

# **HP SureStore E Disk Array FC60**

## **User's Guide**



Edition E1200  
Order No. A5277-90001  
Printed in U.S.A.

---

## Notice

© Hewlett-Packard Company, 1999, 2000. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

---

## Trademark Information

Microsoft, Windows, Windows NT, and Windows 2000 are registered trademarks of the Microsoft Corporation.

---

## Safety Notices



To protect against personal injury and product damage, do not attempt to lift the product without the assistance of another person or lift device.



Components bearing this symbol may be hot to touch.



Components bearing this symbol are fragile. Handle with care.



Components bearing this symbol are susceptible to damage by static electricity. ESD precautions are required.

## Service

Any servicing, adjustment, maintenance, or repair must be performed only by authorized service-trained personnel.

---

## Format Conventions

### Denotes

---

<b>WARNING</b>	<b>A hazard that can cause personal injury</b>
<b>Caution</b>	A hazard that can cause hardware or software damage
<b>Note</b>	Significant concepts or operating instructions
<b>this font</b>	Text to be typed verbatim: all commands, path names, file names, and directory names
<code>this font</code>	Text displayed on the screen

---

## Printing History

1st Edition - September 1999  
2nd Edition - October 1999  
3rd Edition - February 2000  
4th Edition - July 2000  
5th Edition - September 2000  
6th Edition - October 2000  
7th Edition - December 2000

---

## Manual Revision History

### December 2000

Change	Page
Added Figure 43 to clarify operation of the write cache flush thresholds.	134
Added note regarding the impact of LUN binding on performance.	131
Added information on Managing the Universal Transport Mechanism (UTM).	179
Added information on major event logging available with firmware HP08.	188
Added Allocating Space for Disk Array Logs section describing use of environment variable <code>AM60_MAX_LOG_SIZE_MB</code> .	189
Added information on Purging Controller Logs.	192
Added information for RAID 0 support on HP-UX.	45
Changed the required minimum number of disk modules per enclosure from 2 to 4 based on power supply requirements for the disk enclosure.	55

---

## About This Book









This guide is intended for use by system administrators and others involved in operating and managing the HP SureStore E Disk Array FC60. It is organized into the following chapters and section.

<b>Chapter 1, Product Description</b>	Describes the features, controls, and operation of the disk array.
<b>Chapter 2, Array Planning</b>	Guidelines for designing the disk array configuration that best meets your needs.
<b>Chapter 3, Installation</b>	Instruction for moving the disk array.
<b>Chapter 4, Managing the Disk Array on HP-UX</b>	Complete instructions for managing your disk array using the available management software.
<b>Chapter 5, HP-UX Diagnostic Tools</b>	Information on using STM to gather information about disk array status.
<b>Chapter 6, Troubleshooting</b>	Instructions for isolating and solving common problems that may occur during disk array operation.
<b>Chapter 7, Removal and Replacement</b>	Instructions for removing and replacing all customer-replaceable components.
<b>Chapter 8, Reference / Legal / Regulatory</b>	Regulatory, environmental, and other reference information.
<b>GLOSSARY</b>	
<b>Index</b>	

---

## Related Documents and Information

The following items contain information related to the installation and use of the HP SureStore E Disk Array and its management software.

- *HP SureStore E Disk Array FC60 Advanced User's Guide* - this is the expanded version of the book you are reading. Topics that are discussed in more detail in the Advanced User's Guide are clearly identified throughout this book.  
 **Download:** [www.hp.com/support/fc60](http://www.hp.com/support/fc60)
- *HP Storage Manager 60 User's Guide* - this guide describes the features and operation of the disk array management software for Windows NT and Windows 2000. It is included with the A5628A software kit.  
 **Download:** [www.hp.com/support/fc60](http://www.hp.com/support/fc60)
- *HP Storage Manager 60 Introduction Guide* - this guide introduces the disk array management software for Windows NT and Windows 2000. It is included in electronic format on the *HP Storage Manager 60 CD*.  
 **Download:** [www.hp.com/support/fc60](http://www.hp.com/support/fc60)
- *Fibre Channel Mass Storage Adapters Service and User Manual (A3636-90002)* - describes Fibre Channel Mass Storage/9000. It describes installation of the Fibre Channel I/O adaptors into K-, D-, T-, and V-class systems.  
 **Download:** [www.hp.com/essd/efc/A3636A\\_documentation.html](http://www.hp.com/essd/efc/A3636A_documentation.html)
- *Using EMS HA Monitors (B5735-90001)* - contains information about the EMS environment used for hardware monitoring.  
 **Download:** <http://www.docs.hp.com/hpux/ha/>
- *EMS Hardware Monitors User's Guide* - describes how to use the EMS Hardware Monitors to protect your system from undetected failures.  
 **Download:** <http://www.docs.hp.com/hpux/systems/>
- *Diagnostic/IPR Media User's Guide (B6191-90015)* - provides information on using STM, and enabling the EMS Hardware Event Monitors.  
 **Download:** <http://www.docs.hp.com/hpux/systems/>
- *Managing MC/ServiceGuard (B3939-90024)* - provides information on creating package dependencies for hardware resources.  
 **Download:** <http://www.docs.hp.com/hpux/ha/>

---

## 1 Product Description

Product Description . . . . .	16
Operating System Support . . . . .	18
Management Tools . . . . .	18
Features . . . . .	18
High Availability . . . . .	19
Scalable Storage Capacity . . . . .	19
LED Status Monitoring . . . . .	20
EMS Hardware Event Monitoring (HP-UX Only) . . . . .	20
Disk Enclosure Components . . . . .	21
Operation Features . . . . .	23
Disk Enclosure SC10 Modules . . . . .	25
Array Controller Enclosure Components . . . . .	32
Front Cover . . . . .	35
Controller Modules . . . . .	36
Controller Fan Modules . . . . .	38
Power Supply Modules . . . . .	40
Power Supply Fan Module . . . . .	41
Battery Backup Unit . . . . .	43
Disk Array High Availability Features . . . . .	45
RAID Level Comparisons . . . . .	45
Capacity Management Features . . . . .	50

## 2 Array Planning

Overview . . . . .	52
Array Design Considerations . . . . .	53
Array Hardware Configuration . . . . .	53
RAID, LUNs, and Global Hot Spares . . . . .	54
High Availability . . . . .	54
Performance . . . . .	55

---

Storage Capacity .....	57
Expanding Storage Capacity .....	57
Recommended Disk Array Configurations .....	59
Configuration Considerations .....	59
One Disk Enclosure Configuration .....	60
Two Disk Enclosure Configurations .....	62
Three Disk Enclosure Configurations .....	66
Four Disk Enclosure Configurations .....	70
Five Disk Enclosure Configurations .....	74
Six Disk Enclosure Configurations .....	78
Total Disk Array Capacity .....	82
For high-availability, one disk per SCSI channel is used as a global hot spare...	83

### 3 Installation

Overview .....	86
Host System Requirements .....	87
HP-UX .....	87
Windows NT and Windows 2000 .....	88
Configuration Switches .....	89
Disk Enclosure (Tray) ID Switch .....	89
Disk Enclosure DIP Switch .....	89
Fibre Channel Host ID Address Setting .....	93
Attaching Power Cords .....	97
Attaching SCSI Cables and	
Configuring the Disk Enclosure Switches .....	98
Full-Bus Cabling and Switch Configuration .....	99
Split-Bus Switch and Cabling Configurations .....	102
Connecting the Fibre Channel Cables .....	107
Applying Power to the Disk Array .....	109
Solving Common Installation Problems .....	117



---

## 4 Managing the Disk Array on HP-UX

Tools for Managing the Disk Array FC60 . . . . .	120
System Administration Manager (SAM) . . . . .	120
Array Manager 60 . . . . .	120
STM . . . . .	120
Installing the Array Manager 60 Software . . . . .	122
Running Array Manager 60 . . . . .	122
Managing Disk Array Capacity . . . . .	123
Configuring LUNs . . . . .	123
Selecting Disks for a LUN . . . . .	124
Assigning LUN Ownership . . . . .	128
Selecting a RAID Level . . . . .	128
Global Hot Spares . . . . .	129
Setting Stripe Segment Size . . . . .	130
Evaluating Performance Impact . . . . .	131
Adding Capacity to the Disk Array . . . . .	135
Adding More Disk Modules . . . . .	135
Adding Additional Disk Enclosures . . . . .	137
Replacing Disk Modules with Higher Capacity Modules . . . . .	137
Upgrading Controller Cache to 512 Mbytes . . . . .	139
Managing the Disk Array Using SAM . . . . .	141
Checking Disk Array Status . . . . .	142
Assigning an Alias to the Disk Array . . . . .	145
Locating Disk Modules . . . . .	146
Binding a LUN . . . . .	148
Unbinding a LUN . . . . .	152
Replacing a LUN . . . . .	152
Adding a Global Hot Spare . . . . .	154
Removing a Global Hot Spare . . . . .	156
Managing the Disk Array Using Array Manager 60 . . . . .	157
Command Syntax Conventions . . . . .	160

---

Array Manager 60 man pages . . . . .	160
Quick Help . . . . .	160
Selecting a Disk Array and Its Components . . . . .	161
Preparing to Manage the Disk Array . . . . .	162
Checking Disk Array Status . . . . .	163
Managing LUNs . . . . .	170
Managing Global Hot Spares . . . . .	176
Managing Disk Array Configuration . . . . .	178
Managing the Universal Transport Mechanism (UTM) . . . . .	179
Managing Cache Parameters . . . . .	180
Performing Disk Array Maintenance . . . . .	185
Managing Disk Array Logs . . . . .	188
Managing the Disk Array Using STM . . . . .	194
Checking Disk Array Status Information . . . . .	194
Binding a LUN . . . . .	194
Unbinding a LUN . . . . .	195
Adding a Global Hot Spare . . . . .	195
Removing a Global Hot Spare . . . . .	195
Locating Disk Modules . . . . .	196

## 5 HP-UX Diagnostic Tools

Overview . . . . .	198
Support Tools Manager . . . . .	199
STM User Interfaces . . . . .	199
STM Tools . . . . .	203
Using the STM Information Tool . . . . .	204
Interpreting the Information Tool Information Log . . . . .	206
Interpreting the Information Tool Activity Log . . . . .	206
Using the STM Expert Tool . . . . .	207

## 6 Troubleshooting

Introduction . . . . .	212
About Field Replaceable Units (FRUs) . . . . .	213
HP-UX Troubleshooting Tools . . . . .	214
Windows NT and Windows 2000 Troubleshooting Tools . . . . .	214
EMS Hardware Event Monitoring (HP-UX Only) . . . . .	214
Disk Array Installation/Troubleshooting Checklist . . . . .	217
Power-Up Troubleshooting . . . . .	218
Controller Enclosure Troubleshooting . . . . .	219
Introduction . . . . .	219
Controller Enclosure LEDs . . . . .	220
SureStore E Disk System SC10 Troubleshooting . . . . .	222
Disk Enclosure LEDs . . . . .	222
Losing LUN 0 . . . . .	222
Interpreting Component Status Values (HP-UX Only) . . . . .	225
Isolating Causes . . . . .	226

## 7 Removal and Replacement

Overview . . . . .	230
Disk Enclosure Modules . . . . .	232
Disk Module or Filler Module . . . . .	232
Disk Enclosure Fan Module . . . . .	238
Disk Enclosure Power Supply Module . . . . .	240
Controller Enclosure Modules . . . . .	242
Front Cover Removal/Replacement . . . . .	243
Controller Fan Module . . . . .	244
Battery Backup Unit (BBU) Removal/Replacement . . . . .	246
Power Supply Fan Module Removal/Replacement . . . . .	249
Power Supply Module Removal/Replacement . . . . .	251
SCSI Cables . . . . .	253

---

## 8 Reference / Legal / Regulatory

System Requirements	256
Host Systems	256
Supported Operating Systems	256
Fibre Channel Host Adapters	257
Models and Options	258
A5277A/AZ Controller Enclosure Models and Options	258
A5294A/AZ Disk Enclosure SC10 Models and Options	260
Disk Array FC60 Upgrade and Add-On Products	262
PDU/PDRU Products	263
Replaceable Parts	264
A5277A/AZ Controller Enclosure Specifications	266
Dimensions:	266
Weight:	266
AC Power:	267
Heat Output:	267
Environmental Specifications	268
Acoustics	269
A5294A/AZ Disk Enclosure Specifications	270
Dimensions:	270
Weight:	270
AC Power:	271
DC Power Output:	271
Heat Output:	271
Environmental Specifications	272
Acoustics	273
Warranty and License Information	274
Hewlett-Packard Hardware Limited Warranty	274
Software Product Limited Warranty	275
Limitation of Warranty	275
Hewlett-Packard Software License Terms	277

---

Regulatory Compliance.....	280
Safety Certifications:.....	280
EMC Compliance.....	280
FCC Statements (USA Only).....	281
IEC Statement (Worldwide).....	281
CSA Statement (For Canada Only).....	281
VCCI Statement (Japan).....	282
Harmonics Conformance (Japan).....	282
Class A Warning Statement (Taiwan).....	282
Spécification ATI Classe A (France Seulement).....	283
Product Noise Declaration (For Germany Only).....	283
Geräuschemission (For Germany Only).....	284
Declaration of Conformity.....	285

---

---

**1**

**PRODUCT DESCRIPTION**

---

<b>Product Description</b> .....	<b>16</b>
<b>Disk Enclosure Components</b> .....	<b>21</b>
<b>Array Controller Enclosure Components</b> .....	<b>32</b>
<b>Disk Array High Availability Features</b> .....	<b>45</b>
<b>Capacity Management Features</b> .....	<b>50</b>

---

---

## Product Description


The HP SureStore E Disk Array FC60 (Disk Array FC60) is a disk storage system that features high data availability, high performance, and storage scalability. To provide high availability, the Disk Array FC60 uses redundant, hot swappable modules, which can be replaced without disrupting disk array operation should they fail.

The Disk Array FC60 consists of two primary components: an FC60 controller enclosure, and from one to six HP SureStore E Disk System SC10 enclosures (referred to throughout this document as simply disk enclosures). The controller enclosure is responsible for providing overall control of the Disk Array FC60 by managing the communication between the host and the disk enclosures. Host communication is done through dual Fibre Channel arbitrated loops (when dual controller modules are installed). By using Fibre Channel, the Disk Array FC60 achieves a high data rate throughput. High data throughput is maintained to the disks by using up to six Ultra2 SCSI channels to the disk enclosures (one channel for each disk enclosure).

In addition to increased performance, the use of multiple disk enclosures provides scalability — simplifying the process of adding storage capacity as needed. Up to six disk enclosures can be added incrementally as storage demands increase. Each disk enclosure holds up to ten disk modules in capacities of 9.1 Gbyte, 8.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte. A fully loaded system comprising six disk enclosures, each populated with ten 73.4-Gbyte disk modules, achieves a capacity of over 3 Tbytes.

The Disk Array FC60 enclosures are designed for installation in HP's original 19-inch cabinets, which include the C2785A (1.1m), C2786A (1.6m), C2787A (2m), A1896A (1.1 m), and A1897A (1.6 m) and the HP Rack System/E Racks which includes the A490xA and A150xA Rack System/E cabinets. The Disk Array FC60 is also supported in the Rittal 9000 Series racks.





**Figure 1** HP SureStore E Disk Array FC60 (Controller with Six Disk Enclosures)

---

## Operating System Support

The Disk Array FC60 is currently supported on the following operating systems:

- HP-UX 11.0, 11.11, and 10.20
- Windows<sup>®</sup> NT 4.0
- Windows<sup>®</sup> 2000

---

**Note** Some disk array features are specific to each operating system. These features are clearly identified throughout this book.

---

## Management Tools

### HP-UX Tools

The following tools are available for managing the disk array on HP-UX. These tools are included with the disk array.

- Array Manager 60 command line utilities
- SAM
- Support Tools Manager (STM)

### Windows<sup>®</sup> NT and Windows<sup>®</sup> 2000

The following tools are used to manage the Disk Array FC60 on Windows NT and Windows 2000. This tool is not included with the disk array, but must be ordered separately as product A5628A.

- *HP Storage Manager 60 (A5628A)*

## Features

The Disk Array FC60 offers the following features

- High availability
- Scalable storage capacity
- LED status monitoring

- 
- RAID levels 0, 1, 0/1, 3, and 5 (RAID level 3 supported on Windows NT and Windows 2000 only)
  - EMS hardware monitoring (HP-UX only)

## High Availability

High availability is a general term that describes hardware and software systems that are designed to minimize system downtime — planned or unplanned. The Disk Array FC60 qualifies as high-availability hardware, achieving 99.99% availability.

The following features enable high availability:

- Hot-swappable, high-capacity, high-speed disk modules
- Dual Fibre Channel arbitrated loops (FC-AL) connection to the host
- Redundant, hot-swappable, fans and power supplies
- Support for RAID 1, 0/1, and 5
- Remote monitoring and diagnostics
- EMS Hardware event monitoring and real-time error reporting (HP-UX only)

---

**Note** The Disk Array FC60 is designed to operate with either one or two controller modules; however, for data integrity and high availability it is highly recommended that dual controller modules be installed.

---

## Scalable Storage Capacity

The Disk Array FC60 is designed to provide maximum scalability, simplifying the process of adding storage capacity as required. Storage capacity can be added in three ways:

- By adding additional disk modules to a disk enclosure
- By adding additional disk enclosures to the array
- By replacing existing disk modules with higher capacity modules

The controller enclosure supports up to six disk enclosures. Each disk enclosure holds up to ten disk modules in capacities of 9.1 Gbyte, 18.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte. The minimum configuration for the array is one disk enclosure with two 9.1-Gbyte disk modules. The maximum configuration is six disk enclosures with ten 73.4-Gbyte disk

---

modules. This provides a storage capacity range from 36 Gbytes to over 3 Tbytes of usable storage.

## **LED Status Monitoring**

Both the controller enclosure and disk enclosure monitor the status of their internal components and operations. At least one LED is provided for each swappable module. If an error is detected in any module, the error is displayed on the appropriate module's LED. This allows failed modules to be quickly identified and replaced.

## **EMS Hardware Event Monitoring (HP-UX Only)**

The Disk Array FC60 is fully supported by Hewlett-Packard's EMS Hardware Monitors, which allow you to monitor all aspects of product operation and be alerted immediately if any failure or other unusual event occurs. Hardware monitoring is available at no cost to users running HP-UX 11.0 or 10.20.

Hardware monitoring provides a high level of protection against system hardware failure. It provides an important tool for implementing a high-availability strategy for your system. By using EMS Hardware Monitors, you can virtually eliminate undetected hardware failures that could interrupt system operation or cause data loss.

The EMS Hardware Monitor software with the Disk Array FC60 monitor is distributed on the *HP-UX Support Plus CD-ROM* release 9912 and later. Complete information on installing and using hardware event monitoring is contained in the EMS Hardware Monitors User's Guide (B6191-90011). A copy of this book can be accessed from the Systems Hardware, Diagnostics, & Monitoring page of Hewlett-Packard's on-line documentation web site at <http://www.docs.hp.com/hpux/systems/>

The minimum supported version of the Disk Array FC60 hardware monitor (fc60mon) is A.01.04. To verify the version of the monitor installed, type:


```
what /usr/sbin/stm/uut/bin/tools/monitor/fc60mon
```

---

## Disk Enclosure Components

The SureStore E Disk System SC10, or disk enclosure, is a high availability Ultra2 SCSI storage product. It provides an LVD SCSI connection to the controller enclosure and ten slots on a single-ended backplane for high-speed, high-capacity LVD SCSI disks. Six disk enclosures fully populated with 9.1 Gbyte disks provide 0.54 Tbytes of storage in a 2-meter System/E rack. When fully populated with 73.4 Gbyte disks, the array provides over 3 Tbytes of storage. These values represent maximum storage; usable storage space will vary depending on the RAID level used.

The disk enclosure consists of modular, redundant components that are easy to upgrade and maintain. See Figure 2. Disks, fans, power supplies, and Bus Control Cards (BCCs) are replaceable parts that plug into individual slots in the front and back of the disk enclosure. Redundant fans, power supplies, and disk modules can be removed and replaced without interrupting storage operations. In addition, a single disk within a LUN can be replaced while the system is on.



**Figure 2** Disk Enclosure Components, Exploded View

---

## Operation Features

The disk enclosure is designed to be installed in a standard 19-inch rack and occupies 3.5 EIA units (high). Disk drives mount in the front of the enclosure. Also located in the front of the enclosure are a power switch and status LEDs. A lockable front door shields RFI and restricts access to the disk drives and power button (Figure 3 on page 24).

BCCs are installed in the back of the enclosure along with redundant power supplies and fans.

## Status Indicators

LEDs on the front and back of the disk enclosure enable you to quickly identify and replace failed components, thereby preventing or minimizing downtime. "Troubleshooting" on page 211 provides more detailed information about the operation of these LEDs.

Two system LEDs on the front, top right corner of the disk enclosure (A in Figure 3) indicate the status of the disk enclosure. The left LED indicates when power is on or off and the right LED identifies if a fault has occurred. Additional pairs of LEDs above each disk slot (D in Figure 3), indicate disk activity and a fault condition. The left LED (green) indicates disk I/O activity and the right LED goes on if the disk module has experienced a fault. The disk fault LEDs are also used by the management tools to identify a specific disk module by flashing its fault LED.


On the back of the disk enclosure, the following LEDs (K) indicate the status of replaceable components and Fibre Channel link:

- Power supply status and fault LEDs
- Fan status and fault LEDs

Bus Controller Card LEDs:

- BCC Fault LED
- Term Power LED (monitors power on the SCSI bus)
- Full Bus Mode LED
- LVD Mode LED
- Bus Free Status LED

For detailed information on LED operation, refer to "Troubleshooting" on page 211



**Figure 3** Disk Enclosure Front and Back View

**Power Switch**

The power switch (B in Figure 3) interrupts power from the power supplies to the disk enclosure components. Power to the power supplies is controlled by the power cords and the AC source.



---

## Disk Enclosure SC10 Modules

The disk enclosure hot-swappable modules include the following:

- Disks and fillers
- Fans
- Power supplies

### Disks and Fillers

Hot-swappable disk modules make it easy to add or replace disks. Fillers are required in all unused slots to maintain proper airflow within the enclosure.

Figure 4 illustrates the 3.5-inch disks in a metal carrier. The open carrier design allows ten half height (1.6 inch) disks to fit the 19-inch width of a standard rack and meet cooling needs.

---

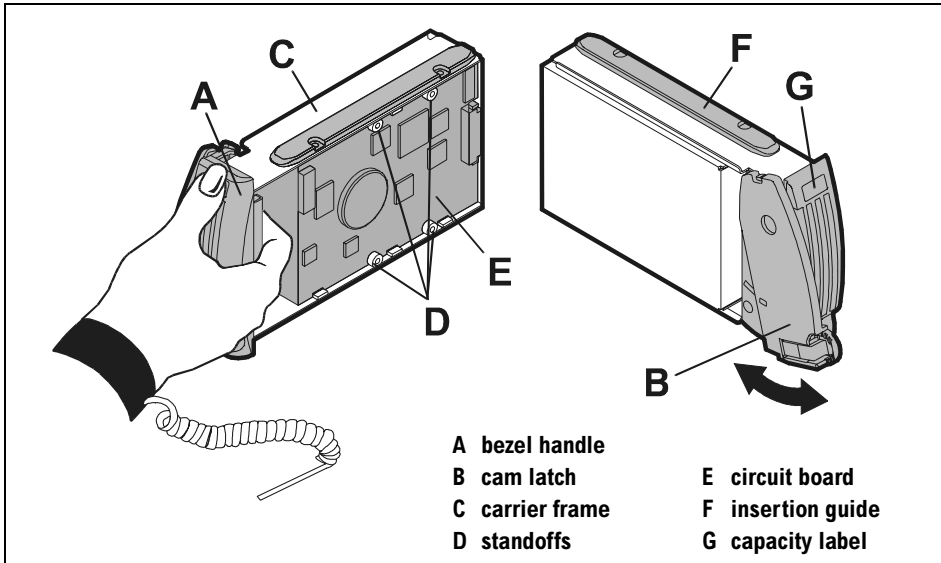
**WARNING**      **Touching exposed circuits on the disk module can damage the disk drive inside. To avoid damage, always handle disks carefully and use ESD precautions.**

---

The following plastic parts of the disk are safe to touch:

- Bezel-handle (A in Figure 4)
- Cam latch (B)
- Insertion guide (F)

Metal standoffs (D) protect exposed circuits against damage when the disk is laid circuit-side down on a flat surface.



