GBIC are Class 1 Laser Rated

EMC specifications

- EN 55022: class A
- EN 50204 RF Immunity 900 MHz Pulse
- EN 61000-4-8 Magnetic Immunity
- EN 61000-4-2: ESD
- EN 61000-4-3: Radiated Immunity
- EN 61000-4-4: EFT
- EN 61000-4-5: Surge
- EN 61000-4-6: Conducted Immunity
- EN 61000-4-11: Voltage Dips and Interruptions

Radio and television interference



Note

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Appendix E Specifications

Environmental and physical

- Humidity: 10-85% non-condensing
- Normal operating temperature at sea level: 5-40° C
- Storage temperature: -25-60° C
- AC input voltage 100-240V rated
- 50/60 Hz

Rack mount dimensions

- Height: 131mm (5.16") 3U
- Width: 447mm (17.60")
- Depth: 585mm (23.03")
- Approximately 39 kg (86 pounds) fully loaded with 24 drives

Floor mount dimensions

- Height: 524mm (20.62") 3U
- Width: 137mm (5.39")
- Depth: 595mm (23.42")
- Approximately 42 kg (92 pounds) fully loaded with 24 drives

Appendix F Part numbers

Each Diamond Storage Array has a master part number to identify it. The floor model base number is on the side of the unit. The rack mount part number is on the top of the unit and near the host interface cards.

Description	All models
System management card (Ethernet)	MGMT-ETHR-001
System management card (RS-232)	MGMT-R232-001
*Cable, 3 meter copper FC (unequalized) (1 Gb FC)	CBL-FCCU-003
*Cable, 10 meter copper FC (unequalized) (1GbFC)	CBL-FCCU-010
*Cable, 10 meter Fibre (1 Gb FC)	CBL-FCFI-010
*Cable, 3 meter optical LCLC (2 Gb FC)	CBL-LCLC-003
*Cable, 10 meter optical LCLC (2 Gb FC)	CBL-LCLC-010
*Cable, 3 meter SCSI VHDC	CBL-V68E-003
*Cable, 30 meter Fibre Channel (1Gb FC)	CBS-FCFI-030
CPCI Filler panel, 3U	CPCI-PANL-003
Blower assembly	DBLA-A001-000
Drive card, dummy	DDSL-0000-000
Dual Disk Drive Sled Assembly, 150 GB	DDSL-150G-000
Dual Disk Drive Sled Assembly, 160 GB	DDSL-160G-M00 DDSL160G-J00
Dual Disk Drive Sled Assembly, 320 GB	DDSL-320G-M00
Host Interface Card (GBIC Fibre Channel)	EPCI-FCCU-001
Host Interface Card (2 gigabit) Class 1 laser rated	EPCI-FCCU-002
Host Interface Card (Ultra 160 SCSI)	DHIC-U160-BLK
Power supply	DPWR-A001-001
*GBIC-DB9 (copper) AMP 788615-1 (1Gb FC)	GBIC-DB90-000
*GBIC-SW SC cnctr Class 1 laser rated (1Gb FC)	GBIC-SWFO-000
*Connector, SFP optical Class 1 laser rated (2 Gb FC)	SFP2-0000-000
*Terminator, 68-pin VHDCI	TERM-V68E-002

**All parts UL recognized*

Appendix G Warranty

Manufacturer limited warranty

Manufacturer warrants to the original purchaser of this product that it will be free from defects in material and workmanship as described in the ATTO Technology website, www.attotech.com. Manufacturer liability shall be limited to replacing or repairing, at its option, any defective product. There will be no charge for parts or labor should Manufacturer determine that this product is defective.

Products which have been subject to abuse, misuse, alteration, neglected, or have been serviced, repaired or installed by unauthorized personnel shall not be covered under this warranty provision. Damage resulting from incorrect connection or an inappropriate application of this product shall not be the responsibility of Manufacturer. Manufacturer's liability is limited to Manufacturer's product(s); damage to other equipment connected to Manufacturer's product(s) will be the customer's responsibility.

This warranty is made in lieu of any other warranty, express or implied. Manufacturer disclaims any implied warranties of merchantability or fitness for a particular purpose. Manufacturer's responsibility to repair or replace a defective product is the sole and exclusive remedy provided to the customer for breach of this warranty. Manufacturer will not be liable for any indirect, special, incidental, or consequential damages irrespective of whether Manufacturer has advance notice of the possibility of such damages. No Manufacturer dealer, agent or employee is authorized to make any modification, extension or addition to this warranty.

Contact ATTO Technology, Inc.

Customer service, sales information and technical support are available by phone Monday through Friday, Eastern Standard Time 8:00 a.m. to 8:00 p.m., or by fax and web site 24-hours a day.

ATTO Technology, Inc.

155 CrossPoint Parkway

Amherst, New York 14068

(716) 691-1999 • voice

(716) 691-9353 • fax

<http://www.attotech.com/diamond>

ATTO Technology can also be reached via e-mail at the following addresses:

Sales Support: sls@attotech.com

Technical Support: techsupp@attotech.com

Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>