**Figure 4** Disk Module

Disks fit snugly in their slots. The cam latch (B in Figure 4) is used to seat and unseat the connectors on the backplane.

A label (G) on the disk provides the following information:

- Disk mechanism height: 1.6 inch (half height) or 1 inch (low profile)
- Rotational speed: 10K RPM and 15K RPM (18 Gbyte only)
- Capacity: 9.1 Gbyte, 18.2 Gbyte, 36.4 Gbyte, or 73.4 Gbyte

A large zero on the capacity label distinguishes a filler from a disk. Fillers are required in all unused slots to maintain proper airflow within the enclosure.

---


**CAUTION** *Fillers must be installed in unused slots to maintain proper cooling within the disk enclosure.*

---

## BCCs

Two Backplane Controller Cards, BCCs, control the disks on one or two buses according to the setting of the Full Bus switch. When the Full Bus switch is set to on, BCC A, in the top slot, accesses the disks in all ten slots. When the Full Bus switch is off, BCC A accesses disks in the even-numbered slots and BCC B accesses disks in the odd-numbered slots.

**Note** In full bus mode, all ten disks can be accessed through either BCC. However, internally each BCC still manages five disks. This means that if the BCC that is not connected to the SCSI channel fails, access to its five disks will be lost. Failure of the BCC that is connected to the SCSI channel will render all ten disks inaccessible.



**Figure 5** BCC

---

Each BCC provides two LVD SCSI ports (B in Figure 5) for connection to the controller enclosure.

The EEPROM on each BCC stores 2 Kbytes of configuration information and user-defined data, including the manufacturer serial number, World Wide Name, and product number.


The following are additional features of the BCC:

- LEDs (C in Figure 5) show the status of the BCC and the bus.
- A rotary switch (D) used to set the enclosure (tray) ID which is used by internal controller operations, and also by the management tools to identify each enclosure.
- DIP switches set disk enclosure options. The only option used by the Disk Array FC60 is the full-bus/split-bus mode.
- Screws (F) prevent the card from being unintentionally disconnected.
- Cam levers (G) assist in installing and removing the BCC from the enclosure, ensuring a tight connection with the backplane.

BCC functions include drive addressing, fault detection, and environmental services.

## Fans

Redundant, hot-swappable fans provide cooling for all enclosure components. Each fan has two internal high-speed blowers (A in Figure 6), an LED (B), a pull tab (C), and two locking screws (D).




**Figure 6** Fan

Internal circuitry senses blower motion and triggers a fault when the speed of either blower falls below a critical level. If a fan failure occurs, the amber fault LED will go on. An alert should also be generated by EMS Hardware Monitoring when a fan failure occurs.

---

## Power Supplies

Redundant, hot-swappable 450-watt power supplies convert wide-ranging AC voltage from an external main to stable DC output and deliver it to the backplane. Each power supply has two internal blowers, an AC receptacle (A in Figure 7), a cam handle (B) with locking screw, and an LED (C). Internal control prevents the rear DC connector from becoming energized when the power supply is removed from the disk enclosure.



**Figure 7** Power Supply

---

**Note** Although it is possible to operate the disk enclosure on one power supply, it is not recommended. Using only one supply creates a single point of failure. If the power supply fails, the entire enclosure is inaccessible. To maintain high availability, both power supplies should be used at all times, and a failed supply should be replaced as soon as possible.

---

---

Power supplies share the load reciprocally; that is, each supply automatically increases its output to compensate for reduced output from the other. If one power supply fails, the other delivers the entire load.

Internal circuitry triggers a fault when a power supply fan or other power supply part fails. If a power supply failure occurs, the amber fault LED will go on. An alert should also be generated by EMS Hardware Monitoring when a power supply failure occurs.

---


## Array Controller Enclosure Components

The array controller enclosure, like the disk enclosure, consists of several modules that can be easily replaced, plus several additional internal assemblies. See Figure 8. Together, these removable modules and internal assemblies make up the field replaceable units (FRUs). Many modules can be removed and replaced without disrupting disk array operation.

The following modules are contained in the controller enclosure:


- Controller modules
- Controller fan module
- Power supply modules
- Power supply fan module
- Battery backup unit






**Figure 8** Controller Enclosure Exploded View

During operation, controller enclosure status is indicated by five LEDs on the front left of the controller enclosure. Faults detected by the controller module cause the corresponding controller enclosure fault LED to go on. Additional LEDs on the individual components identify the failed component. See "Troubleshooting" on page 211 for detailed information on LED operation.



**Figure 9** Controller Enclosure Front View



**Figure 10** Controller Enclosure Rear View


## Front Cover

The controller enclosure has a removable front cover which contains slots for viewing the main operating LEDs. The cover also contains grills that aid air circulation. The controller modules, controller fan, and battery backup unit are located behind this cover. This cover must be removed to gain access to these modules, and also, to observe the controller status and BBU LEDs.

---

## Controller Modules

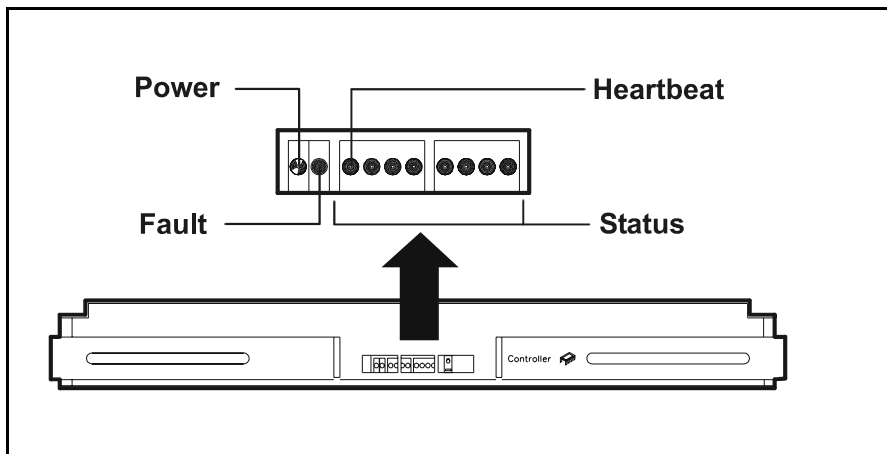
The controller enclosure contains one or two controller modules. See Figure 11. These modules provide the main data and status processing for the Disk Array FC60. The controller modules slide into two controller slots (A and B) and plug directly into the backplane. Two handles lock the modules in place. Each controller slot has a controller letter that identifies the physical location of the controller in the chassis: controller slot A or controller slot B (also known as BD1 and BD2, respectively, as referenced on the back of the controller enclosure).



**Figure 11** Controller Modules

Each controller module has ten LEDs. See Figure 12. One LED identifies the controller module's power status. A second LED indicates when a fault is detected. The remaining eight LEDs provide detailed fault condition status. The most significant LED, the heartbeat, flashes approximately every two seconds beginning 15 seconds after power-on. "Troubleshooting" on page 211 contains additional information on controller LED operation.

The controller module connects to the host via Fibre Channel, and to the disk enclosures via LVD SCSI. Each controller must have a unique host fibre ID number assigned using the ID switches on the back of the controller modules. See "Installation" on page 85 for more information on setting host IDs.



**Figure 12** Controller Module LEDs


---

## Controller Memory Modules

Each controller module contains SIMM and DIMM memory modules. Two 16-Mbyte SIMMs (32 Mbytes total) store controller program and other data required for operation. The standard controller module includes 256-Mbytes of cache DIMM, which is upgradeable to 512 Mbytes. The cache may be configured as either two 128-Mbyte DIMMs, or a single 256-Mbyte DIMM. Cache memory serves as temporary data storage during read and write operations, improving I/O performance. When cache memory contains write data, the Fast Write Cache LED, on the front of the controller enclosure is on. See Figure 13.

## Controller Fan Modules

The controller fan module is a single removable unit containing dual cooling fans and temperature monitoring logic. See Figure 13. It includes five LEDs that indicate overall system status and controller fan status. The fans provide cooling by pulling air in through ventilation holes, moving it across the controller cards, and exhausting it out the ventilation holes in the fan assembly. The dual fans provide a redundant cooling system to both controller modules. If one fan fails, the other continues to operate and provide sufficient air circulation to prevent the controllers from overheating until the fan module is replaced. The fan module plugs into a slot on the front of the controller enclosure, and has a handle and captive screw for easy service.




**Figure 13** Controller Fan Module

---

## Power Supply Modules

Two separate power supplies provide electrical power to the internal components by converting incoming AC voltage to DC voltage. Both power supplies are housed in removable power supply modules that slide into two slots in the back of the controller and plug directly into the power interface board. See Figure 14.



**Figure 14** Power Supply Modules

Each power supply uses a separate power cord. These two power cords are special ferrite bead cords (part no. 5064-2482) required for FCC compliance. Both power cords can plug into a common power source or each cord can plug into a separate circuit (to provide power source redundancy).



---

Each power supply is equipped with a power switch to disconnect power to the supply. Turning off both switches turns off power to the controller. This should not be performed unless I/O activity to the disk array has been stopped, and the write cache has been flushed as indicated by the Fast Write Cache LED being off.

---

**CAUTION** *The controller power switches should not be turned off unless all I/O activity to the disk array has been suspended from the host. Also, power should not be turned off if the Fast Write Cache LED is on, indicating that there is data in the write cache. Wait until the Fast Write Cache LED goes off before shutting off power to the disk array.*


---

Each power supply is equipped with a power-on LED indicator. If the LED is on (green) the supply is providing dc power to the controller. If the LED is off, there is a malfunction or the power has been interrupted. The system Power Fault LED on the front of the controller enclosure works in conjunction with the Power Supply LEDs. If both power supplies are on, the system Power Fault LED will be off. If either power supply is off or in a fault state, the system Power Fault LED goes on. When both power supplies are off or not providing power to the enclosure, the system power LED on the front of the controller enclosure will be off.

## Power Supply Fan Module

Like the controller fan, the power supply fan module (Figure 15) is a single removable unit that contains dual cooling fans. Dual fans provide a redundant cooling system to both power supply modules. If one fan fails, the other will continue to operate. A single fan will provide sufficient air circulation to prevent the power supplies from overheating until the entire power supply fan module can be replaced. The power supply fan module plugs directly into a slot on the back of the controller enclosure, between the power supplies. It has a locking lever that allows it to be unlatched and removed.


The power supply fan can be hot swapped, provided the exchange is performed within 15 minutes. This time limit applies only to the total time during which the fan is out of the enclosure, beginning when you remove the failed unit and ending when you re-seat the new one. This does not include the time it takes you to perform this entire procedure (including checking LEDs).



**Figure 15** Power Supply Fan Module

## Battery Backup Unit

The controller enclosure contains one removable battery backup unit (BBU) that houses two rechargeable internal batteries (A and B) and a battery charger board. The BBU plugs into the front of the controller enclosure where it provides backup power to the controller's cache memory during a power outage. The BBU will supply power to the controllers for up to five days (120 hrs). All data stored in memory will be preserved as long as the BBU supplies power. When power to the disk array is restored, the cache data will be written to disk.



**Figure 16** Battery Backup Unit

**CAUTION** *During a power outage, do not remove the controller or the BBU. Removing either of these modules can compromise data integrity.*

---

The BBU contains four LEDs that identify the condition of the battery. Internally, the BBU consists of two batteries or banks, identified as bank “A” and bank “B.” During normal operation both of the Full Charge LEDs (Full Charge-A and Full Charge-B) are on and the two amber Fault LEDs are off. If one or both of the Fault LEDs are on, refer to "Troubleshooting" on page 211 for information on solving the problem. The Full Charge LEDs flash while the BBU is charging. It can take up to seven hours to fully charge a new BBU.

### **Battery Operation and Replacement**

Replace the BBU every two years or whenever it fails to hold a charge, as indicated by the BBU Fault LEDs. The service label on the BBU provides a line for recording the date on which the BBU was serviced. Check this label to determine when to replace the BBU. When a BBU is replaced, it may require up to seven hours to fully charge. The Full Charge LEDs flash while the BBU is charging, and remain on when the BBU is fully charged.

If you replace the BBU and still experience battery-related problems (such as a loss of battery power to the controllers or batteries not charging properly), the controller enclosure may have some other internal component failure. In this case contact your HP service engineer.

### **Battery Operation for No Data Loss**

The BBU protects the write cache (data which has not been written to disk) for at least 120 hours (five days) in case of a power failure. When power to the disk array is restored, data in the cache is written to the disks and no data loss occurs. However, if the system is to be powered off for more than 120 hours, it is imperative that a proper shutdown procedure be executed or data may be lost. The following are recommendations:

- Battery status must always be checked and replaced when a failure is indicated.
- Never remove the BBU without first performing a proper shutdown procedure.
- For a planned shutdown, make sure that all data has been written to disks before removing power. This is indicated by the Fast Write Cache LED which will be off when there is no longer any data in write cache. See Figure 13.
- If the BBU is removed, do not shut off power to the array unless the Fast Write Cache LED is off. Data in write cache will be posted to disk 10 seconds after the BBU is removed.

---

## Disk Array High Availability Features

High availability systems are designed to provide uninterrupted operation should a hardware failure occur. Disk arrays contribute to high availability by ensuring that user data remains accessible even when a disk or other component within the Disk Array FC60 fails. Selecting the proper Fibre Channel topology and system configuration can protect against the failure of any hardware component in the I/O path to the disk array by providing an alternate path to all user data.

The Disk Array FC60 provides high availability in the following ways:

- Supported RAID levels 1, 0/1, 3, and 5 all use data redundancy to protect data when a disk failure occurs. RAID 0 is supported but it does not offer data redundancy and should not be used in high-availability environments.
- Global hot spare disks serve as automatic replacements for failed disks.
- Alternate hardware paths to user data protects against I/O path failures.
- Redundant, hot-swappable hardware components can be replaced without interrupting disk array operation.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

### RAID Level Comparisons

To help you decide which RAID level to select for a LUN, the following tables compare the characteristics for the supported RAID levels. Where appropriate, the relative strengths and weakness of each RAID level are noted.

---

**Note** RAID 3 is supported on Windows NT and Windows 2000 only.

---

---

**Table 1** RAID Level Comparison: Data Redundancy Characteristics

RAID Level	Disk Striping	Mirroring	Parity	Handle multiple disk failures?
RAID 0	Yes	No	No	No. RAID 0 offers no data redundancy or protection against disk failure. RAID 0 should only be used for non-critical data. The failure of a single disk in a RAID 0 LUN will result in the loss of all data on the LUN.
RAID 1	No	Yes	No	No
RAID 0/1	Yes	Yes	No	Yes, providing both disks in a mirrored pair do not fail.
RAID 3	Yes	No	Yes	No
RAID 5	Yes	No	Yes	No

**Table 2** RAID Level Comparison: Storage Efficiency Characteristics

RAID Level	Storage Efficiency
RAID 0	100%. All disk space is use for data storage.
RAID 1 and 0/1	50%. All data is duplicated, requiring twice the disk storage for a given amount of data capacity.
RAID 3 and 5	One disk's worth of capacity from each LUN is required to store parity data. As the number of disks in the LUN increases, so does the storage efficiency.  3-disk LUN: 66% 4-disk LUN: 75% 5-disk LUN: 80% 6-disk LUN: 83%

**Table 3** RAID Level Comparison: Relative Performance Compared to an Individual Disk\*

LUN Configuration	Relative Read Performance for Large Sequential Access	Relative Write Performance for Large Sequential Access
RAID 0	The read and write performance of a RAID 0 LUN increases as the multiple of the number of disks in the LUN. For example, a 4-disk RAID 0 LUN will achieve close to four times the performance of a single disk.	
RAID 1 mirrored pair	Up to 2.0 > than single disk	Equal to single disk
RAID 0/1 group with 10 disks	Up to 10.0 > than single disk	Up to 5.0 > than single disk
RAID 0/1 group with 6 disks	Up to 6.0 > than single disk	Up to 3.0 > than single disk
RAID 3 group with 5 disks	Up to 4.0 > than single disk	Up to 1.25 > than single disk
RAID-5 group with 5 disks	Up to 4.0 > than single disk	Up to 1.25 > than single disk

\* Compares the relative read and write performance for array configurations with the performance of a single stand-alone disk whose performance is 1.0. The read and write performance shown is the theoretical maximum performance relative to individual disk performance. The performance numbers are not based on read/write caching. With caching, the performance numbers for RAID 5 writes improve significantly.

---

**Table 4** RAID Level Comparison: General Performance Characteristics

<b>RAID Level</b>	<b>General Performance Characteristics</b>
RAID 0	<ul style="list-style-type: none"><li>– Simultaneous access to multiple disks increases I/O performance. In general, the greater the number of mirrored pairs, the greater the increase in performance.</li></ul>
RAID 1	<ul style="list-style-type: none"><li>– A RAID 1 mirrored pair requires one I/O operation for a read and two I/O operations for a write, one to each disk in the pair.</li><li>– The disks in a RAID 1 mirrored pair are locked in synchronization, but the disk array can read data from the module whose read/write heads are the closest.</li><li>– RAID 1 read performance can be twice that of an individual disk. Write performance can be the same as that of an individual disk.</li></ul>
RAID 0/1	<ul style="list-style-type: none"><li>– Simultaneous access to multiple mirrored pairs increases I/O performance. In general, the greater the number of mirrored pairs, the greater the increase in performance.</li></ul>
RAID 3	<ul style="list-style-type: none"><li>– Provides high read throughput for large sequential I/Os.</li><li>– Write performance is limited by the need to perform four I/O operations per write request.</li><li>– Because some I/O operations occur simultaneously, performance depends on the number of disks in the LUN. Additional disks may improve performance.</li><li>– The I/O performance of RAID 5 benefits significantly from write caching.</li></ul>
RAID 5	<ul style="list-style-type: none"><li>– Provides high read throughput for small block-size requests (2 KB to 8 KB).</li><li>– Write performance is limited by the need to perform four I/O operations per write request.</li><li>– Because some I/O operations occur simultaneously, performance depends on the number of disks in the LUN. Additional disks may improve performance.</li><li>– The I/O performance of RAID 5 benefits significantly from write caching.</li></ul>



**Table 5** RAID Level Comparison: Application and I/O Pattern Performance Characteristics

RAID level	Application and I/O Pattern Performance
RAID 0	RAID 0 is a good choice in the following situations: <ul style="list-style-type: none"> <li>– Data protection is not critical. RAID 0 provides no data redundancy for protection against disk failure.</li> <li>– Useful for scratch files or other temporary data whose loss will not seriously impact system operation.</li> <li>– High performance is important.</li> </ul>
RAID 1	RAID 1 is a good choice in the following situations: <ul style="list-style-type: none"> <li>– Speed of write access is important.</li> <li>– Write activity is heavy.</li> <li>– Applications need logging or recordkeeping.</li> <li>– Daily updates need to be stored to a database residing on a RAID 5 group. The database updates on the RAID 1 group can be copied to the RAID 5 group during off-peak hours.</li> </ul>
RAID 0/1	RAID 0/1 is a good choice in the following situations: <ul style="list-style-type: none"> <li>– Speed of write access is important.</li> <li>– Write activity is heavy.</li> <li>– Applications need logging or recordkeeping.</li> <li>– Daily updates need to be stored to a database residing on a RAID 5 group. The database updates on the RAID 1 group can be copied to the RAID 5 group during off-peak hours.</li> </ul>
RAID 3	RAID 3 is a good choice in the following situations: <ul style="list-style-type: none"> <li>– Applications using I/O large sequential transfers of data, such as multimedia applications.</li> <li>– Applications on which write operations are 33% or less of all I/O operations.</li> </ul>
RAID 5	RAID 5 is a good choice in the following situations: <ul style="list-style-type: none"> <li>– Multi-tasking applications using I/O transfers of different sizes.</li> <li>– Database repositories or database servers on which write operations are 33% or less of all I/O operations.</li> <li>– Multi-tasking applications requiring a large history database with a high read rate.</li> <li>– Transaction processing is required.</li> </ul>

---

## Capacity Management Features

The disk array uses a number of features to manage its disk capacity efficiently. The use of LUNs allow you to divide the total disk capacity into smaller, more flexible partitions. Caching improves disk array performance by using controller RAM to temporarily store data during I/Os.

### Note

---

#### Differences in HP-UX and Windows NT/2000 Capacity Management

Capacity management on Windows NT and Windows 2000 offers some unique features. Refer to the *HP Storage Manager 60 Introduction Guide* for information on Windows-specific features. Some of the terms used in the *HP Storage Manager 60* software differ from those used here. These terms are also listed in the *HP Storage Manager 60 Introduction Guide*.

---



This topic is discussed in greater detail in the *Advanced User's Guide*, which can be downloaded from <http://www.hp.com/support/FC60>.

---

Overview .....	52
Array Design Considerations .....	53
Recommended Disk Array Configurations .....	59

---

---

## Overview

This chapter provides information to assist you in configuring the Disk Array FC60 to meet your specific storage needs. Factors to be considered when configuring the disk array include high availability requirements, performance, storage capacity, and future expandability. This chapter discusses configuration features of the Disk Array FC60 as it relates to these requirements.

**Note** Planning the disk array configuration is typically done before the disk array hardware is received. This information is then used during installation of the disk array to create the desired system configuration. After the disk array is installed, the information in this chapter will help you verify your configuration.

---

---

## Array Design Considerations

The Disk Array FC60 provides the versatility to meet varying application storage needs. To meet a specific application need, the array should be configured to optimize the features most important for the application. Array features include:

- High availability
- Performance (high I/O transfer rates)
- Storage capacity (optimize for lowest cost/Mbyte)
- Scalability

Optimizing array operation for a specific feature is managed by configuring various array installation and operation options. These options include: hardware configuration, RAID level and LUN creation, the number of SC10 disk enclosures used, and system support software. It should be noted that optimizing the disk array for one feature may compromise another. For example, optimizing for maximum performance may increase the cost per megabyte of storage.

### Array Hardware Configuration

Array configuration options that affect high availability, performance, and storage capacity include Fibre Channel connections, disk enclosure bus configuration, and internal Ultra2 SCSI channels. This information is presented first, because it is the basis for some of the array configuration planning.

#### Fibre Channel Connection

If the controller enclosure has both controller modules installed, dual Fibre Channels can be connected to the controller enclosure (if only one controller module is installed, only one Fibre Channel cable can be connected). Using dual Fibre Channel-AL connections increases the data throughput and provides for higher data availability.

#### Ultra2 SCSI Channel Operations

The disk array controller enclosure provides six Ultra2 SCSI channel connections for up to six disk enclosures. Six separate SCSI channels provide configuration flexibility. Disk

---

enclosures can be added incrementally (up to six) as storage requirements grow. Multiple SCSI channels also increase data throughput. This increased data throughput occurs as a result of the controller's ability to transfer data simultaneously over multiple data paths (channels). The more channels used, the faster the data throughput.

### **Disk Enclosure Bus Configuration**

The disk enclosure can connect to either one or two SCSI channels, depending on its bus configuration. Disk enclosure design allows the backplane bus to be split from a single bus into two separate SCSI busses. When the backplane is operating as a single bus it is referred to as full-bus mode; when the bus is split into two separate buses it is referred to as split-bus mode. When in full-bus mode, one SCSI channel connects to all of the ten disk modules. If the enclosure is configured for split-bus mode, five disk modules are connected to each of the two separate buses and a separate SCSI channel connects to each of the BCCs. See "Installation" on page 85 for more information.

When using split-bus mode, the maximum number of disk enclosures is limited to three (each disk enclosure uses two channel connections). If the storage capacity needs to be increased by adding more disk enclosures, the array will need to be reconfigured. Reconfiguring the array requires shutting down the host system, powering down the array, installing additional disk enclosures, and reconfiguring and recabling all enclosures.

### **RAID, LUNs, and Global Hot Spares**

In addition to the above hardware configuration considerations, the RAID level and LUN structure has considerable impact on high availability, performance, and storage capacity. For information on how RAID level, LUNs, and hot spares affect the performance of a disk array, see "Disk Array High Availability Features" on page 45.

### **High Availability**

If your application requires high availability, you should implement the options discussed here. The Disk Array FC60 is fully qualified to run MC/ServiceGuard and MC/LockManager. To work in these environments, a high availability configuration must be used. To configure

---

the array for high availability, there must be no single points of failure. This means that the configuration must have at least these minimum characteristics:

- Two controllers connected to separate Fibre Channel loops (using separate Fibre Channel host I/O adaptors)
- Two disk enclosures (minimum)
- Eight disk modules, four in each disk enclosure (minimum)
- LUNs that use only one disk per disk enclosure.

With its dual controllers, the Disk Array FC60 provides two independent I/O paths to the data stored on the array. Data is transmitted from the array controllers to the disks through up to six Ultra2 SCSI channels connected to the disk enclosures. Any of several RAID levels (1, 0/1, or 5) can be selected, however, for the greatest high availability, It is recommended that RAID level 1 be used for optimum high availability.

---

**Note** The Disk Array FC60 is designed to operate with either one or two controller modules, however for high availability it is highly recommended that two controller modules be installed.

---

## Performance

The maximum aggregate performance that can be sustained by the disk array is approximately 170 megabytes per second (using dual Fibre Channel-AL connections). This performance can be achieved by configuring at least four disk modules per Ultra2 SCSI bus utilizing the six Ultra2 SCSI channels. This can be accomplished in two ways. One way is to configure six disk enclosures, one each per Ultra2 SCSI channel (disk enclosure full-bus mode). In each of these enclosures, configure at least four disk modules.

Adding more disk modules to each of these disk enclosure will increase storage capacity, but will not appreciably increase the sequential throughput. Additional capacity may be a worthwhile addition, since in many computing environments, capacity, not access speed, is the limiting factor.

Another way to configure for maximum performance is to connect three disk enclosures to the controller enclosure and configure these enclosures for split-bus operation. Then connect an Ultra2 SCSI channel to each split bus (two channels per disk enclosure). Each

---

of the buses must be configured with at least four disk modules (eight disk modules per disk enclosure). This configuration also offers full sequential performance and is more economical to implement.

To scale up sequential transfer performance from the host, configure additional disk arrays. This will increase the total I/O bandwidth available to the server.

Performance can also be measured by the number of I/O operations per second a system can perform. I/Os per second are important in OLTP (on-line transaction processing) applications. To maximize I/Os per second, configure the maximum number of disk modules. For the same capacity, you may elect to use a larger number of 9.1 Gbyte disk modules instead of a smaller number of higher capacity disk modules to obtain optimal I/Os per second.

---

**Note** For the maximum I/Os per second, configure RAID 0/1 and the maximum number of disk modules.

---



---

## Storage Capacity

For configurations where maximum storage capacity at minimum cost is a requirement, consider configuring the disk array in RAID 5 (using the maximum number of data drives per parity drives) and only supplying one or two hot spare drives per disk array. Also, purchase the lowest cost/Mbyte drive available (typically the largest capacity drives available at the time of purchase). This configuration allows the maximum volume of storage at the lowest cost. Disk arrays configured in this way will need to be carefully monitored to make sure that failed disks are promptly replaced.

## Expanding Storage Capacity

The disk array is designed to meet a range of capacity requirements. It can be configured with up to six disk enclosures, and from 4 to 60 disk modules. The disk array can be purchased with enough capacity to meet your current storage requirements. As your storage needs grow, you can easily add more capacity to the Disk Array FC60.

There are several ways to increase the storage capacity of the disk array. It can be increased by replacing existing smaller capacity disk modules with larger capacity disk modules, by adding more disk modules to the disk enclosures, or by adding additional disk enclosures.

The best method for expansion is to install all six disk enclosures at initial installation (full-bus configuration). Then, install only the required capacity (number of disk modules), leaving empty disk slots for future expansion. Additional disk modules can be installed as the requirement for additional capacity grows. This method allows for greater flexibility in LUN creation and does not require that the system be shut down for expansion. Adding disk enclosures to the array after initial installation requires that the system be shut down during the installation. All the remaining configuration expansion methods require that the data to the array be suspended (the system shut down) to add additional disk enclosure storage to the array.

---

**Note** For maximum performance, all six SCSI channels from the controller enclosure must be connected with a minimum of two disk modules per channel (disk enclosure configured for full bus mode).

---

To expand a split-bus configuration (adding more disk enclosures) will require shutting the host system down. If the initial installation included only one or two disk enclosures, then

---

another, two or one disk enclosures, respectively, can be added by using split-bus mode. However, if you are adding up to four, five, or six enclosure, the enclosures configuration will need to be switched from split-bus to full-bus (refer to “Disk Enclosure Bus Configuration” section, earlier in this chapter, for additional information).

**Note** Typically adding only one enclosure does not provide any options for creating LUNs. It is best to expand the array with at least a minimum of two disk enclosures at a time. Two additional drives could then be configured as RAID 1 or 0/1.

Installing a five disk enclosure array limits expansion to six disk enclosures. Adding one additional enclosure does not provide any versatility for creating LUNs (unless all data is removed and the LUNs are rebuilt).

---

If the initial installation uses split-bus disk enclosure (split-bus accepts three disk enclosures maximum) and expansion requires adding four or more enclosures, the existing disk enclosures will need to be reconfigured for full-bus mode and the additional number of enclosures installed into the rack and required number of full-bus enclosures. This expansion requires disk enclosures be recabled using full bus enclosures. As in all cases of adding disk enclosures to the array, the system has to be shut down for the expansion. Determine the RAID level and how the LUNs will be created for the expansion

If the initial installation consisted of one or more full-bus configured disk enclosures, then additional full-bus configured disk enclosures can be added to the array. The system should be shut down for the addition of the enclosures. Determine the RAID level and how the LUNs will be created for the additional storage.

To scale up sequential performance, first make sure that the configuration includes both controllers modules. Maximum sequential transfer performance will be reached with approximately 20 disk modules simultaneously transferring data. To achieve additional sequential transfer performance, you will need to add a second disk array and more disk modules.

To increase I/Os per second performance, add disk modules. Transaction performance is directly related to the number of disk modules installed in the disk array.

---

## Recommended Disk Array Configurations

This section presents recommended configurations for disk arrays using one to six disk enclosures. Configurations are provided for achieving high availability/high performance, and maximum capacity. The configuration recommended by Hewlett-Packard is the high availability/ high performance configuration, which is used for factory assembled disk arrays (A5277AZ). The configurations identify the number of disk enclosures, cable connections, disk enclosure bus modes, RAID level, and LUN structure.

Most of the configurations offer the highest level of availability, which means they are capable of surviving the failure of a single disk (provided LUNs are created with one disk module per disk enclosure), SCSI channel, disk enclosure, or controller module. The only configurations that do not offer the highest level of availability are the single disk enclosure configuration, and the two enclosure high capacity configuration. These configurations cannot survive the failure of an entire disk enclosure, so they should not be used in environments where high availability is critical.

The configurations list maximum disk capacity and usable disk capacity (with ten disk modules installed). Configurations based on RAID 1 have less usable disk capacity than RAID 5, but I/O performance is optimized when using RAID 1. Although the recommended configurations presented here all contain ten disk modules, a disk enclosure can contain four, eight, or ten disk modules.

---

**Note**      The terms “LUN” and “volume group” are used interchangeably in the text and figures in this section.

---

### Configuration Considerations

The following factors should be considered when using any of the recommended configurations.

- **Multiple Hosts** - A single host system is shown, but configurations can be adapted to create multi-host, high availability systems.
- **Global hot spares** - although none of the configurations use global hot spares, their use is recommended to achieve maximum protection against disk failure.

- 
- **Split bus operation** - With three or fewer disk enclosures, increased performance can be achieved by operating the disk enclosures in split bus mode, which increases the number of SCSI busses available for data transfer. However, operating the disk enclosures in split bus mode may make it more difficult to expand the capacity of the array. In a split bus configuration, it may be necessary to take down the host, back up data, and rebind LUNs when adding disk enclosures. If you anticipate the need to expand your disk array, you may want to consider selecting a configuration that uses more enclosures operating in full bus mode. In addition to simplifying expansion, this type of configuration also gives you greater flexibility when creating LUNs.
  - **Segment size** - the recommended segment size is 16 Kbyte for RAID 5, and 64 Kbyte for RAID 1 and RAID 0/1.
  - **Maximum LUNs** - A maximum of 30 LUNs can be configured on the disk array.

## One Disk Enclosure Configuration

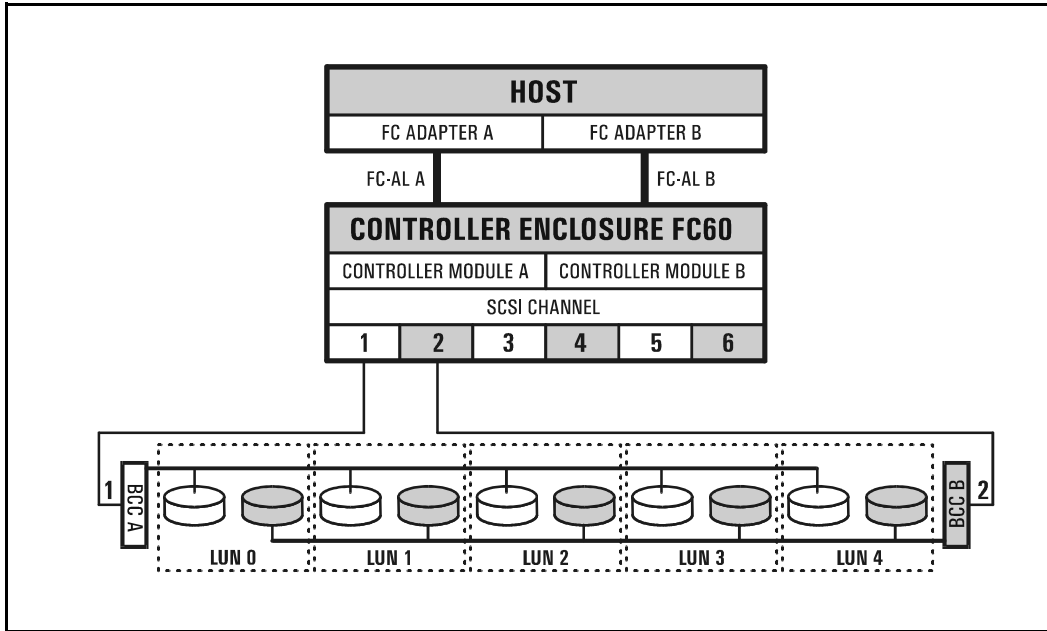
---

**Note** A single disk enclosure configuration is not recommended for environments where high availability is critical. For optimum high availability, at least two disk enclosures are required. This protects against the failure of a single disk enclosure.

---

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - One disk enclosure with ten 73 GBytes disk modules
  - Disk enclosure configured for split-bus mode (two SCSI channels)
- **LUN Configuration**
  - Five RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is on a separate SCSI channel
- **Data Availability**
  - Not recommended for maximum high availability.
  - Handles a single disk failure, single BCC failure, a single channel failure, or a single controller failure
  - Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.

- **Disk Capacity**
  - Maximum capacity 730 GBytes
  - Usable capacity 365 GBytes



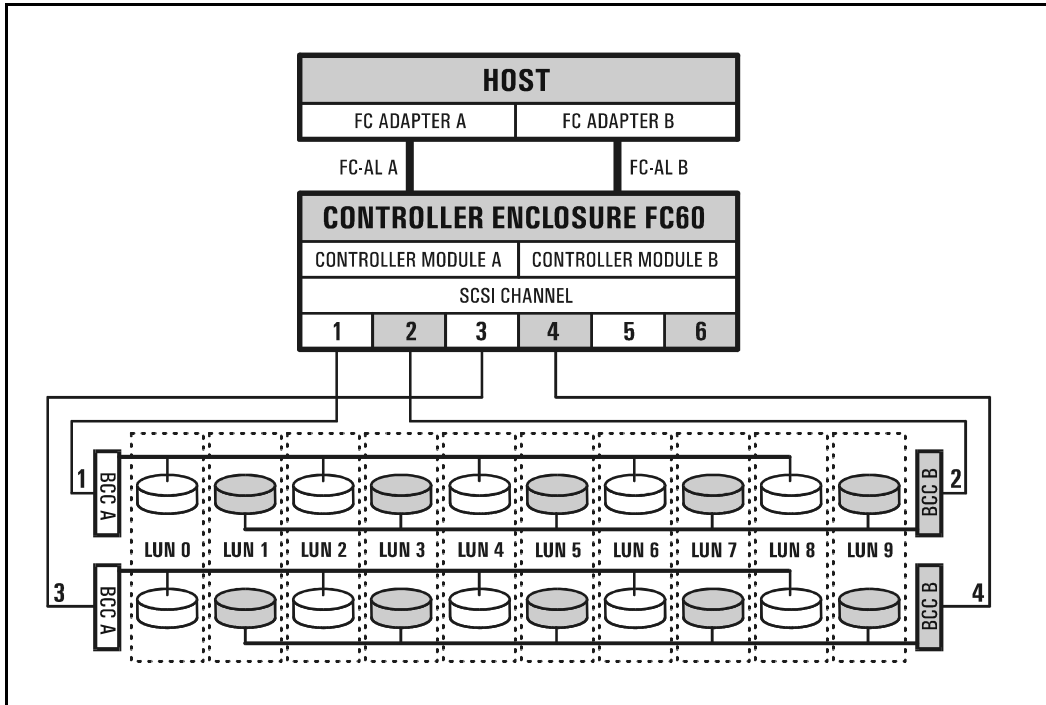
**Figure 17** One Disk Enclosure Array Configuration

---

## Two Disk Enclosure Configurations

### High Availability/ High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Two disk enclosures with ten 73 GByte disk modules (20 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - Ten RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, BCC failure, single channel failure, or a single controller failure.
  - Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.
- **Disk Capacity**
  - Maximum capacity 1460 GBytes
  - Usable capacity 730 GBytes



**Figure 18** Two Disk Enclosure High Availability/ High Performance Configuration

---

## Maximum Capacity

---

**Note** This configuration is not recommended for environments where high availability is critical. To achieve high availability each disk in a LUN should be in a different disk enclosure. This configuration does not achieve that level of protection.

---

- **Hardware Configuration**

- Two disk array controllers connected directly to host Fibre Channel adapters
- Two disk enclosures with ten 73 GByte disks each (20 disks total)
- Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)

- **LUN Configuration**

- Five RAID 5 LUNs, each comprising four disks (3 data + 1 parity)
- Each disk in a LUN is on a separate SCSI bus.

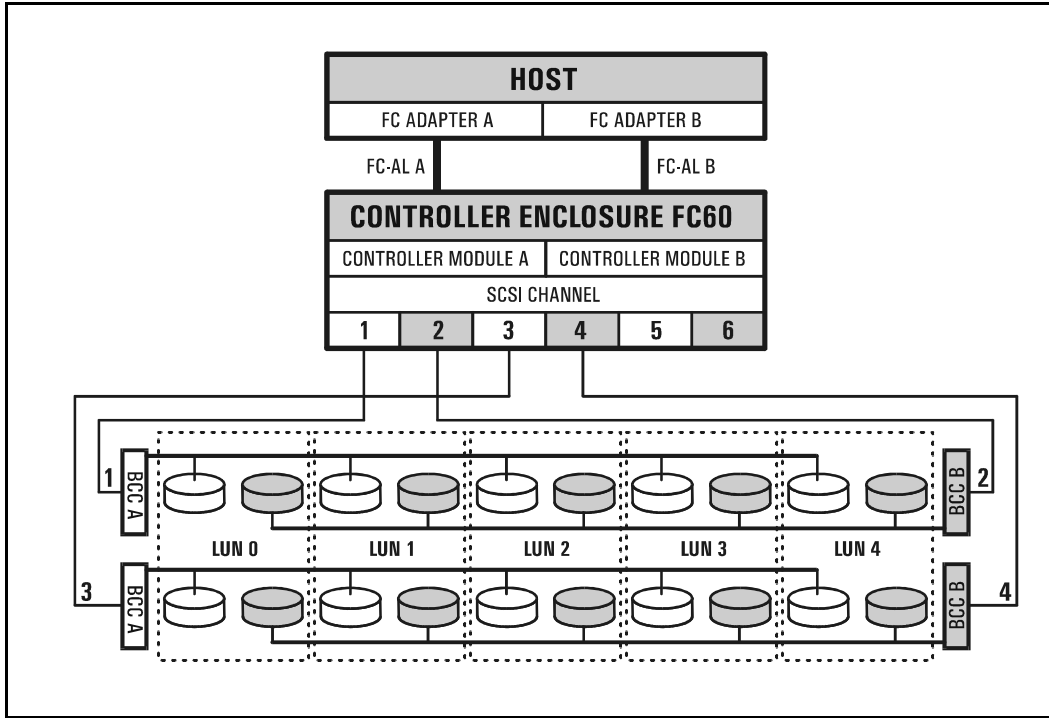
- **High Availability**

- Handles the failure of a single disk, single controller, or a single channel
- Does not handle a disk enclosure failure, consequently this configuration is NOT recommended for critical high availability installations.
- Expansion requires powering down the disk array, removing terminators and/or cables from the enclosures, and cabling additional disk enclosures.

- **Disk Capacity**

- Maximum capacity 1460 GBytes
- Usable capacity 1095 GBytes






**Figure 19** Two Disk Enclosure Maximum Capacity Configuration

---

## Three Disk Enclosure Configurations

### High Availability/ High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Three disk enclosures with ten 73 GByte disks each (30 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - 15 RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, a single controller, a single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array, recabling the array to a full bus configuration, rebinding the LUNs, and restoring all data
- **Disk Capacity**
  - Maximum capacity 2190 GBytes
  - Usable capacity 1095 GBytes




**Figure 20** Three Disk Enclosure High Availability/ High Performance Configuration

---

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Three disk enclosures with ten 73 GByte disks each (30 disks total)
  - Disk enclosures configured for split-bus mode (two SCSI channels per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising three disks (2 data + 1 parity).
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single controller, single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array, recabling the array to a full bus configuration, rebinding the LUNs, and restoring all data
- **Disk Capacity**
  - Maximum capacity 2190 GBytes
  - Usable capacity 1460 GBytes




**Figure 21** Three Disk Enclosure Maximum Capacity Configuration

---

## Four Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Four disk enclosures with ten 73 GByte disks each (40 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising four disks (2+2)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array and adding additional disk enclosures and cables
- **Disk Capacity**
  - Maximum capacity 2920 GBytes
  - Usable capacity 1460 GBytes




**Figure 22** Four Disk Enclosure High Availability/High Performance Configuration

---

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Four disk enclosures with ten 73 GByte disks each (40 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising four disks (3 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single controller, single channel, single BCC, or a single disk enclosure failure
  - Expansion requires powering down the disk array and adding additional disk enclosures and cabling
- **Disk Capacity**
  - Maximum capacity 2920 GBytes
  - Usable capacity 2190 GBytes





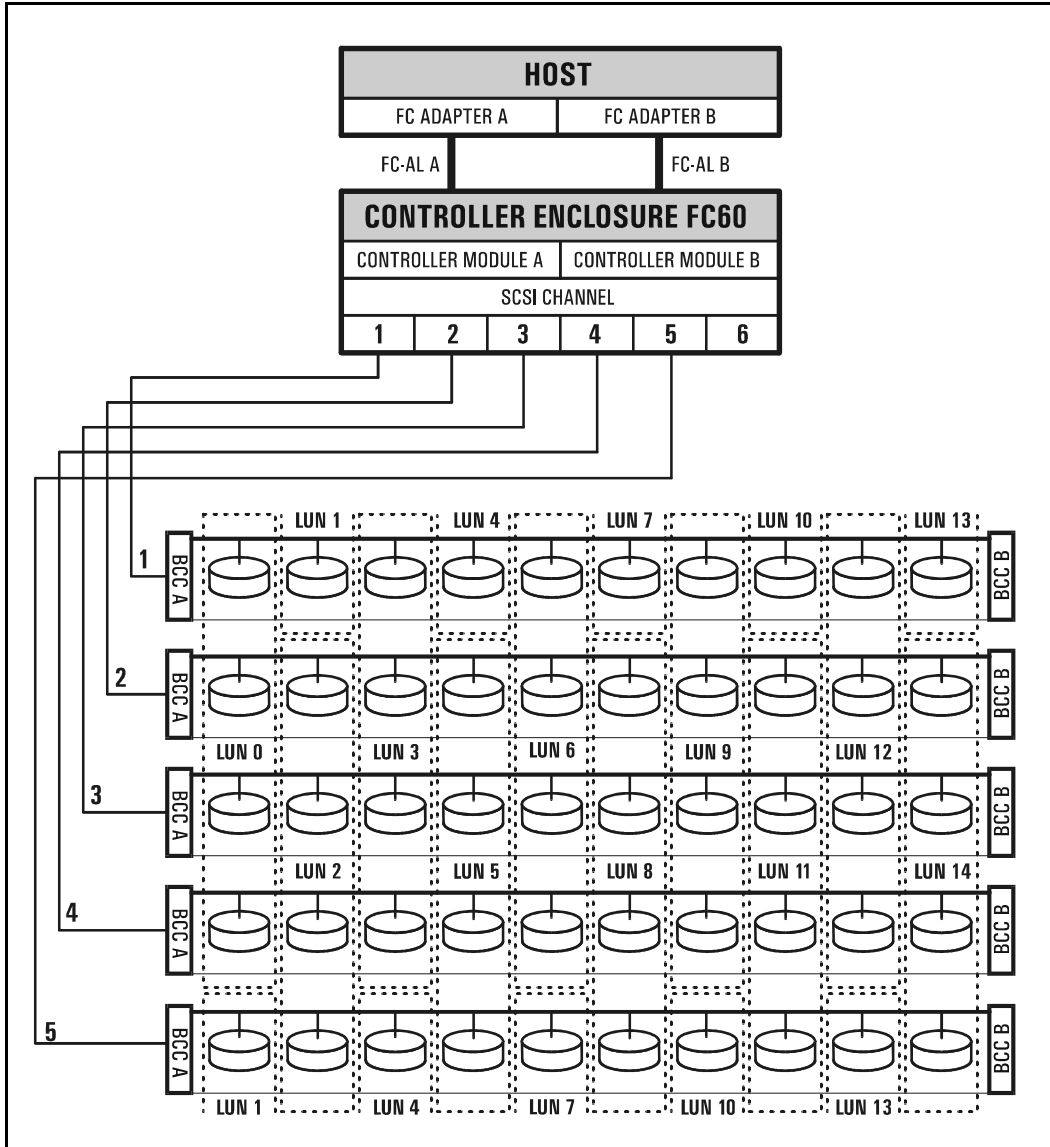
**Figure 23** Four Disk Enclosure Maximum Capacity Configuration

---

## Five Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Five disk enclosures with ten 73 GByte disks each (50 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising four disks (2+2)
  - Five RAID 1 LUNs, each comprising two disks (1+1)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array and adding an additional disk enclosure and cabling
- **Disk Capacity**
  - Maximum capacity 3650 GBytes
  - Usable capacity 1825 GBytes

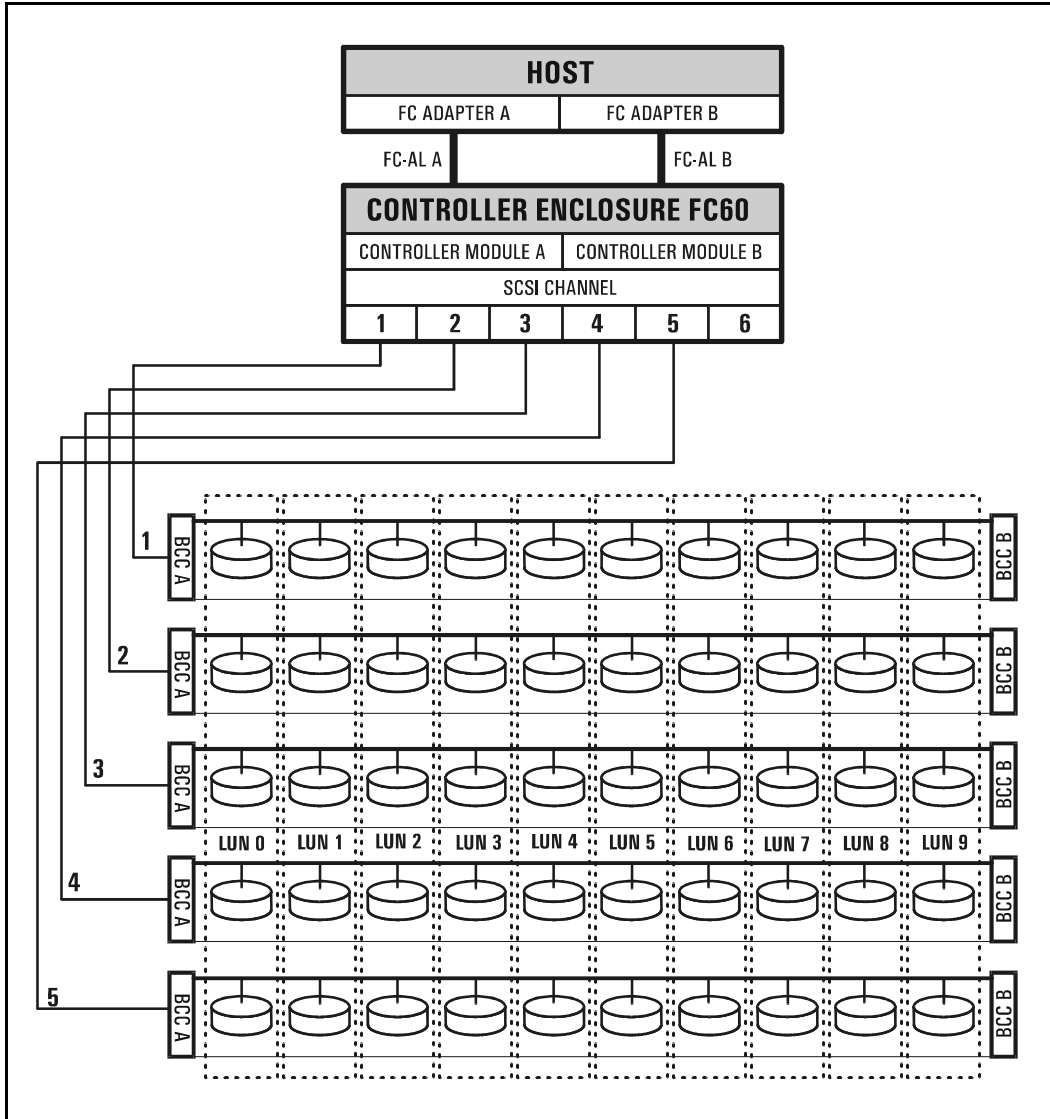


**Figure 24** Five Disk Enclosure High Availability/High Performance Configuration

---

## Maximum Capacity

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Five disk enclosures with ten 73 GByte disks each (50 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising five disks (4 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure.
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
  - Expansion requires powering down the disk array, and adding an additional disk enclosure and cabling
- **Disk Capacity**
  - Maximum capacity 3650 GBytes
  - Usable capacity 2920 GBytes




**Figure 25** Five Disk Enclosure Maximum Capacity Configuration

---

## Six Disk Enclosure Configurations

### High Availability/High Performance

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Six disk enclosures with ten 73 GByte disks each (60 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 0/1 LUNs, each comprising six disks (3+3)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, single BCC failure, single channel failure, or a single controller failure
- **Disk Capacity**
  - Maximum capacity 4380 GBytes
  - Usable capacity 2190 GBytes




**Figure 26** Six Disk Enclosure High Availability/High Performance Configuration

---

## **Maximum Capacity**

- **Hardware Configuration**
  - Two disk array controllers connected directly to host Fibre Channel adapters
  - Six disk enclosures with ten 73 GByte disks each (60 disks total)
  - Disk enclosures configured for full-bus mode (one SCSI channel per enclosure)
- **LUN Configuration**
  - Ten RAID 5 LUNs, each comprising six disks (5 data + 1 parity)
  - Each disk in a LUN is in a separate enclosure
- **High Availability**
  - Handles a single disk failure, single disk enclosure/BCC failure, single channel failure, or a single controller failure
- **Disk Capacity**
  - Maximum capacity 4380 GBytes
  - Usable capacity 3650 GBytes





**Figure 27** Six Disk Enclosure High Maximum Capacity Configuration

---

## **Total Disk Array Capacity**

The total capacity provided by the disk array depends on the number and capacity of disks installed in the array, and the RAID levels used. RAID levels are selected to optimize performance or capacity.

Table 6 lists the total capacities available when using fully loaded disk enclosures configured for optimum performance. Table 7 lists the same for optimum capacity configurations.

The capacities listed reflect the maximum capacity of the LUN. The actual storage capacity available to the operating system will be slightly less, as some capacity is consumed when binding the LUN and creating the file system.

For high-availability, one disk per SCSI channel is used as a global hot spare.

**Table 6** Capacities for Optimized Performance Configurations

Number of disk enclosures	RAID Level	No. of LUNs	Disks per LUN	Total Capacity (with indicated disks)			
				9.1 GB	18.2 GB	36.4 GB	73 GB
2 (split bus)	1	8	2	72.8 GB	145.6 GB	291.2 GB	584 GB
3 (split bus)	1	12	2	109.1 GB	218.4 GB	436.8 GB	876 GB
4 (full bus)	0/1	9	4 (2+2)	163.8 GB	327.6 GB	655.2 GB	1314 GB
6 (full bus)	0/1	9	6 (3+3)	245.7 GB	491.4 GB	982.8 GB	1971 GB

**Table 7** Capacities for Optimized Capacity Configurations

Number of disk enclosures	RAID Level	No. of LUNs	Disks per LUN	Total Capacity (with indicated disks)			
				9.1 GB	18.2 GB	36.2 GB	73 GB
3 (split bus)	5	8	3 (2D+1P)	145.6 GB	291.2 GB	582.4 GB	1168 GB
4 (full bus)	5	9	4 (3D+1P)	245.7 GB	491.4 GB	982.8 GB	1971 GB
5 (full bus)	5	9	5 (4D+1P)	327.6 GB	655.2 GB	1310.4 GB	2628 GB
6 (full bus)	5	9	6 (5D+1P)	409.5 GB	819 GB	1638 GB	3285 GB



<b>Configuration Switches . . . . .</b>	<b>89</b>
<b>Attaching Power Cords . . . . .</b>	<b>97</b>
<b>Attaching SCSI Cables and Configuring the Disk Enclosure Switches . . . . .</b>	<b>98</b>
<b>Connecting the Fibre Channel Cables . . . . .</b>	<b>107</b>
<b>Applying Power to the Disk Array . . . . .</b>	<b>109</b>
<b>Solving Common Installation Problems. . . . .</b>	<b>117</b>

---

---

## Overview

Installation of the Disk Array FC60 must be performed by a service-trained representative. This chapter contains information that may be useful during the ongoing operation and management of the disk array. Tasks such as setting configuration switches, connecting cabling, and powering the disk array on and off are discussed in this chapter.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

## Host System Requirements

### HP-UX

The Disk Array FC60 is supported on the following host configurations:

- **Supported host platforms:** Table 8 lists the supported host platforms. The table also identifies which platforms support the disk array as a boot device when running HP-UX 11.x
- **Supported HP-UX versions:** 11.0, 11.11, and 10.20

---

**Note** The Disk Array FC60 is supported on HP-UX 11.11 with array firmware release HP03 only. Later versions of disk array firmware are not currently supported when running HP-UX 11.11.

---

- **Software requirements:** The Disk Array FC60 management software must be installed on the host. This software is distributed on *HP-UX Support Plus CD-ROM*, release 9909 or later.

**Table 8 Supported Host Platform Information**

Supported Host	Boot Support on HP-UX 11.x?	Fibre Channel I/O Adapter
K-class	Yes	A3404A
V-class	Yes	A5158A, A3740A
L-class	Yes	A5158A, A3740A
D-class	Yes	A3591B
N-class	Yes	A5158A, A3740A
R-class	Yes	A3591B
T-class	No	A3636A

---

**Table 8 Supported Host Platform Information (cont'd)**

<b>Supported Host</b>	<b>Boot Support on HP-UX 11.x?</b>	<b>Fibre Channel I/O Adapter</b>
C-class	No	A5158A on HP-UX 11.x A3740A on HP-UX 10.20
A4xx-A5xx class	Yes	A5158A

## **Windows NT and Windows 2000**

For information on Windows NT and Windows 2000 host system requirements, see the *HP Storage Manager 60 User's Guide* included with the *HP Storage Manager 60* software (A5628A).



---

## Configuration Switches

This section describes the configuration switches on the controller enclosure and the disk enclosures. Configuration switch settings must be set to support the configuration (full-bus or split-bus) being installed (as planned for from chapter 2, Array Planning). Controller enclosure and the disk enclosures configuration switches include:

- Disk enclosure (Tray) ID switches
- Disk enclosure DIP switches
- Controller enclosure Fibre Channel host ID address switches

### Tips for Setting Configuration Switches

- All switch settings should be the same on both BCC modules on the disk enclosures.
- Disk enclosures must be power cycled before the new switch setting will take effect.
- If a BCC is installed that has settings different than its peer, a buzzer will sound. The settings must be corrected to use the disk enclosure.

### Disk Enclosure (Tray) ID Switch

The disk enclosure ID switch sets the Enclosure ID (or Tray ID). This value is used by FC60 diagnostic software and the controller to identify the disk enclosure for status monitoring operations. Each enclosure in the array must be set to a different Enclosure ID. To set the ID, set the rotary ID switches on the back of both BCC modules to the same number. See Figure 28. Typically, the disk enclosure closest to the controller enclosure is set to “0”, the next one down to “1”, and so on down the rack. See "Selecting Disks for a RAID 0/1 LUN" on page 124 for additional disk module addressing information.)

### Disk Enclosure DIP Switch


A DIP switch located on the back of each BCC is used to set configuration options for BCC operation. See Figure 28. These DIP switches contain five switch segments. Switch segment 1 is used to select either split-bus or full-bus operation. See "Full-Bus/Split-Bus

---

(Switch 1) Configuration" on page 92. The remaining four switch segments are not used for disk array operation and must always be set to "0."

**Note** Note that one BCC is inverted with respect to the other. Thus, the settings on one BCC appear as inverted and in reverse order from the other.

---



**Figure 28** Disk Enclosure Configuration Switches

---

**Table 9** Disk Enclosure Switches

Switch	Setting	Operation
1 - Full-/Split Bus Mode	1	Full-Bus Mode
	0	Split-Bus Mode
2 - Stand-Alone/Array Mode	0	Always set to Off (Array Mode)
3 - Bus Reset on Power Fail	0	Must be set to 0
4 - High/Low Bus Addressing	0	Set to 0 (Low addressing)
5 - Not Used	0	Not used; must be set to 0

### Full-Bus/Split-Bus (Switch 1) Configuration

The disk enclosure's internal bus connects the disk drives together and to the BCCs. This bus can be set for two modes of operation: full-bus or split-bus mode. When switch 1 is set to Full Bus (1), all ten disk can be accessed together on a single SCSI bus. In this mode, one SCSI channel from the FC60 controller is connected to all ten disk modules through either BCC (typically BCC A).

---

**Note** In full bus mode, all ten disks can be accessed through either BCC. However, internally each BCC still manages five disks. This means that if the BCC that is not connected to a SCSI channel fails, access to its five disks will be lost. Failure of the BCC that is connected to a SCSI channel will render all ten disks inaccessible.

---

When switch 1 is set to 0 (split-bus mode), the disk enclosure is split into two separate buses with five disks to each bus. In split-bus mode, each bus is controlled by a separate BCC. BCC A (the top BCC) controls the even drive slots (0, 2, 4, 6, 8) and BCC B (the bottom BCC) controls the odd bus slots (1, 3, 5, 7, 9). In split-bus mode, two SCSI channels are connected to the disk enclosure, one to a SCSI connector on BCC A and a second a SCSI connector on BCC B.

Each slot in the disk enclosure is assigned a SCSI ID. A range of IDs is assigned to the odd disk slots, and a different range to the even slots (see Table 10). The two ID ranges include a low range of IDs (0, 1, 2, 3, and 4) and a high range of IDs (8, 9, 10, 11, and 12). (BCCs are also provided addresses as shown in Table 10). Note that the SCSI IDs do not correspond to the physical slot number.

The assignment of the SCSI IDs differs depending on whether the enclosure is operating in full-bus or split-bus mode. When full-bus mode is selected, the low ID range (0 - 4) is assigned to the even disk slots, and the high range (8 - 12) is assigned to the odd slots. See Table 10. When the disk enclosure is in split-bus mode, the low ID range is assigned to both the even slots and the odd slots. This is possible because the two busses are isolated within the enclosure. See "Selecting Disks for a RAID 0/1 LUN" on page 124 for additional disk module addressing information.

**Table 10** Disk Slot, Full-Bus/Split-Bus Mode, SCSI IDs

	Physical Disk Slot #s										BCC	
	0	1	2	3	4	5	6	7	8	9	A	B
<b>Full-Bus Mode SCSI IDs</b>												
BCC A (or B)	0	8	1	9	2	10	3	11	4	12	14	15
<b>Split-Bus Mode SCSI IDs</b>												
BCC-A (Even Slots)	0	-	1	-	2	-	3	-	4	-	15	-
BCC-B (Odd Slots)	-	0	-	1	-	2	-	3	-	4	-	15

## Fibre Channel Host ID Address Setting

The controller enclosure has two Fibre Channel connections, one for each controller module. Each controller module communicates with the host through its Fibre Channel loop connection. Internally the two controllers are connected to a common Fibre Channel loop. Because they are connected to a common loop, each controller module must have a unique Fibre Channel Host ID.

---


The Fibre Channel Host IDs are selected by two separate DIP switches located on the back of the controller enclosure. See Figure 29. Switch Host ID BD1 SW1 selects the address for controller module A (Fibre Channel connector J3) and Host ID BD2 SW2 selects the address for controller module B (Fibre Channel connector J4).

Each Fibre Channel Host ID DIP switch contains a bank of seven switches that select the address using a binary value, 000 0000 (0) through 111 1111 (126). To set an address, set the switches in the up position for “1” or down for “0” (refer to Table 11 for binary switch settings). Figure 29 illustrates the loop ID switch set to 42 (0101010).

---

**Note** The factory default Fibre Channel Host ID loop address settings for controller A (BD1) is 5, and the default for controller B (BD2) is 4. If you must change the default values, make sure you select a unique value for each ID. The Fibre Channel Host ID loop address settings must be unique for each controller, regardless of the external Fibre Channel topology.

---



**Figure 29** Fibre Channel Connectors and Fibre Channel Host (Loop) ID Switches

**Note** Occasionally two or more ports in an arbitrated loop will arbitrate simultaneously. Priorities are decided according to the loop IDs. The higher the loop ID, the higher the priority.

**Table 11** Fibre Channel Addresses

Decimal	Binary	Decimal	Binary	Decimal	Binary	Decimal	Binary
0	000 0000	38	010 0110	76	100 1100	114	111 0010
1	000 0001	39	010 0111	77	100 1101	115	111 0011
2	000 0010	40	010 1000	78	100 1110	116	111 0100
3	000 0011	41	010 1001	79	100 1111	117	111 0101
4	000 0100	42	010 1010	80	101 0000	118	111 0110
5	000 0101	43	010 1011	81	101 0001	119	111 0111
6	000 0110	44	010 1100	82	101 0010	120	111 1000
7	000 0111	45	010 1101	83	101 0011	121	111 1001
8	000 1000	46	010 1110	84	101 0100	122	111 1010
9	000 1001	47	010 1111	85	101 0101	123	111 1011
10	000 1010	48	011 0000	86	101 0110	124	111 1100
11	000 1011	49	011 0001	87	101 0111	125	111 1101
12	000 1100	50	011 0010	88	101 1000	126	111 1110
13	000 1101	51	011 0011	89	101 1001	127	111 1111
14	000 1110	52	011 0100	90	101 1010		
15	000 1111	53	011 0101	91	101 1011		
16	001 0000	54	011 0110	92	101 1100		
17	001 0001	55	011 0111	93	101 1101		
18	001 0010	56	011 1000	94	101 1110		
19	001 0011	57	011 1001	95	101 1111		
20	001 0100	58	011 1010	96	110 0000		
21	001 0101	59	011 1011	97	110 0001		
22	001 0110	60	011 1100	98	110 0010		
23	001 0111	61	011 1101	99	110 0011		
24	001 1000	62	011 1110	100	110 0100		
25	001 1001	63	011 1111	101	110 0101		
26	001 1010	64	100 0000	102	110 0110		
27	001 1011	65	100 0001	103	110 0111		
28	001 1100	66	100 0010	104	110 1000		
29	001 1101	67	100 0011	105	110 1001		
30	001 1110	68	100 0100	106	110 1010		
31	001 1111	69	100 0101	107	110 1011		
32	010 0000	70	100 0110	108	110 1100		
33	010 0001	71	100 0111	109	110 1101		
34	010 0010	72	100 1000	110	110 1110		
35	010 0011	73	100 1001	111	110 1111		
36	010 0100	74	100 1010	112	111 0000		
37	010 0101	75	100 1011	113	111 0001		



---

## Attaching Power Cords

Each enclosure (controller and disk enclosures) contains dual power supplies that must be connected to the power source (PDU). When connecting power cords for high availability installations, connect one enclosure power cord to one source (PDU) and the other power cord to an alternate source (PDU). To complete the power connection, follow the steps below.

1. Set the power switch on the disk enclosures to OFF. The switch is located at the front, top, right corner of the enclosure.
2. Set the two power switches on the back of the controller enclosure to OFF (0).
3. Connect power cords to both power supplies for each enclosure. Attach one end of the power cord to the enclosure power supply receptacle and the other end to the PDU/PDRU receptacle.

---

**Note** The controller enclosure requires special ferrite bead power cords (part no. 5064-2482) for FCC compliance.

---

When connecting power cords, connect the power cords from the right-side power supplies to the PDUs/PDRUs on the right side of the rack and connect the power cords from the left-side power supplies to the PDUs/PDRUs on the left side of the rack. This method ensures that each enclosure power supply is connected to a separate power source to ensure high availability.

4. Connect the PDUs/PDRUs into their respective power sources.

---

## Attaching SCSI Cables and Configuring the Disk Enclosure Switches



**NOTE!** It is critical that all SCSI cables be tightened securely. Use the following steps to ensure the cable connectors are seated properly.

1. Connect the cable to the enclosure connector and tighten the mounting screws finger tight.
2. Push on the connector and retighten the mounting screws. Repeat once more.
3. Use a flat blade screwdriver to tighten the screw 1/4 turn. Be sure the screw is not cross-threaded.

There are two types of cabling configurations for the array: full-bus and split-bus. The configuration you select is determined by a number of factors which are described in Chapter 2, Array Planning.

When you have determined which configuration you will use, refer to "Full-Bus Cabling and Switch Configuration" on page 99" or "Split-Bus Switch and Cabling Configurations" on page 102".

### **A Word About SCSI Terminators**

Terminators need to be installed on both the disk enclosures and the controller enclosure. On the controller enclosure, any unused SCSI connector must be terminated. On the disk enclosure, not every connector requires termination. If a SCSI cable is connected to a BCC connector, the other connector on the BCC must be terminated. If there is no SCSI cable connected to the BCC, no terminator is required on either connector. A BCC that is operating in split-bus mode will have a cable connected to each BCC. In this case, the unused connector on each BCC must be terminated. See Figure 31 or Figure 33.

---

## Full-Bus Cabling and Switch Configuration

Cabling for a full bus configuration requires connecting one SCSI cable from the controller to the disk enclosure and setting the configuration switches. Figure 31 illustrates full-bus cabling connections for a six disk enclosure array. It is possible to configure any number of disk enclosures, from one to six, using this method. However, full bus is typically used when four or more disk enclosures are installed in the array.

---

**Note** For consistency and ease of management, it is recommended that you observe the following rules:

- Connect channel 1 on the controller enclosure to the disk enclosure mounted directly beneath the controller enclosure. Connect channel 2 to the next disk enclosure down, and so on down the rack.
- Install cables along the left side of the disk enclosure connectors (as viewed from the back of the rack) and install terminators along the right side connectors.

---

To configure the disk array for full-bus operation:


1. Connect one end of a SCSI cable to the controller enclosure SCSI channel 1 connector, J12 JVD ARRAY SCSI 1. See Figure 31.
2. Connect the other end of the SCSI cable to the disk enclosure BCC A, SCSI connector B.

---

**Note** The lower BCC is inverted from the upper BCC. Make sure you take this into account when setting the switches.

---

3. Set the Enclosure (Tray) ID rotary switches on both BCCs to “0”. See Figure 30.
4. Set the Enclosure DIP switches on both BCCs to the following settings:
  - a. Set segment 1 to “1”
  - b. Set the remaining four segments to “0”. See Figure 30.
5. Connect a terminator to SCSI connector A on BCC A (see Figure 31, “T”).
6. Repeat the preceding steps for each disk enclosure, using the next sequential SCSI channel and incrementing the Disk Enclosure (Tray) ID setting for each enclosure.



**Figure 30** Full Bus BCC Configuration Switch Settings




Figure 31 Full-Bus Cabling

---

## Split-Bus Switch and Cabling Configurations

Split-bus cabling requires two SCSI cables from each disk enclosure to the controller enclosure. Split-bus cabling is typically used for installations with 3 or fewer disk enclosures. Cabling for a split-bus configuration is shown in Figure 33.


- 
- Note** For consistency and ease of management, it is recommended that you observe the following rules:
- Connect channels 1 and 2 on the controller enclosure to the disk enclosure mounted directly beneath the controller enclosure. Connect channels 3 and 4 to the next disk enclosure down, and channels 5 and 6 to the next enclosure down.
  - Install cables along the left side of the disk enclosure connectors (as viewed from the back of the rack) and install terminators along the right side connectors.
- 

To configure the disk array for split-bus operation:


1. Connect one end of a SCSI cable to the controller enclosure SCSI channel 1 connector, J12 JVD ARRAY SCSI 1 (see Figure 33).
2. Connect the other end of the SCSI cable to the disk enclosure BCC A, SCSI connector B.
3. Connect a second SCSI cable to the controller enclosure SCSI channel 2 connector, J11 JVD ARRAY SCSI 2.
4. Connect the other end of the SCSI cable to the disk enclosure BCC B, SCSI connector A .

- 
- Note** The lower BCC is inverted from the upper BCC. Make sure you take this into account when setting the switches.
- 

5. Set the Enclosure (Tray) ID rotary switches on both BCCs to “0”. See Figure 32.
6. Set all segments of the Enclosure DIP switches on both BCCs to “0”.
7. Connect a SCSI terminator to the remaining SCSI connector on each BCC.
8. Repeat the preceding steps for the remaining disk enclosures, using the next sequential SCSI channels and incrementing the Disk Enclosure (Tray) ID setting for each enclosure.



**Figure 32** Split-Bus Configuration Switch Settings




**Figure 33** Split-Bus Cabling




## Bus Addressing Examples

Each disk module within the disk array is identified by its channel number and SCSI ID. These values will differ depending on which type of bus configuration is used for the disk enclosures. See "How are disk modules in the array identified?" on page 125 for more information. Figure 34 is an example of split-bus addressing. Figure 35 is an example of full-bus addressing.



**Figure 34** Split-Bus Addressing Example



**Figure 35** Full-Bus Addressing Example

---

## Connecting the Fibre Channel Cables

Fibre Channel cables provide the I/O path to the disk array. The Fibre Channel cable connects the controller enclosure directly to a host, or to a hub.

For operation on HP-UX, the host must contain an HP Fibre Channel Mass Storage/9000 adapter. For HP-UX supported adapters, installation, and configuration information for the HP Fibre Channel adapters, refer to the *Hewlett-Packard Fibre Channel Mass Storage Adapter Service and User Manual (J3636-90002)* supplied with the adapter, or you can download this document from <http://www.docs.hp.com/>. This document also contains information on verifying the Fibre Channel connection to the disk array and information on troubleshooting the Fibre Channel loop.

For information on Fibre Channel host adapters supported on Windows NT and Windows 2000, check the Host Adapter folder on the *HP Storage Manager 60 CD*.

To connect the Fibre Channel cables to the array, complete the following steps:

---

**Note** It is advisable to select Fibre Channel cables that are as close as possible to the proper length for the installation.

---


1. Connect a fibre optic cable to a host or hub. Locate the appropriate connector on the adapter card. Remove the optical cover and connect the fibre optic cable. Refer to the appropriate system documentation for additional information.

---

**Note** When plugging the Fibre Cable connector into the adapter, push the plug into the connector until it snaps into the connector. There is a slight amount of slack in the connection which may prevent the connector from making an optimal connection. After plugging the plug into the connector, grab the cable just behind the plug and push lightly inward to ensure a good optical connection.

---

2. Plug a media interface adapter (MIA) into the Fibre Channel connectors on the controller enclosure (Figure 36, connectors J3 and J4). Make sure that the RFI gaskets are installed between the MIA and the controller connector. Tighten the two screws in each MIA.



**Figure 36** MIA, RFI Gasket, and Fibre Channel Installation

3. Connect the Fibre Channel connectors to the controller MIAs.
  - a. Remove the optical protectors from the ends of the MIAs and the Fibre Channel cables (Figure 36).
  - b. Insert the Fibre Channel connectors into the MIAs. The fibre optic cable connector is keyed to install only one way.

---

## Applying Power to the Disk Array

Once the hardware installation is complete, the disk array can be powered up. It is important that the proper sequence be followed when powering up the components of the disk array. To ensure proper operation, power should be applied to the disk enclosures first and then to the controller enclosure, or all components can be powered up simultaneously. This gives the disks time to spin up, ensuring that the disks are detected when the controller comes on line.

Once power is applied to all components, proper operation of the disk array is verified by checking the LEDs on both the controller enclosure and disk enclosure.

---

**Note**      **What happens if the disk enclosures are powered up before the controller enclosure?**

When the disk array controller enclosure is powered up, it checks for the presence of the disks in the disk enclosures. If the disk enclosures have not yet been powered on, the controller will not detect the disks. This will result in no LUNs appearing because none of the disks are visible to the controller. In this case it will be necessary to power down the disk array and power it up again in the proper sequence.

---


To power up the disk array:

---

**Note**      The following procedure assumes the disk array components will be powered up individually. All disk array components can also be powered up simultaneously using a single rack power switch.

---

1. Open all front doors on the disk enclosures.
2. Power up each disk enclosure using the disk enclosure power switch (A in Figure 37).




**Figure 37** Disk Enclosure Power Switch and System LEDs

3. Check the LEDs on the front of the disk enclosures (see Figure 39). The System Power LED (B in Figure 37) should be on and the Enclosure Fault LED (C) should be off.

It is normal for the Enclosure Fault LED (amber) to go on momentarily when the enclosure is first powered on. However, if the Enclosure Fault LED remains on, it indicates that a fault has been detected. Refer to "Troubleshooting" on page 211 for additional information.

4. Power up the controller enclosure using both power switches on the back of the enclosure (see Figure 38).



**Figure 38** Controller Enclosure Power Switches

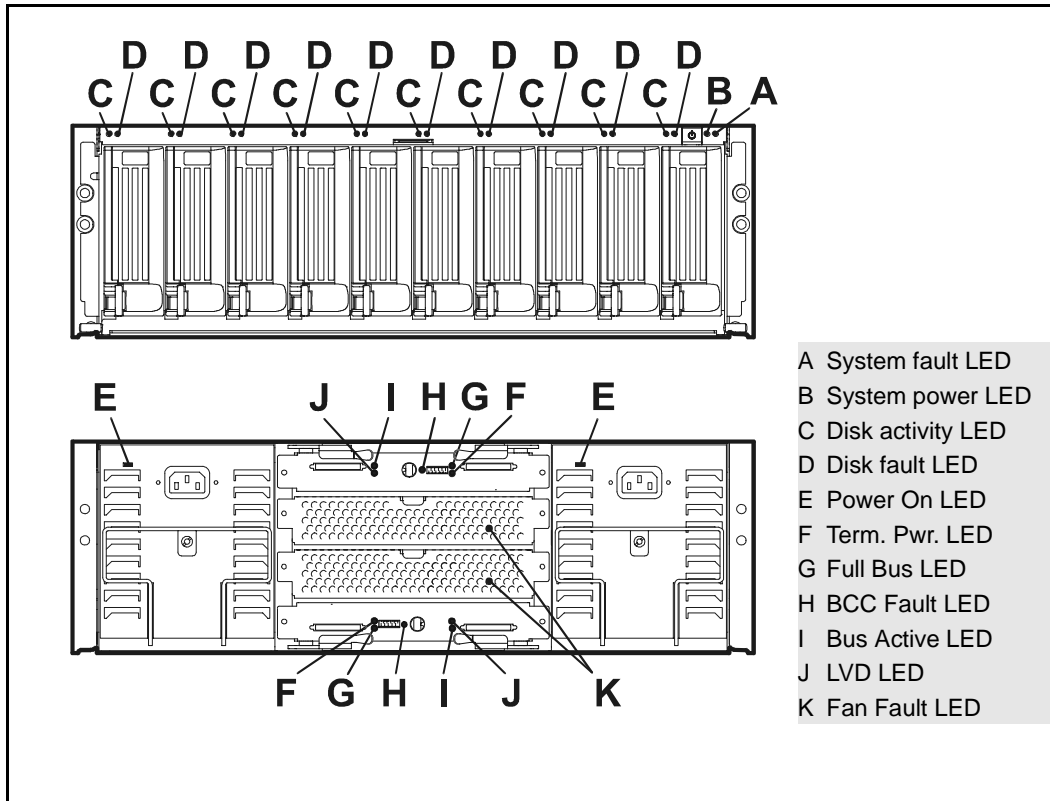
5. Check the controller enclosure LEDs (see Figure 40). The Power LED should be on and the remaining LEDs should be off. If any fault LED is on, an error has been detected. Refer to "Troubleshooting" on page 211, for additional information.
6. Close and lock the disk enclosure doors.
7. If the host was shutdown to install the array, boot the host.
8. Perform an ioscan to verify that the host sees the array.

---

**Table 12** Normal LED Status for the Disk Enclosure

Module	LED	Normal State
Front Enclosure	System Fault	Off
	System Power	On (green)
	Disk Activity	Flashing (green) when disk is being accessed.
	Disk Fault LED	Off
Power Supply	Power Supply	On (green)
BCC Module	Term. Power	On (green)
	Full Bus	Off (if split bus) On (green - if single bus)
	BCC Fault	Off
	Bus Active	On (green bus is available for use) Off (Isolator chip disabled & bus not avail.)
	LVD	On (green bus operating in LVD mode) Off (bus operating in single ended mode)
	Fan Module	Fan






**Figure 39** Disk Enclosure LEDs

Table 13 Normal LED Status for Controller Enclosure

Module	LED	Normal State
Controller Enclosure	Power On	On (green)
	Power Fault	Off
	Fan Fault	Off
	Controller Fault	Off
	Fast Write Cache	On (green) while data is in cache
Controller	Controller Power	On (green)
	Controller Fault	Off
	Heartbeat	Blink (green)
	Status	Green There are 8 status LEDs. The number and pattern of these LEDs depend on how your system is configured.
Controller Battery	Fault - B	Off
	Full Charge - B	On (green) <sup>1</sup>
	Fault - A	Off
	Full Charge - A	On (green) <sup>1</sup>
Controller Power Assembly	Power 1	On (green)
	Power 2	On (green)
Controller Fan Assembly	Fan Power	On (green)
	Fan Fault	Off
<p>1. Both Full Charge A and Full Charge B LEDs are ON after batteries are fully charged. The LEDs flash while charging is in progress, and remain on when charging is complete. Any time power to the disk array controller enclosure is cycled, the BBU will recharge for 15 minutes. During this time, the Full Charge LEDs will flash.</p>		



**Figure 40** Controller Enclosure LEDs

---

## Powering Down the Array

When powering down the disk array, the controller enclosure should be powered down before the disk enclosures.

To power down the disk array:

1. Stop all I/Os from the host to the disk array.
2. Wait for the Fast Write Cache LED to go off, indicating that all data in cache has been written to the disks.
3. Power down the controller enclosure.
4. Power down the disk enclosures.

---

**Note** If necessary, power to all components of the disk array can be shut off simultaneously using a single cabinet power switch or circuit breaker. In this situation, stop all I/O to the disk array and then wait two minutes for all cache data to be written to disk before shutting off power. The Fast Write Cache LED will remain off when all cache data is written.

---

For optimum data protection, WCE (Write Cache Enable) should be disabled on all disks in the array.

---

## Solving Common Installation Problems

**Problem.** When performing an `ioscan`, the host sees the disk array controllers but none of the disks or LUNs in the disk array.

**Solution.** This is typically caused by powering on the disk array controller enclosure before powering on the disk enclosure(s). Turn off power to all disk array components, power on the disk enclosures and wait for the disks to spin up, then power on the controller enclosure.

**Problem.** When installing a factory-assembled disk array, Array Manager 60 does not see the array. The `amdsp -R` command does not detect the array, although an `ioscan` displays the disk array LUNs.

**Solution.** This can result if any of the disk array enclosures are not powered on, or they have been disconnected or removed from the cabinet. Factory-assembled disk arrays have an entire LUN structure created on them. If any of the disk enclosures are missing when the disk array is powered on, the LUNs will appear but they will all have a size of zero bytes. Make sure all disk enclosures are present and powered on before the controller enclosure is powered on.

**Problem.** The `ioscan` results shows extra LUNs on the disk array. There are several LUN 0 entries for the disk array that I did not bind.

**Solution.** The volume set addressing used by the disk array creates virtual SCSI busses to support addressing up to 32 LUNs. These LUNs do not really exist on the disk array, but are displayed in the `ioscan` output.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

---

**Problem.** The Array Manager 60 commands are not working. When I attempt to execute a command I get a message indicating that AM60Srvr is not running.

**Solution.** If the AM60Srvr daemon has stopped, no Array Manager 60 commands can execute. AM60Srvr can be restarted using the following command:

**`/opt/hparray/bin/AM60Srvr`**

---

# 4

## MANAGING THE DISK ARRAY ON HP-UX

---

Tools for Managing the Disk Array FC60 .....	120
Installing the Array Manager 60 Software .....	122
Managing Disk Array Capacity .....	123
Adding Capacity to the Disk Array .....	135
Upgrading Controller Cache to 512 Mbytes.....	139
Managing the Disk Array Using SAM .....	141
Managing the Disk Array Using Array Manager 60.....	157
Managing the Disk Array Using STM .....	194

---

---

## Tools for Managing the Disk Array FC60

**Note** On Windows NT and Windows 2000, the disk array is managed using the *HP Storage Manager 60* software. See the *HP Storage Manager 60 User's Guide* for information on managing the disk array on Windows NT and Windows 2000.

---

There are three tools available for managing the Disk Array FC60 on HP-UX: the HP-UX System Administrator Manager (SAM), Array Manager 60, and Support Tools Manager (STM). The various management tasks and the tools that can be used to perform them are shown in Table 14.

### System Administration Manager (SAM)

The common tasks involved in managing the disk array can be performed using the HP-UX System Administration Manager, or SAM. This allows you to manage the disk using the same tool you use to manage your host system.

Not all tasks can be performed using SAM. In general, the more common tasks are available from SAM while the more advanced tasks are performed using Array Manager 60.

### Array Manager 60

Array Manager 60 comprises a set of utilities that provide complete capability for managing the disk array. In addition to the common tasks available through SAM, Array Manager 60 offers more advanced management capability. Although you will likely use SAM to do much of the disk array management, you will need to use Array Manager 60 to perform any tasks that are not supported by SAM.

### STM

STM, the Support Tools Manager, provides a third alternative for managing your disk array. STM can be used to perform many common disk array management tasks. See "HP-UX Diagnostic Tools" on page 197 for complete information on using STM.



**Table 14** Management Tools and Tasks

Task	Tool		
	SAM	Array Manager 60	STM
Checking disk array status	Yes	Yes (amdsp)	Yes
Managing LUNs	Yes	Yes (amcfg)	Yes
Managing global hot spares	Yes	Yes (ammgr)	Yes
Assigning an alias to the disk array	Yes	Yes (ammgr)	No
Managing cache memory	No	Yes (ammgr)	No
Managing the rebuild process	No	Yes (amutil)	No
Synchronizing controller date and time	No	Yes (ammgr)	No
Locating disk array components	Yes	Yes (amutil)	Yes
Performing a parity scan	No	Yes (ammgr)	Yes
Displaying disk array serial numbers	Yes	Yes (amdsp)	No
Disabling disk Write Cache Enable (WCE)	No	Yes (amutil)	No
Managing disk array logs	No	Yes (amlog)	Yes

---

## Installing the Array Manager 60 Software

The Array Manager 60 software must be installed on the host to which the disk array is connected. The software should have been installed with disk array hardware. However, if you decide to move the disk array to a different host, you will need to install the software on the new host.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

### Note

---

#### **Must Array Manager 60 be installed to use SAM?**

Yes. The Array Manager 60 software is required to manage the HP SureStore E Disk Array FC60 even when using SAM. When using SAM, the Array Manager 60 utilities are invoked to perform the management tasks.

---

## Running Array Manager 60

Array Manager 60 runs as a daemon process (AM60Srvr) and is started automatically when the system is started. If for some reason Array Manager 60 stop running, an error message will be displayed when you attempt to execute any of the commands.

If you suspect that Array Manager 60 is not running, type the following command:

```
ps -ef | grep AM60Srvr
```

AM60Srvr is the Array Manager daemon and it should be displayed as running. If AM60Srvr is not running, start it using the following command:

```
/opt/hparray/bin/AM60Srvr
```

---

## Managing Disk Array Capacity

During installation, a LUN structure is created on the disk array. This structure may meet your initial storage requirements, but at some point additional capacity may be required. This involves adding disks and binding new LUNs.

Careful LUN planning will ensure that you achieve the desired levels of data protection and performance from your disk array. Altering the characteristics of an existing LUN such as its size and RAID level requires unbinding and rebinding the LUN — a time consuming process. With careful planning you can avoid having to reconfigure the disk array when adding disk array capacity.

### Configuring LUNs

The primary task in managing disk array capacity is creating the LUN structure you need. Factors such as capacity, RAID level, and stripe segment size must be determined for each LUN. The LUN configuration you select should be based on system environment and application needs.

The process of binding LUNs is much like adding new disks to your system. Each LUN appears to the host as a single hard disk drive. However, unlike conventional hard disks, the disk array gives you the flexibility to control the size and operating characteristics of each LUN .

When binding LUNs, consider the following:

- Any size limitations imposed by HP-UX or LVM.
- The maximum numbers of LUNs that can be created on the disk array is 32. Because LUN numbering begins at 0, the highest LUN number is 31.
- Your backup strategy. If you do unattended backup to a device such as tape, you may want to avoid creating a LUN that is larger than the capacity of the media. This allows you to backup an entire LUN without changing media.

---

## Selecting Disks for a LUN

When binding a LUN, you must select the disks that will be used. The capacity of the LUN is determined by the number and capacity of the disks you select, and the RAID level.

When selecting disks for a LUN, consider the following:

- To maximize high availability, select disks in different disk enclosures or on different channels. Multiple disks in the same enclosure make a RAID 5 LUN vulnerable to an enclosure failure. A RAID 1 or 0/1 LUN can survive an enclosure failure, as long as both disks of a mirrored pair are not in the same enclosure. If you attempt to select disks in the same enclosure or on the same channel when binding a LUN, you will be warned that doing so may compromise high availability.
- Select disks of the same capacity. Binding a LUN using different size disks will result in unused capacity on the larger disks. For example, binding a 4-disk LUN that includes two 9-Gbyte disks and two 18-Gbyte disks will result in a LUN with the capacity of four 9-Gbyte disks. Only 9 Gbytes of the larger disks will be used when creating the LUN.
- Consider using more disks for RAID 5 LUNs. This increases both performance and storage efficiency. Because RAID 5 uses only one disk's worth of capacity for parity, more disks in the LUN will result in a greater percentage of the disk capacity used for data. For example, in a 4-disk RAID 5 LUN, 75% of the capacity is used for data (3 of 4). Adding another disk to the LUN will increase the percentage of disk capacity used for data to 80% (4 of 5).


## Selecting Disks for a RAID 0/1 LUN

The order in which you select disks is important when creating a RAID 0/1 LUN. The first half of the disks you select will be the primary disks, and the second half of the disks will be the disk mirrors. To maintain data availability, the disk mirrors must be in a different enclosure than the primary disks. This applies regardless of which tool you are using to bind a RAID 0/1 LUN.

For example, in Figure 41 a 4-disk RAID 0/1 LUN is being bound using one disk enclosure on channel 1, and a second disk enclosure on channel 2. The correct order for selecting disks is 1:2, 1:3, 2:2, 2:3. This selection order creates mirrored pairs of 1:2/2:2 and 1:3/2:3. This maintains high availability because the primary disks are on channel 1, and the mirror disks are on channel 2.

---

Selecting disks in the incorrect order of 1:2, 2:2, 1:3, and 2:3 results in mirrored pairs of 1:2/1:3 and 2:2/2:3. This puts the primary disk and mirror disk of each pair in the same enclosure, making the LUN vulnerable to an enclosure failure.



**Figure 41** Disk Selection for a RAID 0/1 LUN

### How are disk modules in the array identified?

When performing tasks such as binding LUNs, disk modules are identified by a numbered pair of the form “channel number:SCSI ID”. For example, 1:3 identifies the disk module on channel 1 with a SCSI ID of 3. Status information may also include the enclosure number and slot number.

- **Channel number** indicates the SCSI channel that the disk enclosure containing the disk module is connected to. There are six channels available (1 - 6). If the disk enclosure is using full-bus configuration, all disk modules within the enclosure will be on the same SCSI channel. If the disk enclosure is configured for split-bus operation, disks in the even-numbered slots will be on one channel, and disks in the odd-numbered slots will be on a different channel.
- **SCSI ID** identifies the ID assigned to the disk module. SCSI ID's are determined by the slot in which the disk module is installed in the disk enclosure. If the disk enclosure is using full-bus configuration, the even-numbered slots are assigned SCSI IDs 0 - 4, and the odd-numbered slots are assigned IDs 8 - 12. (The gap in addresses is necessary for

---

internal management of enclosure components.) If the disk enclosure is configured for split-bus operation, the both the even-numbered slots and the odd-numbered slots are assigned IDs of 0 - 4.

- When viewing status information for the disk array, you may also see the disk **enclosure number** and **slot number** displayed. These parameters identify the physical location of the disk module within the disk array.


The disk enclosure number is assigned during installation using the disk enclosure ID switches located on the rear of the enclosure. Each disk enclosure in the disk array must have a unique enclosure ID number. Disk enclosures should be numbered from the top down, with the enclosure mounted highest in the rack being 0.

The slot number indicates the physical enclosure slot in which the disk module is installed. This number is related to, but is not the same as the disk SCSI ID. Slots are assigned SCSI IDs differently depending on whether the disk enclosure is configured for full or split bus operation as described above.

Figure 42 illustrates the four components involved in disk module addressing.

### **Increasing LUN Capacity**

The capacity of an existing LUN cannot be increased dynamically. If you need a larger LUN, you must unbind the LUN and then rebind it using more or larger disks. This requires you to back up any important data on the LUN, and then restore the data once the new, larger LUN has been bound.



**Figure 42** Disk Module Addressing Parameters

---

## Assigning LUN Ownership

When a LUN is bound, you must identify which disk array controller (A or B) owns the LUN. The controller that is assigned ownership serves as the primary I/O path to the LUN. The other controller serves as the secondary or alternate path to the LUN. If there is a failure in the primary I/O path and alternate links are configured, ownership of the LUN automatically switches to the alternate path, maintaining access to all data on the LUN.

When assigning LUN ownership, consider the following:

- To ensure optimum performance, LUN ownership should be balanced between controllers. This ensures that one controller is not overloaded with I/O requests, while the other is idle.
- Controller ownership can be changed on an existing LUN without impacting its operation. If you find that there is an imbalance between I/O path loads, you can change the ownership of one or more LUNs to solve the problem.

---

### Note

#### **Does the primary path selected using LVM impact LUN ownership?**

Yes. The primary I/O path established using LVM defines the owning controller for the LUN. This may override the controller ownership defined when the LUN was bound. For example, if controller A was identified as the owning controller when the LUN was bound, and LVM subsequently established the primary path to the LUN through controller B, controller B becomes the owning controller.

---

## Selecting a RAID Level

---

**CAUTION** *RAID 0 does not provide data redundancy. It should only be used in situations where high performance is more important than data protection. The failure of any disk within a RAID 0 LUN will cause the loss of all data on the LUN.. RAID 0 should only be used for non-critical data that could be lost in the event of a hardware failure.*

---

The RAID level you select for each LUN determines the method used to store data on the disks. Once selected, the RAID level cannot be changed dynamically. If you want to change



---

the RAID level used by a LUN, you must unbind the LUN and rebind it using the new RAID level.

With the exception of RAID 0, all RAID levels supported by the disk array provide protection against disk failure. However, there are differences in performance and storage efficiency between RAID levels. The two factors that will influence the RAID level you select are performance and storage efficiency.

- **Performance** - The performance characteristics differ for each RAID level. The types of applications you are running and the I/O activity associated with them should influence the RAID level you select for a LUN. The performance characteristics for each RAID level are summarized in "RAID Level Comparisons" on page 45.
- **Storage Efficiency** - the storage efficiency can range from 50% for RAID 1 and 0/1 up to > 80% for RAID 5. The higher the efficiency, the less cost per megabyte for storing your data.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

## Global Hot Spares

Global hot spares provide additional protection against disk failures. The number of global hot spares you use will reflect how much protection you need. Each global hot spare you add will provide protection against the failure of a single disk. In addition to restoring hardware redundancy, a global hot spare also restores disk array performance that may be diminished while the disk array is operating in degraded mode with a failed disk.

When adding global hot spares, consider the following:

- It is recommended that you add one global hot spare per disk channel. This provides maximum protection against disk failure. A global hot spare can be used to replace any failed disk within the array, regardless of what channel it is on.
- When a failed disk is replaced, the data will be copied from the global hot spare disk to the replacement disk and the global hot spare will again be available as protection against another disk failure.

- 
- If you choose to limit the number of global hot spares, make sure you are able to respond quickly to replace a failed disk. If an operator is always available to replace a disk, you may not need the added protection offered by multiple global hot spares.

## Setting Stripe Segment Size

Another factor you may have to consider is the stripe segment size you use for a LUN. The stripe segment size determines how much data is written to a disk before moving to the next disk in the LUN to continue writing. For example, if the stripe segment size is set to 4 Kbytes, the disk array will write 4 Kbytes of data to disk 1, then 4 Kbytes of data to disk 2, then 4 Kbytes of data to disk 3, and so on. When creating a LUN, the default value used for the stripe segment size is the value currently set for the cache page size (4 Kbytes or 16 Kbytes). The default cache page size is 4 Kbytes.

When setting stripe segment size, consider the following:

- Stripe segment size can affect disk array performance. The smaller the stripe segment size, the more efficient the distribution of data read or written across the stripes in the LUN. However, if the stripe segment is too small for a single I/O operation, the operation requires access to two disk. Called a stripe crossing, this action reduces performance. The optimum stripe segment size is the smallest size that will rarely force I/Os to a second disk.
- The stripe segment size can be set to any multiple of the cache page size setting. For example, if the cache page size is set to 4 Kbytes, the stripe segment size can be set to 4 Kbytes, 8 Kbytes, 16 Kbytes, 32 Kbytes, etc.

---

## Evaluating Performance Impact

Several disk array configuration settings have a direct impact on I/O performance of the array. When selecting a setting, you should understand how it may affect performance. Table 15 identifies the settings that impact disk array performance and what the impact is.

---

**Note** The LUN binding process impacts disk array performance. While a LUN is being bound, benchmarking tools should not be used to evaluate performance. Wait until all LUNS are bound to get an accurate indication of disk array performance.

---

Table 15 Performance Impact of Configuration Settings

<p><b>Setting:</b> RAID level</p> <p><b>Function:</b> Sets the RAID level used by the LUN.</p> <p><b>Performance Impact:</b> The RAID level selected impacts the entire performance profile for the LUN. Read I/O and write I/O performance are directly influenced by the RAID level. See "RAID Level Comparisons" on page 45 for more information on the performance characteristics of each RAID level.</p>
<p><b>Setting:</b> Stripe segment size</p> <p><b>Function:</b> Sets the number of blocks of data the controller will write to or read from a single disk before switching to the next disk in the LUN.</p> <p><b>Performance Impact:</b> Optimum performance is typically achieved when the segment size matches the I/O size. In this case, only one disk is required to service an I/O, leaving the remaining disk in the LUN available for other I/Os.</p> <p>A large segment size provides good read performance in most RAID levels. The controller will have to access fewer disks to retrieve the data, leaving the remaining drives available for other I/O operations. Large segment sizes are typically useful for applications that require high I/O throughput.</p> <p>A small logical unit segment size is useful for most RAID 5 write applications, because the controller firmware is capable of performing group writes (writing of data simultaneously to multiple disks, while calculating the parity for the stripe, as opposed to the single-threaded read-modify-write). Small segment sizes are typically useful for applications that require large numbers of small I/Os to be processed quickly.</p>

---

Table 15 Performance Impact of Configuration Settings (cont'd)

**Setting:** Cache flush threshold (default 80%)

**Function:** Sets the level at which the disk array controller begins flushing write cache content to disk media. The setting is specified as the percentage of total available cache that can contain write data before flushing begins. The cache flush threshold can be set independently for each controller. Note that available cache is reduced by half with cache mirroring enabled. For example, if the controller has 256 Mbytes of cache and cache mirroring is enabled, only 128 Mbytes of cache is available for each controller. The remaining 128 Mbytes is being used to mirror the other controller. See Figure 43.

**Performance Impact:** This setting impacts write I/O performance by increasing or decreasing the number of write cache hits. (A "hit" is an I/O that can be serviced from cache rather than accessing the disk media.)

A higher cache threshold value increases the amount of write data that will be stored in cache before flushing to the disk is initiated. This increases the possibility that subsequent writes may be serviced from cache rather than accessing the disk, reducing the number of writes to the disk media. This improves performance for write I/Os. However, a higher value reduces the amount of cache available for reads, which may reduce the performance for read I/Os.

Lowering the cache threshold value has the opposite effect. Because less memory is used for write cache before flushing begins, the number of hits on write cache goes down with a subsequent decrease in write performance. However, read performance may improve because more memory is available for read cache.

Table 15 Performance Impact of Configuration Settings (cont'd)

**Setting:** Cache flush limit (default 100%)

**Function:** Determines how much data will remain in write cache when flushing stops. It is expressed as a percentage of the cache flush threshold. For optimum performance this value is set to 100% by default. This ensures that the entire amount of cache specified by the cache flush threshold will contain write cache data, increasing the number of write cache hits. The cache flush threshold can be set independently for each controller.

A value of 100% percent can be used for this setting because the cache flushing algorithm typically overshoots a bit in both directions. Assuming a cache flush threshold of 80% (default), the write cache contents will actually rise slightly past the threshold before flushing begins - to 82% for example. The cache flush limit of 100% (default) indicates that flushing should stop when the cache contents reaches 80%, but the flushing continues until the cache reaches a slightly lower level - 78% for example. This has the effect of maintaining a constant 80% level of data in write cache. See Figure 43.

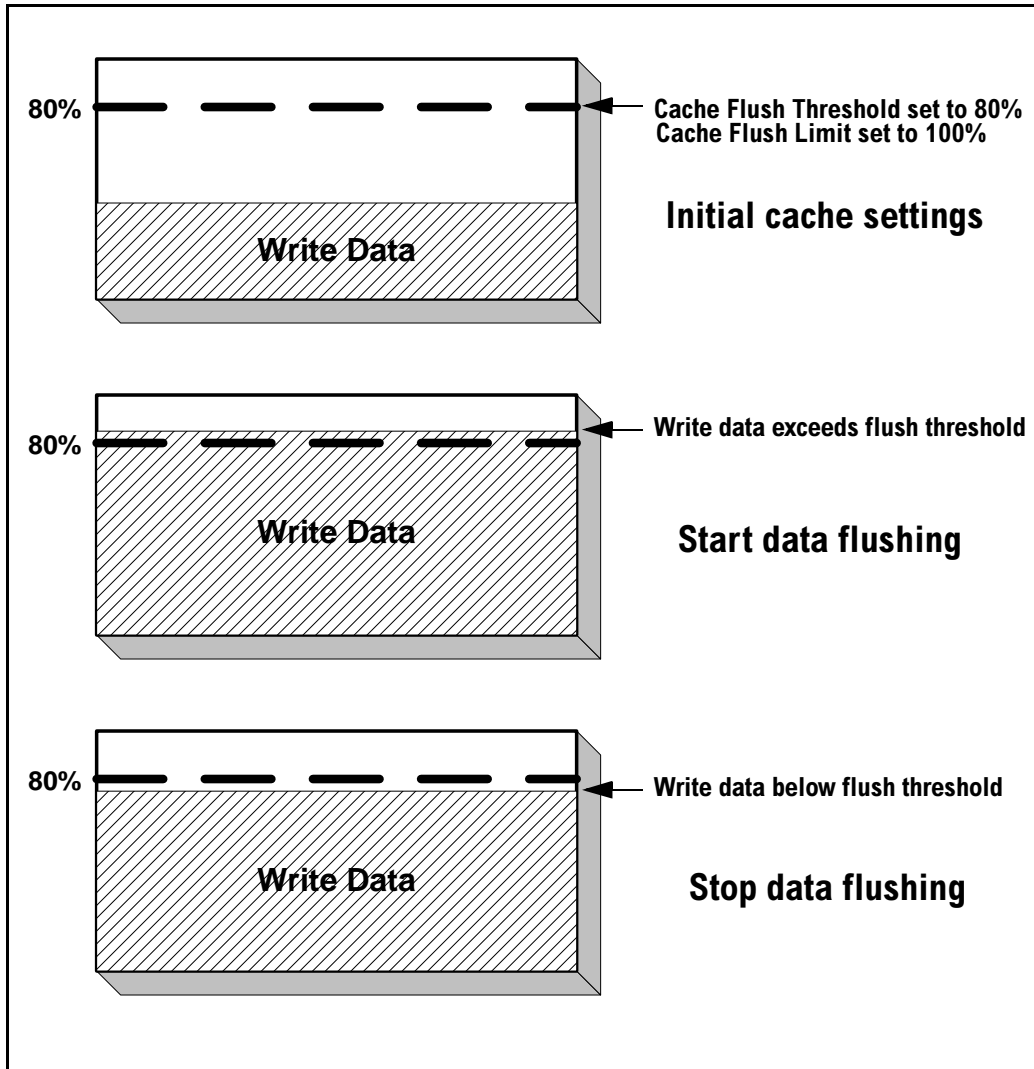
**Performance Impact:** A high value for the flush limit increases the amount of data remaining in write cache at the end of the flush. This improves the possibility of hits on write cache. Because less data is written to the disk during each flush cycle, data flushing will occur more often, increasing disk activity. In most situations this value should be left at the default value of 100%.

A lower value reduces the amount of write data remaining in cache, thereby reducing the number of write hits. More data is flushed during each cycle, reducing the number of flushes. There will be fewer cache flush operations, but each one will last longer.

**Setting:** Cache page size

**Function:** Sets the number of blocks transferred into cache by the controller at one time.

**Performance Impact:** A larger cache block size increases the probability of data near the accessed block also being available in cache. However, a larger cache block size also fills up the cache sooner. A small cache block size should be used for systems that require transaction-processing requests, or I/O streams that are typically small and random. Larger cache blocks should be used for large I/O, sequential, high bandwidth applications.



**Figure 43** Cache Flush Threshold Example

---

## Adding Capacity to the Disk Array

As your system storage requirements grow, you may need to increase the capacity of your disk array. Disk array capacity can be increased in any of the following ways:

- You can add new disk modules to the disk array if there are empty slots in the disk enclosures.
- You can add additional disk enclosures to the disk array.
- You can replace existing disk modules with higher capacity modules.

### Adding More Disk Modules

If there are empty slots in the disk enclosures, the easiest way to increase capacity is to add more disk modules and create additional LUNs. The disks you add will typically be used to bind LUNs, so make sure you add enough disk modules to create the desired structure. Disk modules can be added to the disk array while it is online without disrupting operation.

---

**CAUTION** *Ensure that there is no important data on the disk module you are installing in the disk array. Any data on the disk will be lost when it is installed.*

---

To increase capacity by adding disk modules:

1. Install the disk modules in available slots in the disk enclosures. See "Disk Enclosure Modules" on page 232 for instructions on installing disk modules.

---

**Note** Each disk in the LUN must be in a different enclosure/channel to maintain high availability.

---

- 
2. Bind a LUN with the new disks using the management tool of your choice:
    - To use SAM, see "Binding a LUN" on page 148
    - To use Array Manager 60, see "Binding a LUN" on page 170
    - To use STM, see "Binding a LUN" on page 194

---

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

---

3. Disable WCE on the disk modules using the following command:

**amutil -w <ArrayID>**

See "Disabling Disk Module Write Cache Enable (WCE)" on page 183 for more information.

## Moving Disks from One Disk Array to Another

---

**CAUTION** *Before moving disks from one array to another, ensure that there is no important data on the disks. All data on the disks will be lost when they are installed in the new disk array.*

---

If you have more than one HP SureStore E Disk Array FC60, you can move disks from one array to another to balance capacity. The disks can be installed in the new array with power on and they will be treated as new disks. Any important data on the disks should be backed up before the disks are moved.

---

**Note** **Can I move a disk module within the array with power on?**  
You should not move disks within the array with power on. Moving a disk from one slot to another will erase the data on the disk module.

---



---

## Adding Additional Disk Enclosures

Adding additional disk enclosures is another way to increase the capacity of the disk array. Up to six disk enclosures can be added to a disk array. Because it requires shutting down and possibly reconfiguring the disk array, adding new disk enclosures should be done by a trained-service representative.

An alternative to adding new disk enclosures is to move enclosures and their associated disks from one Disk Array FC60 to another. The LUNs on all disks in the enclosure should be unbound before moving it. See "Moving Disks from One Disk Array to Another" on page 136 for more details.



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

## Replacing Disk Modules with Higher Capacity Modules

If you have no available slots for adding new disk modules, you can increase the capacity of the disk array by replacing existing disk modules with higher capacity disk modules. This is typically done on a LUN-by-LUN basis and requires backing up and restoring all data on the LUN.

To increase capacity by replacing disk modules:

1. Identify the LUN impacted by the replacement of the disk modules. Make sure you replace all the disks in the LUN with the higher-capacity disks.
2. Backup all data on the LUN.
3. From the host, stop all I/O activity to the LUN unmount the file system.
4. Unbind the LUN using the management tool of your choice:
  - To use SAM, see "Unbinding a LUN" on page 152
  - To use Array Manager 60, see "Unbinding a LUN" on page 174
  - To use STM, see "Unbinding a LUN" on page 195
5. Replace the disk modules. See "Disk Enclosure Modules" on page 232 for instructions on replacing disk modules.

---

6. Bind a LUN with the new disks using the management tool of your choice:

- To use SAM, see "Binding a LUN" on page 148
- To use Array Manager 60, see "Binding a LUN" on page 170
- To use STM, see "Binding a LUN" on page 194

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

---

7. Mount the file system on the LUN and restore the data to the LUN.

8. Disable WCE on the disk modules using the following command:

**`amutil -w <ArrayID>`**

See "Disabling Disk Module Write Cache Enable (WCE)" on page 183 for more information.

---

## Upgrading Controller Cache to 512 Mbytes

Controller cache can be upgraded from the standard 256 Mbytes of cache to 512 Mbytes. This provides improved I/O performance for write operations. See Table 36 on page 262 for cache upgrade kit part numbers.

---

**CAUTION** *The cache upgrade kit must be installed by service-trained personnel only. Attempting to install the cache upgrade kit without the proper training may damage the disk array controllers.*

---

### Determining What Type of Cache Memory is Installed on the Controllers

If there is currently 256 Mbytes of cache installed on the controllers, you must determine if this memory is configured as two 128 MB DIMMs, or as a single 256 MB DIMM. This will influence the number of upgrade kits you must order.

You can determine the cache configuration by checking the Manufacturing Part Number (MFG P.N.) on the product label on the front of the controller.

1. Remove the controller enclosure front panel.
2. Check the product label on the front of the controller. Depending on the part number, the cache is configured in one of the following ways:
  - MFG P.N. A5278-62006 uses two 128 MB DIMMs. It will be necessary to order two 256 DIMMs for each controller.
  - MFG P.N. A5278-62016 uses one 256 MB DIMM. It will be necessary to order one 256 DIMM for each controller.
3. Once you know the cache configuration, refer to Table 16 to determine how many cache upgrade kits are required.

---

**Table 16** Controller Cache Upgrade Kit Selection

<b>Initial controller configuration</b>	<b>Cache Upgrade Kit(s)</b>
Dual controllers, each with two 128 MB DIMMs	Two A5279A kits
Dual controllers, each with one 256 MB DIMM	One A5279A kit
Single controller with two 128 MB DIMMs	One A5279A kit
Single controller with one 256 MB DIMM	One A5279A kit (only one of the DIMMs will be used)

---

## Managing the Disk Array Using SAM

Most of the tasks involved in everyday management of the disk array can be performed using SAM. Using SAM you can:

- Check disk array status
- Bind and unbind LUNs
- Add and remove global hot spares

---

**Note**      **Does it make any difference which controller I select when managing the disk array?**

When using SAM, you must select one of the controllers on the disk array you want to manage. Typically you can select either controller to perform an operation and the outcome will be the same. An exception is when you are binding a LUN. The controller you select will be assigned ownership of the LUN. To maintain I/O performance, you should divide the ownership of the LUNs on the disk array between the two controllers.

---

---

## Checking Disk Array Status

All aspects of disk array operation are continually monitored and the current status is stored for viewing. You can selectively view the status of any portion of the disk array configuration.

To view disk array status:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each HP FC60 disk array controller, and for each LUN on the disk array. The LUNs are listed beneath the owning controller.

Hardware Path	Description
10/0.6.0	SEAGATE ST39173WC
10/12/5.2.0	Toshiba CD-ROM SCSI Drive
8/12.8.0.0.0.0.0	HP FC60 Disk Array Controller A
8/12.8.0.0.0.0.0	HP FC60 LUN 0 (RAID 0/1 -- 6 Disks)
8/12.8.0.0.0.1.7	HP FC60 LUN 15 (RAID 0/1 -- 12 Disks)
8/12.8.0.1.0.0	HP FC60 Disk Array Controller B
8/12.8.0.1.0.2.3	HP FC60 LUN 19 (RAID 0/1 -- 8 Disks)
8/12.8.0.6.0.0	HP FC60 Disk Array Controller B
8/12.8.0.6.0.0.0	HP FC60 LUN 0 (RAID 5 -- 3 Disks)
8/12.8.0.6.0.0.3	HP FC60 LUN 3 (RAID 5 -- 5 Disks)
8/12.8.0.6.0.0.4	HP FC60 LUN 4 (RAID 1 -- 2 Disks)
8/12.8.0.6.0.0.5	HP FC60 LUN 5 (RAID 1 -- 2 Disks)
8/12.8.0.6.0.1.2	HP FC60 LUN 10 (RAID 1 -- 2 Disks)
8/12.8.0.7.0.0	HP FC60 Disk Array Controller A
8/12.8.0.7.0.0.2	HP FC60 LUN 2 (RAID 5 -- 2 Disks)
8/12.8.0.7.0.0.6	HP FC60 LUN 6 (RAID 5 -- 3 Disks)

**Note! The center portion of the screen has been removed for clarity.**

- 
3. Select a controller for the appropriate disk array from the Disk Devices list.
  4. Select the Actions menu, and the View More Information... menu option. The Main Status screen is displayed showing the overall status of the disk array.
  5. Click the appropriate button to view the detailed status for the corresponding disk array component. If you need any help in interpreting the status information, access the online help.

---

Array SN: 000200A0B805E798

Alias: opt1

Array Warnings

NO WARNINGS

General disk array status displayed here

Click a Button for Details

LUN STATUS	■
DISK MODULE STATUS	■
CONTROLLER ENCLOSURE STATUS	■
DISK ENCLOSURE 0 STATUS	■
DISK ENCLOSURE 1 STATUS	■
DISK ENCLOSURE 2 STATUS	■
DISK ENCLOSURE 3 STATUS	■
DISK ENCLOSURE 4 STATUS	■
DISK ENCLOSURE 5 STATUS	■

OK Help

Click here for detailed status of indicated component



---

## Interpreting Status Indicators

A common set of colored status indicators are used to convey the current operating status of each disk array component. The colors are interpreted as follows:

Green	The component is functioning normally. On a disk it also indicates that the disk is part of a LUN.
Red	The component has failed or is not installed.
Blue	Used only for disks, indicates the disk is functioning normally and is a member of the hot spare disk group.
White	Used only for disks, indicates the disk is functioning normally and is a member of the unassigned disk group.

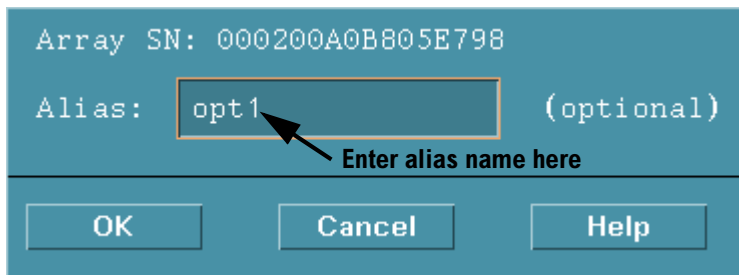
## Assigning an Alias to the Disk Array

If you have many disk arrays to manage, you may find it useful to assign an alias name to the disk arrays to help you in identifying them. The naming strategy you use may reflect the physical location of the disk array, or its function. For example, a disk array located in the data center might be assigned an alias of `data_center_1`.

To assign an alias:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.

- 
- Select the Actions menu, the Disk Array Maintenance menu option, then Modify Array Alias. The Modify Array Alias screen is displayed.



Array SN: 000200A0B805E798

Alias:  (optional)

Enter alias name here

OK Cancel Help

- Enter the name in the Alias field. An alias can contain up to 16 of the following characters: upper case letters, numbers, pound sign (#), period (.), and underscore (\_). All other characters are invalid.
- Click OK.

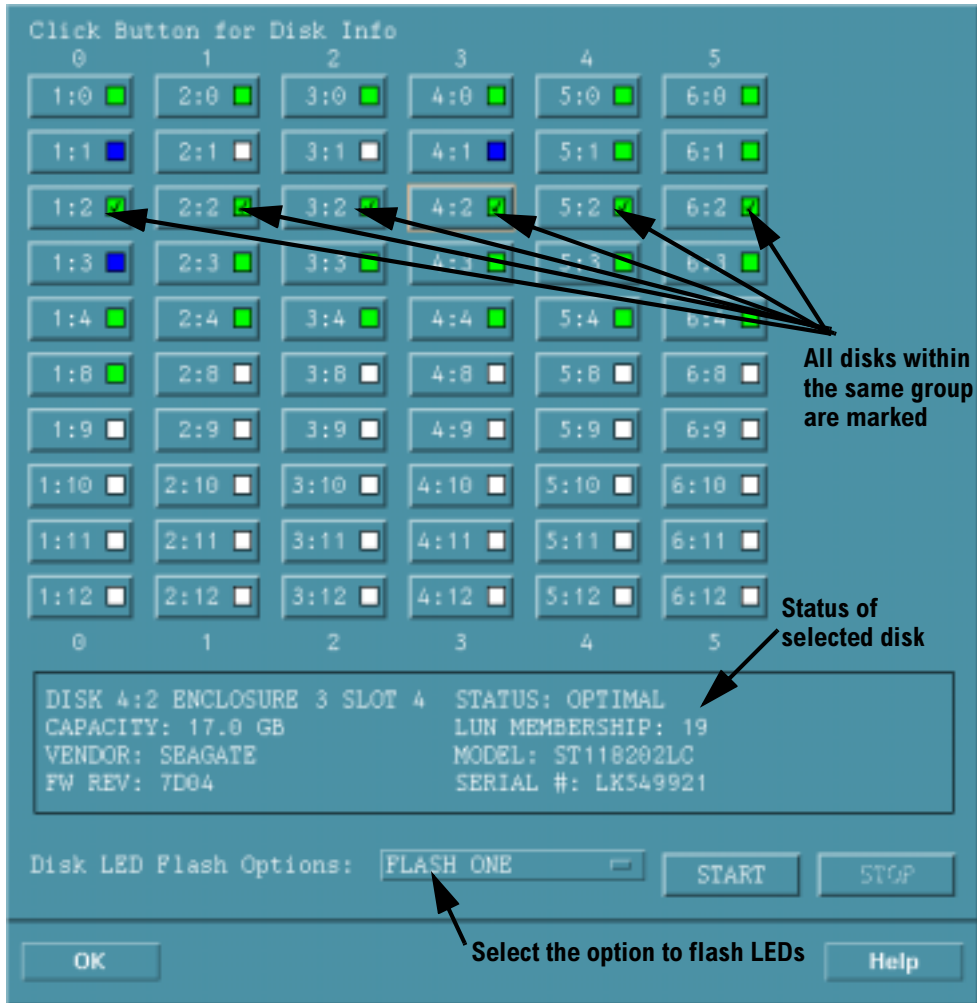
## Locating Disk Modules

To assist you in locating disk modules on the disk array, you can flash the amber Fault LED on one or multiple disks. This is useful if you want to identify the disks that make up a LUN, or a failed disk that must be replaced. It is advisable to positively identify a failed disk before removing it from the disk array. Removing the wrong disk could cause the LUN to become unavailable.

To locate a disk module:

- On the main SAM screen, double-click the Disks and File Systems icon.
- On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
- Select a controller for the appropriate disk array from the Disk Devices list.
- Select the Actions menu, and the View More Information... menu option. The Main Status screen is displayed.

5. Click the Disk Module Status button. The Disk Status window is displayed.



- 
6. Select the disk you want to identify. A check mark will appear on the selected disk and all the other disks in the same disk group. For example, if the selected disk is part of a LUN, all disks within the LUN will be checked. If the disk is a global hot spare, all global hot spares will be checked.
  7. Click on the Disk LED Flash Options button and select the desired option:
    - Flash One will flash the Fault LED on the selected disk only.
    - Flash Group will flash the Fault LEDs on all disk within the group
    - Flash Enclosure will flash the Fault LEDs on all disks in the same enclosure as the selected disk
    - Flash All flashes the Fault LEDs on all disks within the entire disk array
  8. Click Start to begin flashing the Fault LEDs
  9. Click Stop to stop flashing the Fault LEDs, or click OK to close the window and stop the flashing.

## Binding a LUN

A number of settings allow you to define the LUN configuration. Before binding a LUN, make sure you understand what each of the settings does and how it will impact LUN behavior. See "Configuring LUNs" on page 123 for more information.

The LUN binding process is designed to maintain high-availability by recommending that you avoid creating a LUN that includes more than one disk per enclosure. If you attempt to do so, you will be alerted that the resulting LUN will be vulnerable to an enclosure failure.

---

**Note****How long does it take to bind a LUN?**

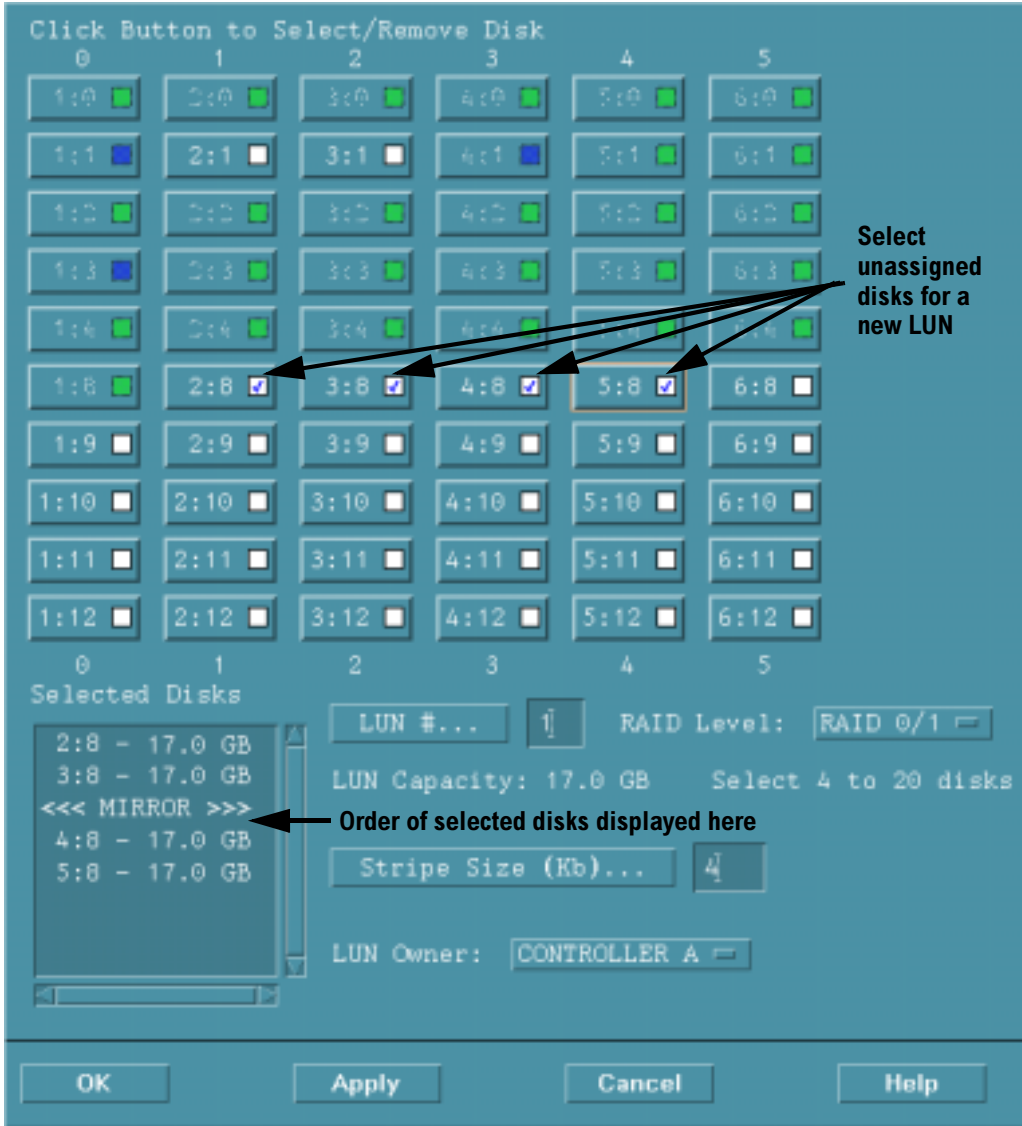
The time it takes to bind a LUN depends on the size of the LUN you are creating. The larger the LUN, the longer it takes. For example, binding a RAID 5 LUN with 18-Gbyte disks can take up to 60 minutes.

---

---

To bind a LUN:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Bind LUN. The LUN Management screen is displayed.



- 
5. Click the LUN # button and select the desired number for the LUN. You can also enter the LUN number directly in the field.
  6. Click the RAID Level button and select the desired RAID level for the LUN.
  7. Select the disks to include in the LUN. All unassigned disks are identified with a status of white.

As you select disks, the capacity of the LUN is calculated and displayed, and the disks are added to the Selected Disks field. When creating a 0/1 LUN, the Selected Disks field identifies the data disks in the upper portion of the field and the associated mirror disks in the lower portion (below the <<<MIRROR>>> separator). See "Selecting Disks for a RAID 0/1 LUN" on page 124 for more information.

8. Click the Stripe Size button and select the stripe element size. You can also enter the stripe size in the field. The stripe size must be a multiple of the current cache page size. If you enter a value that is not a multiple of the cache page size, the value will be rounded up to the nearest correct value.
9. In the LUN Owner field, select the controller that will own the LUN. LUN ownership should be balanced between controllers for optimum performance.
10. Click OK to bind the LUN and exit the screen, or click Apply if you have additional LUNs to bind.

---

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files for the LUN. This makes the LUN usable by the operating system.

---

---

## Unbinding a LUN

Unbinding a LUN makes its capacity available for the creation of a new LUN. All disks in the LUN are returned to the Unassigned disk group when the LUN is unbound.

---

**CAUTION** *All data on a LUN is lost when it is unbound. Make sure you backup any important data on the LUN before unbinding it.*

*Unbinding a LUN will have the same effect on the host as removing a disk. Before unbinding a LUN, check your operating system documentation for any additional information or steps that may be required before unbinding a LUN.*

---

To unbind a LUN:

1. On the main SAM screen, select Disks and File Systems.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller, and each LUN on the disk array. The LUNs are listed beneath the owning controller.
3. From the Disk Devices list, select the LUN to be unbound.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Unbind.
5. Select "OK" to confirm the request.

## Replacing a LUN

It is possible to replace an existing LUN. This operation unbinds the LUN, and then rebinds it using any new parameters specified. This is useful if you wish to change the settings of an existing LUN, such as its size or RAID level, using a single command.

---

**CAUTION** *All data on a LUN is lost when it is replaced. Make sure you backup any important data on the LUN before replacing it.*

---



---

**Note** **Can I replace any LUN on the disk array?**

Yes. In addition, the replace command is the only way you can alter the configuration of LUN 0.

---

LUN 0 is unique in that it must exist on the disk array to permit communication with the host. Consequently, you cannot unbind LUN 0. If you want to alter LUN 0, you must use the replace command.

To replace a LUN:

1. On the main SAM screen, select Disks and File Systems.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller, and each LUN on the disk array. The LUNs are listed beneath the owning controller.
3. From the Disk Devices list, select the LUN to be unbound.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Replace. The LUN Management screen is displayed.
5. Click the RAID Level button and select the desired RAID level for the LUN.
6. Select the disks to include in the LUN. The disks that were members of the original LUN are marked for inclusion. Remove any of these disks, or add additional disks as desired.
7. Click the Stripe Size button and select the stripe element size. You can also enter the stripe size in the field. The stripe size must be a multiple of the current cache page size. If you enter a value that is not a multiple of the cache page size, the value will be rounded up to the nearest correct value.
8. In the LUN Owner field, select the controller that will own the LUN. LUN ownership should be balanced between controllers for optimum performance.
9. Click OK to bind the new LUN and exit the screen.

---

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files on the LUN. This makes the LUN usable by the operating system.

---

---

## Adding a Global Hot Spare

Global hot spares provide an additional level of protection for the data on your disk array. A global hot spare automatically replaces a failed disk, restoring redundancy and protecting against a second disk failure. For maximum protection against disk failures it is recommended that you add a global hot spare for each channel. For more information on using global hot spares, see "Global Hot Spares" on page 129.

A global hot spare is added using an unassigned disk. If there are no unassigned disks available, you cannot add a global hot spare unless you install a new disk or unbind an existing LUN.

---

**CAUTION** *If you have disks of different capacities in your disk array, always select disks of the largest capacity for your global hot spares. This ensures that any disk failure is protected. See "" on page 49 for more information on selecting disks for global hot spares.*

*If a larger disk is used to replace a smaller disk that has failed, the difference in capacity will be unused. For example, if an 18-Gbyte disk is used as a global hot spare for a 9-Gbyte disk, the remaining 9 Gbytes of the global hot spare will be unused.*

---

To add a global hot spare:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Add Hot Spare. The Add Hot Spare screen is displayed.

Click a Button to Add a Hot Spare Disk

0	1	2	3	4	5
1:0 <input type="checkbox"/>	2:0 <input type="checkbox"/>	3:0 <input type="checkbox"/>	4:0 <input type="checkbox"/>	5:0 <input type="checkbox"/>	6:0 <input type="checkbox"/>
1:1 <input type="checkbox"/>	2:1 <input checked="" type="checkbox"/>	3:1 <input checked="" type="checkbox"/>	4:1 <input type="checkbox"/>	5:1 <input type="checkbox"/>	6:1 <input type="checkbox"/>
1:2 <input type="checkbox"/>	2:2 <input type="checkbox"/>	3:2 <input type="checkbox"/>	4:2 <input type="checkbox"/>	5:2 <input type="checkbox"/>	6:2 <input type="checkbox"/>
1:3 <input type="checkbox"/>	2:3 <input type="checkbox"/>	3:3 <input type="checkbox"/>	4:3 <input type="checkbox"/>	5:3 <input type="checkbox"/>	6:3 <input type="checkbox"/>
1:4 <input type="checkbox"/>	2:4 <input type="checkbox"/>	3:4 <input type="checkbox"/>	4:4 <input type="checkbox"/>	5:4 <input type="checkbox"/>	6:4 <input type="checkbox"/>
1:8 <input type="checkbox"/>	2:8 <input type="checkbox"/>	3:8 <input type="checkbox"/>	4:8 <input type="checkbox"/>	5:8 <input type="checkbox"/>	6:8 <input type="checkbox"/>
1:9 <input type="checkbox"/>	2:9 <input type="checkbox"/>	3:9 <input type="checkbox"/>	4:9 <input type="checkbox"/>	5:9 <input type="checkbox"/>	6:9 <input type="checkbox"/>
1:10 <input type="checkbox"/>	2:10 <input type="checkbox"/>	3:10 <input type="checkbox"/>	4:10 <input type="checkbox"/>	5:10 <input type="checkbox"/>	6:10 <input type="checkbox"/>
1:11 <input type="checkbox"/>	2:11 <input type="checkbox"/>	3:11 <input type="checkbox"/>	4:11 <input type="checkbox"/>	5:11 <input type="checkbox"/>	6:11 <input type="checkbox"/>
1:12 <input type="checkbox"/>	2:12 <input type="checkbox"/>	3:12 <input type="checkbox"/>	4:12 <input type="checkbox"/>	5:12 <input type="checkbox"/>	6:12 <input type="checkbox"/>
0	1	2	3	4	5

Unassigned disks selected as hot spares

DISK 3:1 ENCLOSURE 2 SLOT 2 STATUS: UNASSIGNED  
 CAPACITY: 17.0 GB MEMBERSHIP: UNASSIGNED  
 VENDOR: SEAGATE MODEL: ST118202LC  
 FW REV: 7D04 SERIAL #: LK556121

OK Apply Cancel Help

- 
5. Select the disk to be used as a global hot spare.

Only unassigned disks, identified by a white status indicator, are available for selection as hot spares.

6. Click OK to add the global hot spare and exit the screen, or click Apply if you want to add more global hot spares.

## Removing a Global Hot Spare

If you need to increase the available capacity of your disk array, you can do so by removing a global hot spare. The global hot spare will be added to the unassigned disk group and can be used to bind a LUN. However, the loss of the global hot spare may make your disk array more vulnerable when a disk failure occurs.

To remove a global hot spare:

1. On the main SAM screen, double-click the Disks and File Systems icon.
2. On the Disks and File Systems screen, double-click the Disk Devices icon. The Disk Devices list is displayed. There is an entry for each disk array controller.
3. Select a controller for the appropriate disk array from the Disk Devices list.
4. Select the Actions menu, the Disk Array Maintenance menu option, and then Remove Hot Spare. The Delete Hot Spare screen is displayed showing all current global hot spares, which are identified by a blue status indicator.
5. Select the global hot spare you want to remove.
6. Click OK to remove the global hot spare and exit the screen, or click Apply if you want to remove additional global hot spares.

## Managing the Disk Array Using Array Manager 60

The Array Manager 60 command line utilities allow you to configure, control, and monitor all aspects of disk array operation. Array Manager 60 is intended for performing the more advanced tasks involved in managing the disk array.

The Array Manager 60 utilities and the tasks they are used to perform are summarized in Table 17 and Table 18.

---

**Note** You must log in as superuser or root to use the Array Manager 60 utilities to manage the disk array.

---

**Table 17** Array Manager 60 Task Summary

Task	Command
Capacity Management	
Binding a LUN	amcfg -L <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID>
Unbinding a LUN	amcfg -D <LUN> [-force] <ArrayID>
Calculating LUN Capacity	amcfg -C -d channel:ID,channel:ID... -r <RAIDlevel> -s <stripe> <ArrayID>
Adding a Global Hot Spare	ammgr -h channel:ID <ArrayID>
Removing a Global Hot Spare	ammgr -d channel:ID <ArrayID>
Identifying Disks	amdsp -d <ArrayID>
Changing LUN Ownership	amcfg -M <LUN> -c <cntrlrID> <ArrayID>
Replacing a LUN	amcfg -R <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>..... -r <RAIDlevel> <options> <ArrayID>
Disk Array Status	
Displaying Status Information	amdsp <option> <ArrayID>
Listing Disk Array IDs	amdsp -i
Rescanning for Disk Arrays	amdsp -R

**Table 17** Array Manager 60 Task Summary (cont'd)

Task	Command
Disk Array Configuration	
Assigning an Alias to the Disk Array	ammgr -D <ArrayAlias> <ArrayID>
Setting Cache Page Size	ammgr -p {4   16} <ArrayID>
Setting the Cache Flush Threshold	ammgr -T <cntrlrID>:<percent> <ArrayID>
Setting the Cache Flush Limit	ammgr -L <cntrlrID>:<percent> <ArrayID>
Disabling Disk Module Write Cache Enable (WCE)	amutil -w <ArrayID>
Synchronizing the Controller Date and Time	ammgr -t <ArrayID>
Managing the Universal Transport Mechanism (UTM)	ammgr -u -U [-I <LUN>] <ArrayID>
Disk Array Maintenance	
Locating Disk Modules	amutil <option> <ArrayID>
Checking Rebuild Progress	amdsp -r <ArrayID>
Changing Rebuild Priority Settings	amutil -R <LUN> -f <Freq> -a <Amt> <ArrayID>
Performing a Parity Scan	ammgr -P <LUN>
Viewing Disk Array Logs	amlog [-s <StartTime>] [-e <EndTime>] [-d <LogDir>] [-a <ArrayID>]
Flushing Disk Array Log Contents	amutil -l <ArrayID>
Purging Controller Logs	amutil -p <ArrayID>
Resetting Battery Age	ammgr -b

**Table 18** Array Manager 60 Command Summary

Command	Tasks
amcfg	<ul style="list-style-type: none"> <li>Binding a LUN</li> <li>Unbinding a LUN</li> <li>Changing LUN Ownership</li> <li>Replacing a LUN</li> <li>Calculating LUN Capacity</li> </ul>
ammgr	<ul style="list-style-type: none"> <li>Adding a Global Hot Spare</li> <li>Removing a Global Hot Spare</li> <li>Assigning an Alias to the Disk Array</li> <li>Setting Cache Page Size</li> <li>Setting the Cache Flush Threshold</li> <li>Setting the Cache Flush Limit</li> <li>Synchronizing the Controller Date and Time</li> <li>Performing a Parity Scan</li> <li>Resetting Battery Age</li> <li>Managing the Universal Transport Mechanism (UTM)</li> </ul>
amdsp	<ul style="list-style-type: none"> <li>Identifying Disks</li> <li>Displaying Status Information</li> <li>Listing Disk Array IDs</li> <li>Rescanning for Disk Arrays</li> <li>Checking Rebuild Progress</li> </ul>
amutil	<ul style="list-style-type: none"> <li>Locating Disk Modules</li> <li>Changing Rebuild Priority Settings</li> <li>Disabling Disk Module Write Cache Enable (WCE)</li> <li>Flushing Disk Array Log Contents</li> </ul>
amlog	<ul style="list-style-type: none"> <li>Viewing Disk Array Logs</li> </ul>

---

## Command Syntax Conventions

The following symbols are used in the command descriptions and examples in this chapter.

**Table 19** Syntax Conventions

Symbol	Meaning
< >	Indicates a variable that must be entered by the user.
	Only one of the listed parameters can be used (exclusive OR).
[ ]	Values enclosed in these braces are optional.
{ }	Values enclosed in these braces are required.

## Array Manager 60 man pages

Online man pages are included for each Array Manager 60 command. The man page includes detailed information about the command and its usage.

**To access the man page for an Array Manager 60 command, type:**

```
man <command_name>
```

Substitute one of the Array Manager 60 utility names for **command\_name**. For example, to access the amdsp man page, type:

```
man amdsp
```

## Quick Help

A quick listing of the syntax and available options for a command can be displayed by using the “?” option with the command. For example, for quick information about the ammgr utility, type:

```
ammgr -?
```



## Selecting a Disk Array and Its Components

When using Array Manager 60, you must select the disk array you will be managing. In addition, many commands also require you to identify the controller, disk, or LUN within the disk array that will be impacted by the command.

The command parameters used to select a disk array and its internal components are listed and described in Table 20.

---

**Note**      **Does it make any difference which controller I select?**

There are two paths to the disk array — one for each controller. Typically you can select either controller to perform an operation and the outcome will be the same. An exception is when you are binding a LUN. The controller you select will be assigned ownership of the LUN. To maintain I/O performance, you should divide the ownership of the LUNs on the disk array between the two controllers.

---

**Table 20** Command Parameters for Identifying Disk Array Components

Command parameter	Description
ArrayID	Identifies the disk array. The ArrayID may be either the disk array ID (its S/N), or an alias name assigned to the array. Each disk array must have a unique ArrayID. The disk array S/N and alias can be determined using the <code>amdsp -i</code> command.
cntrlrID	Identifies the controller within the disk array controller subsystem. Values can be either A or B, depending on which slot the controller is installed in.
channel:ID	Identifies a disk within the disk array. The channel value indicates the SCSI channel that the disk enclosure containing the disk is connected to. The ID value indicates the SCSI ID assigned to the slot in which the disk is installed.  Valid channel values are 1 - 6, and valid SCSI ID values are 0 - 4, 8 - 12.
LUN	Identifies a specific LUN on the disk array. Valid LUN values are 0 - 31.

---

## Preparing to Manage the Disk Array

Before you begin using Array Manager 60 to manage your disk array for the first time, you may want to perform the following procedure. It will locate all the disk arrays on the host and allow you to assign an alias to each one. A short, meaningful alias should be easier to remember than the disk array ID when using the Array Manager 60 commands.

1. Rescan for all disk arrays on the host by typing:

**amdsp -R**

2. List all the disk arrays by typing:

**amdsp -i**

Each disk array will be listed along with its ID.

3. Assign an alias to each disk array by typing:

**ammgr -D <ArrayAlias> <ArrayID>**

**ArrayAlias** is the name you want to assign to the disk array. You can use up to 16 of the following characters: letters, numbers, pound sign (#), period (.), and underscore (\_). All other characters are invalid.

## Checking Disk Array Status

An important part of managing the disk array involves monitoring its status to ensure it is working properly. Changes in disk array status may indicate a possible hardware failure, so it is important to check disk array status regularly.

All aspects of disk array operation are continually monitored and the current status is stored for viewing. You can selectively view the status of any portion of the disk array configuration.



This topic is discussed in greater detail in the *Advanced User's Guide*, which can be downloaded from <http://www.hp.com/support/FC60>.

## Displaying Status Information

To display status information, type:

```
amdsp <option> <ArrayID>
```

- **option** identifies the type of status information to display. Table 21 lists the command options and the status information each displays.

**Table 21** Command Options for Displaying Disk Array Status

Option	Status Information Displayed
none	General array information. This includes disk capacity usage. Sample output is shown in Figure 44.
-l [LUN]	LUN status. Status for all LUNs on the disk array will be displayed, unless a specific LUN is specified using the LUN option. If UTM is enabled, it will be displayed in the LUN output.
-g	Disk group status.
-d	Disk status. Complete information will be displayed for each disk in the disk array.
-c	Controller status. Information is displayed for both disk array controllers and for all disk system BCC modules.
-s	Array status. Information about the entire array is displayed.
-h	Hardware status. Information is displayed for all disk array hardware components, including fans and power supplies.

---

**Table 21** Command Options for Displaying Disk Array Status (cont'd)

Option	Status Information Displayed
-a	All status. This option displays all the information returned by the preceding options.
-p <devicefile>	Hardware path information. Displays hardware path information for the controller corresponding to the specified device file.
-r	Rebuild status. Display the progress of any rebuilds occurring on the disk array.
-A	Array server status.

### Command Example

The following example displays hardware information for disk array Array1.

**amdsp Array1**

A sample out put from this command is shown in Figure 44. The important fields are identified with the normal values.

**Figure 44** Disk Array Sample Status Output (amdsp)

```

Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
-----
Array State         = READY
Server name         = speedy
Array type          = 3
Mfg. Product Code  = 348-0040789

--- Disk space usage -----
Total physical      = 271.4 GB
Allocated to LUNs  = 135.4 GB
Used as Hot spare   = 0.0 GB
Unallocated (avail for LUNs) = 0.0 GB
-----

Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
-----
LUN      Status      Capacity  Ctrl  RAID  Segment  Disks
-----
  0  OPTIMAL          16.9 GB   A     1      4      1:0
                                     3:0
  1  OPTIMAL          16.9 GB   B     1      4      2:0
                                     4:0
  .
  .
31  UTM:GOOD
-----
LUN  WCE  RCD  CME  CWOB      WCA  RCA  CMA
---  ---  ---  ---  ---      ---  ---  ---
  0   X   X   X   X          X   X   X
  1   X   X   X   X          X   X   X
  .
  .
Total capacity of LUNs on controller A = 84.6 GB
Total capacity of LUNs on controller B = 50.8 GB
Total capacity of all configured LUNs = 135.4 GB

```

**LUN Information - the status of each LUN should be OPTIMAL**

**LUN 31 is the default UTM LUN. GOOD status indicates the UTM is enabled and operating normally.**

**Cache Status for each LUN - all LUNs should have the status values shown**

**LUNs should typically be distributed between both controllers for best performance. LUN ownership can be reassigned to balance LUN usage.**

---

```
Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
-----
```

```
SCSI Channel:ID    = 1:0
Enclosure           = 0
Slot                = 0
Disk State          = OPTIMAL
Disk Group and Type = 060E86000238C6360F LUN
Capacity            = 17.0 GB
Manufacturer and Model = SEAGATE ST318203LC
Serial Number       = LRB61150
Firmware Revision   = HP01
.
.
.
```

**Disk Information - the status of each disk should be Optimal for each disk assigned to a LUN.**

Total capacity of all installed physical disks = 271.4 GB

Hot Spare Activity

-----  
None

```
Vendor ID           = HP
Product ID          = A5277A
Array ID            = 000A00A0B80673A6
Array alias         = Array1
-----
```

Disk Group 060E86000738C6395B:

```
Disk Group Type     = LUN
Number of LUNs      = 1
LUN(s)              = 6
Remaining Capacity  = 0.0 MB
RAID Level           = 1
Segment Size        = 4 KB
```

**Disk Group Information - this is another representation of the LUN information**

```
Disks:
1:3  3:3
.
.
.
```

Total remaining capacity for LUN disk groups = 0.0 MB

Information for Controller A - 000A00A0B80673A6:

```

Controller Status = GOOD
Controller Mode   = ACTIVE
Vendor ID        = HP
Product ID       = A5277A
Serial Number    = 1T00310110
Firmware Revision = 03010108
Boot Revision    = 03010002
HP Revision      = HP01
Loop ID         = 5
AL_PA           = 0xE0
Preferred AL_PA  = 0xE0
Controller Date  = 05/08/2000
Controller Time  = 13:38:53
Cache Battery Age = 0 to 89 days
NVSRAM Checksum = 0x353181DF
    
```

**Controller Information - make sure the following conditions are met:**

- Both controllers should be ACTIVE
- The Loop ID must be unique for each controller
- The three levels of firmware revisions must be identical for each controller

Information for Controller B - 000A00A0E80673A6:

```

Controller Status = GOOD
Controller Mode   = ACTIVE
Vendor ID        = HP
Product ID       = A5277A
Serial Number    = 1T90510185
Firmware Revision = 03010108
Boot Revision    = 03010002
HP Revision      = HP01
Loop ID         = 4
AL_PA           = 0xE1
Preferred AL_PA  = 0xE1
Controller Date  = 05/08/2000
Controller Time  = 13:38:54
Cache Battery Age = 0 to 89 days
NVSRAM Checksum = 0x353181DF
    
```

Information for Disk System 1 (USSA02010592), Controller B:

```

SCSI Channel      = 2
Thumbwheel Setting = 0
Controller Status = GOOD
Vendor ID        = HP
Product ID       = A5294A
Serial Number    = USSA02010649
Firmware Revision = HP04
    
```

**Disk Enclosure Information - make sure the following conditions are met:**

- The status of both controllers (BCCs) should be GOOD
- The Thumbwheel Setting must be the same for each controller

Information for Disk System 1 (USSA02010592), Controller A:

```

SCSI Channel      = 1
Thumbwheel Setting = 0
Controller Status = GOOD
Vendor ID        = HP
Product ID       = A5294A
Serial Number    = USSA02010592
    
```

- The Thumbwheel Setting should correspond to the enclosure position in the rack. Uppermost enclosure set to 0, next one down set to 1, etc.

---

```
Vendor ID          = HP
Product ID         = A5277A
Array ID           = 000A00A0B80673A6
Array alias        = Array1
```

```
-----
Overall State of Array      = READY
Array configuration:
  Cache Block Size         = 4 KB / 4 KB
  Cache Flush Threshold    = 80 % / 80 %
  Cache Flush Limit        = 100 % / 100 %
  Cache Size                = 256MB / 256MB
```

**Cache settings**

```
Vendor ID          = HP
Product ID         = A5277A
Array ID           = 000A00A0B80673A6
Array alias        = Array1
```

```
-----
Array Controller Subsystem:
```

```
  Controller A: GOOD
  Controller B: GOOD
    PS 1: GOOD
    PS 2: GOOD
  Fan 1: GOOD
  Fan 2: GOOD
  Temp Sensor: GOOD
  Battery: GOOD
```

```
Disk System 1 - USSA02010592:
```

```
BCC Controller B: GOOD
BCC Controller A: GOOD
  PS 1: GOOD
  PS 2: GOOD
  Fan 1: GOOD
  Fan 2: GOOD
  Temp Sensor: GOOD
```

```
Disk System 2 - USSA02010595:
```

```
BCC Controller B: GOOD
BCC Controller A: GOOD
  PS 1: GOOD
  PS 2: GOOD
  Fan 1: GOOD
  Fan 2: GOOD
  Temp Sensor: GOOD
```

```
No LUNs are currently rebuilding
```

**Component status - all hardware components should be GOOD**



---

## Listing Disk Array IDs

You may find it useful to list the disk arrays recognized by the host. The list will include both the disk array ID (or S/N) and alias name if one has been assigned. This is a quick way to determine the ID of each disk array on your system.

---

**Note**      **What if a disk array is not listed?**

If the list does not reflect the current disk arrays on your system, rescan for disk array as described in the next procedure. This will update the disk array list.

---

To list all the disk arrays, type:

```
amdsp -i
```

---

**Note**      To avoid having to use the array ID when managing the disk array, you may want to create an alias. You can also export the array ID to a variable and use the variable. For example, to export the ID for array 000000057D22 to a variable named `id`, type the following:

```
# export <id>=000000057D22
```

You can then use `$id` in place of `<array-id>`

---

## Rescanning for Disk Arrays

If the configuration of the disk arrays on your system has changed, you can rescan the disk arrays to update the information. This command causes Array Manager 60 to rescan the system for new or modified disk arrays connected to the host.

To rescan for disk arrays, type:

```
amdsp -R
```

When the rescan is complete, the command prompt will return. The command does not return a message indicating the operation is complete.

---

## Managing LUNs

Using Array Manager 60 you can perform the following tasks:

- Binding and unbinding LUNs
- Calculating LUN capacity
- Changing LUN ownership
- Replacing a LUN

### Binding a LUN

Binding LUNs is one of the most common tasks you will perform in managing the disk array. A number of settings allow you to define the LUN configuration. Before binding a LUN, make sure you understand what each of the settings does and how it will impact LUN behavior. See "Configuring LUNs" on page 123 for factors that may influence how you bind a LUN.

Before binding a LUN, you may want to calculate how much capacity the LUN will provide with the disks and settings you intend to use. Refer to "Calculating LUN Capacity" on page 173 for instructions on performing this task.

---

#### Note

##### How long does it take to bind a LUN?

The time it takes to bind a LUN depends on the size of the LUN you are creating. The larger the LUN, the longer it takes. For example, binding a RAID 5 LUN with 18-Gbyte disks can take up to 60 minutes.

---

To bind a LUN, type:

```
amcfg -L <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>,<channel:ID>.....  
-r <RAIDlevel> <options> <ArrayID>
```

- **cntrlrID** - identifies the controller (A or B) that will assume ownership of the LUN.
- **LUN** - number assigned to the LUN. It must not conflict with an existing LUN number.
- **channel:ID** - identifies the channel and SCSI ID of each disk included in the LUN. Only unassigned disks can be used for the LUN. Multiple disks must be specified with no spaces between each disk. See "Selecting Disks for a RAID 0/1 LUN" on page 124 for more information if you are binding a RAID 0/1 LUN.

- **RAIDlevel** - RAID level used for the LUN. Valid RAID levels are 0, 1, and 5. RAID 0 support requires firmware version HP08 or later. A RAID 0/1 LUN is created by selecting RAID 1 with more than two disks.
- **<options>** - options giving you control over how the LUN is configured. Table 22 lists valid options and what they do.

---

**Note** After binding a LUN, you must execute the `insf -e` command to install special device files for the LUN. This makes the LUN usable by the operating system.

---

**Table 22** Command Options for Binding a LUN

Option	Description	Default
-c capacity	<p>Specifies the capacity of the LUN. Capacity can be specified in megabytes (M) or gigabytes (G) by appending the appropriate letter to the capacity value.</p> <p>By default, the LUN will use the entire capacity available from all the disks. It is possible to specify a lower value for LUN capacity, but this will result in unused disk capacity. Any capacity not included in the LUN will be inaccessible and is essentially wasted capacity.</p>	<p>If not specified, the LUN capacity will equal the total available capacity of the specified disks.</p> <p>If a capacity value is specified but no letter is appended, gigabytes will be assumed.</p>

**Table 22** Command Options for Binding a LUN (cont'd)

Option	Description	Default
-force	Allows you to bind a LUN using two or more disks in the same enclosure or on the same channel. This option allows you to override the high-availability protection designed into the LUN binding process. Using this option you can specify more than one disk per enclosure or channel. You can also use this option to create a RAID 5 LUN that includes more than six disks.	If not specified, you cannot bind a LUN using multiple disks in the same enclosure. You also cannot bind a RAID 5 LUN with more than six disks.
-s segmentsize	Indicates the stripe segment size in Kbytes. The stripe segment size must be a multiple of the current cache page size setting. For example, if the cache page size is set to 4 Kbytes (8 blocks), valid settings would be 4, 8, 16, 32, 64, 128, etc.	If not specified, the stripe segment size defaults to the same block size used for the cache pages (4 Kbytes or 16 Kbytes). The default cache page size is 4 Kbytes.

### Command Examples

The following example binds a 5-disk, RAID 5 LUN on disk array 0000005EBD20. The disk array includes five disk enclosures, each on its own channel. The LUN uses a stripe segment size of 16 Kbytes, is assigned LUN number 2, and is owned by controller A. Note that each disk is in a different enclosure for high availability, and that there are no spaces between the individual disk parameters.

```
amcfg -L A:2 -d 1:1,2:2,3:1,4:4,5:3 -r 5 -s 16 0000005EBD20
```

The following example binds a RAID 0/1 LUN on disk array rack\_51. The LUN is owned by controller B, is assigned LUN number 4, and uses a stripe segment size of 4 Kbytes. Note that the disks selected create mirrored pairs that are in separate enclosures (1:3, 2:3 and 1:4, 2:4). This maintains high availability. Although RAID 1 is specified, the inclusion of more than two disks causes the disk array to create a RAID 0/1 LUN.

```
amcfg -L B:4 -d 1:3,1:4,2:3,2:4 -r 1 -s 4 rack_51
```

---

## Identifying Disks

Binding a LUN requires the use of unassigned disks. If you are not sure which disks are unassigned, you can determine which disks are available.

To identify unassigned disks, type

```
amdsp -d <ArrayID>
```

The status of all disks in the array will be returned. The information includes the disk group the disk is a member of. Disks in the Unassigned disk group can be used for binding a LUN. You may want to print the information or write down the unassigned disks before you begin binding the LUN.

## Calculating LUN Capacity

Before actually binding a LUN, you may want to calculate how much capacity will be available for a given set of configuration values. This command does not bind a LUN; it simply returns the capacity that would be available if you had bound a LUN using the values provided. You can alter the settings to see how LUN capacity is affected until you get the desired results. Using this command before you bind a LUN will ensure that you achieve the correct capacity.

To calculate the capacity of a LUN, type:

```
amcfg -C -d channel:ID,channel:ID... -r <RAIDlevel>  
-s <stripe> <ArrayID>
```

- **RAIDlevel** - the RAID level used for the LUN. Valid RAID levels are 1, 0/1, and 5.
- **stripe** - the size of the stripe element in Kbytes.

## Command Example

The following example calculates the capacity of a RAID 5 LUN using four disks and a stripe element size of 16 Kbytes.

```
amcfg -C -d 1:1,2:3,3:1,4:2 -r 5 -s 16 0000005EBD20
```

---

## Unbinding a LUN

Unbinding a LUN makes its capacity available for the creation of a new LUN. All disks assigned to the LUN are returned to the Unassigned disk group when the LUN is unbound.

---

**CAUTION** *All data on a LUN is lost when it is unbound. Make sure you backup any important data on the LUN before unbinding it.*

*Unbinding a LUN will have the same impact on the host as removing a disk. Before unbinding a LUN, check your operating system documentation for any additional information or steps that may be required after unbinding a LUN.*

---

To unbind a LUN, type:

```
amcfg -D <LUN> [-force] <ArrayID>
```

- **LUN** - number assigned to the LUN you want to unbind
- The **-force** option is required to delete LUN 0. LUN 0 may only be deleted if a UTM LUN has been enabled via NVSRAM setting. If UTM is disabled, LUN 0 is required for the SCSI command path. The UTM LUN itself may not be deleted using this command.

### Command Example

The following example unbinds LUN 5 on disk array rack\_1.

```
amcfg -D 5 rack_1
```

### Changing LUN Ownership

When binding a LUN, one of the disk array controllers is identified as the owner of the LUN. It is possible to change the ownership of a bound LUN dynamically - that is, without unbinding and rebinding the LUN. You may want to change the ownership of a LUN if you notice an imbalance of I/O loading between controllers. Controller ownership will change automatically if a failure occurs in the primary I/O path to the LUN.

---

**Note****Does the primary path selected using LVM impact LUN ownership?**

Yes. The primary path established using LVM defines the owning controller for the LUN. This may override the controller ownership defined when the LUN was bound. For example, if controller A was identified as the owning controller when the LUN was bound, and LVM subsequently established the primary path to the LUN through controller B, controller B becomes the owning controller.

---

To change the controller ownership of a LUN, type:

```
amcfg -M <LUN> -c <cntrlrID> <ArrayID>
```

- **cntrlrID** identifies the controller (A or B) that is being assigned LUN ownership

**Command Example**

The following example changes the ownership of LUN 3 to controller B on disk array 0000005EBD20.

```
amcfg -M 3 -c B 0000005EBD20
```

**Replacing a LUN**

It is possible to replace an existing LUN. This operation unbinds the LUN and then rebinds it using the new parameters specified. This is useful if you wish to change the settings of an existing LUN, such as its size or RAID level, using a single command.

**CAUTION**

---

*All data on a LUN is lost when it is replaced. Make sure you backup any important data on the LUN before replacing it.*

---

**Note****Can I replace any LUN on the disk array?**

Although you can replace any LUN, the replace command is intended to be used with LUN 0. LUN 0 is unique in that it must exist on the disk array to permit communication with the host. Consequently, you cannot unbind LUN 0. If you want to alter LUN 0, you must use the replace command. For all other LUNs, it is recommended that you use the unbind and bind commands to alter LUN configuration.

---

---

To replace a LUN, type:

```
amcfg -R <cntrlrID>:<LUN> -d <channel:ID>,<channel:ID>.....  
-r <RAIDlevel> <options> <ArrayID>
```

- The parameters and options available when replacing a LUN are the same as those used when binding a LUN. See "Binding a LUN" on page 170.

### Command Examples

The following example replaces existing LUN 0 on disk array 0000005EBD20. The new LUN is RAID 5, uses a stripe segment size of 16 Kbytes, and is owned by controller A. Note that each disk is in a different enclosure for high availability, and that there are no spaces between the individual disk parameters.

```
amcfg -R A:0 -d 1:1,2:2,3:1,4:4,5:3 -r 5 -s 16 0000005EBD20
```

## Managing Global Hot Spares

Global hot spares provide an additional level of protection for the data on your disk array. A global hot spare automatically replaces a failed disk, restoring data redundancy and performance that may suffer while in degraded mode. For maximum protection against disk failure it is recommended that you add one global hot spare per channel.

For more information on using global hot spares, see "Global Hot Spares" on page 129.

---

**CAUTION** *If you have disks of different capacities in your disk array, always select disks of the largest capacity for your global hot spares. This ensures that any disk failure is protected. See "" on page 49 for more information on selecting disks for global hot spares.*

*If a larger disk is used to replace a smaller disk that has failed, the difference in capacity will be unused. For example, if an 18-Gbyte disk is used as a global hot spare for a 9-Gbyte disk, the remaining 9 GB of the global hot spare remain unused.*

---



---

## Adding a Global Hot Spare

A global hot spare is added using an unassigned disk. If there are no unassigned disks available, you cannot add a global hot spare unless you install a new disk or unbind an existing LUN.

To add a global hot spare, type:

```
ammgr -h channel:ID <ArrayID>
```

### Command Example

The following example adds a global hot spare using disk 1:1 on disk array 0000005EBD20.

```
ammgr -h 1:1 0000005EBD20
```

## Removing a Global Hot Spare

If you need to increase the available capacity of your disk array, you can do so by removing a global hot spare. The global hot spare will be added to the unassigned disk group and can be used to bind a LUN. However, the loss of the global hot spare may make your disk array more vulnerable when a disk failure occurs.

To remove a global hot spare, type:

```
ammgr -d channel:ID <ArrayID>
```

### Command Example

The following example removes disk 2:3 as a global hot spare on disk array rack\_51.

```
ammgr -d 2:3 rack_51
```

---

## Managing Disk Array Configuration

### Assigning an Alias to the Disk Array

If you have many disk arrays to manage, you may find it useful to assign an alias name to each disk array to help you identify them. A short, meaningful alias should be easier to remember than the disk array ID when using the Array Manager 60 commands. The naming strategy you use may reflect the physical location of the disk array or its function. Once you have assigned an alias to the disk array, you can use the new name in the <ArrayID> field used to identify the disk array when using a command.

To assign an alias to the disk array, type:

```
ammgr -D <ArrayAlias> <ArrayID>
```

- **ArrayAlias** is the name you want to assign to the disk array. An alias can contain up to 16 of the following characters: letters, numbers, pound sign (#), period (.), and underscore (\_). All other characters are invalid.

### Command Example

The following example assigns an alias of rack\_1 to disk array 0000005EBD20:

```
ammgr -D rack_1 0000005EBD20
```

### Synchronizing the Controller Date and Time

During installation, the date and time maintained internally by the disk array should have been synchronized to the host. This ensures that any information created by the disk array such as log entries reflect the proper time they occurred. If necessary, the disk array date and time can be resynchronized with the host.

To synchronize the controller date and time with the host, type:

```
ammgr -t <ArrayID>
```

---

## Managing the Universal Transport Mechanism (UTM)

On firmware HP08 and later, the Universal Transport Mechanism (UTM) serves as the SCSI communication path between the host and the disk array. In earlier versions of firmware, this communication was done using LUN 0. The UTM is configured as a separate LUN, which is used only for communication and not for storing data. Because it consumes one of the available LUNs, only 31 LUNs are available when using the UTM. Any attempt to perform an I/O to the UTM LUN will be rejected.

The UTM is enabled by default and configured as LUN 31. It is possible to change the LUN used for the UTM to something other than 31. When upgrading to HP08 firmware on a disk array that already has LUN 31 defined, it will be necessary to copy the data a new LUN and delete LUN 31 before downloading the HP08 firmware.

One of the primary benefits of using the UTM is Major Event Logging (MEL). If the UTM is disabled, major event logging will not be performed. Instead, the standard controller logs will be used to gather disk array status. Because of the benefits it offers, it is recommended that the UTM be enabled.

### Enabling and Changing the UTM LUN

LUN 31 is assigned as the default LUN used for the UTM. In most cases this should be acceptable. It is possible to change the LUN used for the UTM if you wish to do so. This command is also used to enable the UTM if it is disabled.

---

**Note**      LUN 0 cannot be used for the UTM.

---

To change the LUN number used for UTM, or to enable the UTM, type:

```
ammgr -U -l <LUN> <ArrayID>
```

- **LUN** identifies the default LUN that will be used for the ATM. This value is stored in controller NVSRAM. If **LUN** is not specified, the default value currently stored in NVSRAM will be used for the ATM. The UTM will be enabled if it is currently disabled.

---

**Note** After executing the above command, the disk array controllers must be manually reset or power cycled before the new setting will be invoked. When the power on completes, execute the following commands:

```
ioscan  
insf -e  
amdsp -R
```

---

## Disabling the UTM

Although it possible to disable the UTM, it is not recommended that you do so. The benefits provided by the UTM, such as major event logging, are not realized when the UTM is disabled. If you want to disable major event logging, you can do so by disable the UTM.

---

**Note** It is necessary to disable the UTM LUN before converting from firmware HP08 to ealier versions of firmware.

---

To disable use of the UTM, type:

```
ammgr -u <ArrayID>
```

---

**Note** After executing the above command, the disk array controllers must be manually reset or power cycled before the new setting will be invoked. When the power on completes, execute the following commands:

```
ioscan  
insf -e  
amdsp -R
```

---

## Managing Cache Parameters

Several cache parameters control how the disk array manages controller cache memory. These setting have a direct influence on I/O performance. Before altering these settings,

---

see Table 15 on page 131 for details on what performance impact altering these settings may have.

### Setting Cache Page Size

Data caching improves disk array performance by storing data temporarily in cache memory. The cache page size can be set to either 4 Kbytes or 16 Kbytes. The default page size is 4 Kbytes. The page size is set for both controllers in the disk array.

To set cache page size, type:

```
ammgr -p {4|16} <ArrayID>
```

#### Command Example

The following example sets the cache page size to 16 Kbytes on disk array rack\_1:

```
ammgr -p 16 rack_1
```

### Setting the Cache Flush Threshold

The cache flush threshold sets the level at which the disk array begins flushing (writing) cache content to the disks. Cache will be flushed until the cache flush limit value is reached. The default value for this setting is 80%. The cache flush threshold can be set independently for each controller.

To set cache flush threshold, type:

```
ammgr -T <cntrlrID>:<percent> <ArrayID>
```

- **percent** indicates the percentage of cache memory that must be full before flushing begins. For example, if set to 80%, cache flushing will begin when 80% of cache memory is full.

#### Command Example

The following example sets the cache flush threshold to 80% for controller A on disk array 0000005EBD20:

```
ammgr -T A:80 0000005EBD20
```

---

## Setting the Cache Flush Limit

Sets the amount of unwritten data to remain in cache after a flush is completed on the given controller. The cache flush limit sets the level at which the disk array stops flushing cache contents to the disks. This value is expressed as a percentage of the current cache flush threshold. The default value for this setting is 100%. The cache flush limit can be set independently for each controller.

To set cache flush limit, type:

```
ammgr -L <cntrlrID>:<percent> <ArrayID>
```

- **percent** indicates the percentage of cache memory that remains full when flushing of cache contents stops. For example, if set to 90%, cache memory will be flushed until 90% of the cache amount defined by the cache flush threshold remains full.

### Command Example

The following example sets the cache flush limit to 100% for controller A on disk array 0000005EBD20:

```
ammgr -L A:100 0000005EBD20
```

---

## Disabling Disk Module Write Cache Enable (WCE)

**Note** To ensure optimum protection against data loss, it is recommended that Write Cache Enable be disabled on all disks in the array. Disabling disk WCE will impact disk array performance. However, it reduces the potential for data loss during a power loss.

---

To ensure optimum data integrity, it is recommended that the Write Cache Enable (WCE) feature be disabled on all disk modules in the array. Each disk has internal write cache memory which is used to store data until it is written to the disk media. (This should not be confused with the array controller write cache, which is protected by the BBU.) If power to a disk is interrupted before data in write cache has been copied to the disk media, the data will be lost. To avoid this situation, the WCE feature should be disabled on all disks in the array. This ensures that data is written directly to the disk media and not stored in disk cache.

This procedure for disabling disk WCE should be executed any time a disk module is replaced or a new disk module is added to the array. This will ensure that WCE is disabled on all disks in the array.

To disable WCE on all disks, type:

```
amutil -w <ArrayID>
```

WCE will be disabled on all disks visible to the array controllers. Depending on the number of disk modules in the array, this command can take up to 60 seconds to complete.

### Command Example

The following example disables WCE on all disks in disk array 0000005EBD20:

```
amutil -w 0000005EBD20
```

---

## Enabling Disk Write Cache Enable (WCE)

---

**CAUTION** *WCE should only be enabled in environments that provide uninterruptible power to the disk array. A loss of power to the disk array may result in data loss with WCE enabled.*

---

If maximum I/O performance is critical, disk WCE can be enabled on all the disks in the array. Disk WCE enhances disk array I/O performance, but increases the possibility of data loss during a power loss.

If WCE is enabled, this procedure for enabling disk WCE should be executed any time a disk module is replaced or a new disk module is added to the array. This will ensure that WCE is enabled on all disks in the array.

To enable WCE on all disks, type:

```
amutil -w on <ArrayID>
```

WCE will be enabled on all disks visible to the array controllers. Depending on the number of disk modules in the array, this command can take up to 60 seconds to complete.

### Command Example

The following example enables WCE on all disks in disk array ARRAY1:

```
amutil -w on ARRAY1
```



## Performing Disk Array Maintenance

At some point during operation of the disk array, you may need to perform maintenance tasks that involve error recovery and problem isolation. This section describes the tasks involved in maintaining the disk array.

### Locating Disk Modules

Array Manager 60 provides the means of identifying a disk module by flashing its amber Fault LED. You can flash the Fault LED on an individual disk, or on all the disks in the array. It is advisable to positively identify a failed disk before removing it from the disk array. Removing the wrong disk could cause the data on the LUN to become unavailable.

To locate a disk module, type:

```
amutil <option> <ArrayID>
```

- **option** identifies the disk modules you want to identify. Table 23 lists valid options.

**Table 23** Command Options for Locating Disk Array Components

Option	Description
-f channel:ID	Flash the Fault LED on the specified disk. A disk is identified by both channel (1- 6) and SCSI ID (0 - 4, 8 - 12).
-F	Flash the Fault LEDs on all the disks in the disk array.
-s	Stop flashing the disk Fault LEDs. This option works for both a single disk and multiple disks.

### Command Example

The following example flashes the Fault LED on disk 3 on channel 2 in disk array data\_center\_1:

```
amutil -f 2:3 data_center_1
```

---

## Managing the Rebuild Process

If a disk fails, the disk array automatically begins the rebuild process the first time an I/O is performed to the LUN, providing that there is a global hot spare available. If no global hot spare is available, the rebuild will not occur until the failed disk has been replaced.

While a rebuild is in process, you can check its progress and change the rate at which the rebuild occurs. A rebuild must be in process to perform either of these tasks.

### Checking Rebuild Progress

To check the progress of a rebuild, type:

```
amdsp -r <ArrayID>
```

The progress of all rebuilds currently taking place on the disk array will be displayed.

### Changing Rebuild Priority Settings

The rebuild process must compete with host I/Os for disk array resources. Two settings allow you to control whether the rebuild process or host I/Os will have greater access to these resources. These settings reflect how important the rebuild process is relative to I/O activity on the LUN being rebuilt.

If you select settings that give higher priority to the rebuild process, it will complete sooner but at the cost of lower I/O performance. Lower priority settings give precedence to host I/Os, which may delay the completion of the rebuild. These settings are used only if host I/Os are occurring during the rebuild.

To change the rebuild priority settings, type:

```
amutil -R <LUN> -f <freq> -a <amount> <ArrayID>
```

- **freq** identifies the rate at which the disk array attempts to execute rebuild commands. Specified in tenths of a second, this value can be set to 1 to 50, or 0.1 seconds to 5.0 seconds. A low setting increases the frequency at which rebuild commands are issued, giving higher priority to the rebuild, but reducing I/O performance. A high value reduces the rebuild command frequency, giving higher priority to host I/Os. The default value for this setting is 1, or 0.1 seconds.

- 
- **amount** identifies the number of blocks to rebuild at a time. This value can be from 1 to 64K and specifies the number of 512-byte blocks processed during each rebuild command. The higher the setting, the more blocks that will be processed, reducing I/O performance. A lower setting gives priority to host I/Os, delaying the completion of the rebuild. The default value for this setting is 64 blocks, or 32 Kbytes of data.

### Command Example

The following example assigns a value of 5 seconds to the rebuild command rate, and sets the data block amount to 16 blocks on LUN 4 on disk array 0000005EBD20. This gives host I/Os higher priority than the default settings.

```
amutil -R 4 -f 5 -a 16 0000005EBD20
```

### Performing a Parity Scan

To verify the integrity of the parity data on a LUN, you can perform a parity scan. This will check each block of data against its parity to ensure that they match.

Like a rebuild, a parity scan competes with host I/Os for disk resources, and can impact host I/O performance. The rebuild priority **amount** setting is used when performing a parity scan. A larger amount value may have a greater impact on I/O performance. To avoid impacting performance, perform a parity scan during periods of low host activity.

To perform a parity scan, type:

```
ammgr -P <LUN> <ArrayID>
```

### Command Example

The following example performs a parity scan on LUN 3 on disk array rack\_1.

```
ammgr -P 3 rack_1
```

### More About Parity Scanning

Parity information is used in RAID 5 LUNs to maintain data redundancy. When a single disk fails in a RAID 5 LUN, the disk array can reconstruct the data on the missing disk by using the parity information.

---

A parity scan compares the data and its associated parity to ensure they match. A parity scan cannot be performed on a LUN that has experienced a disk failure and is operating in degraded mode.

Although RAID 1 LUNs and 0/1 LUNs do not use parity, you can still perform a parity scan on them. The parity scan compares the data on the mirrored disks.

---

**Note**      **What should I do if parity errors are detected?**  
If errors are detected during a parity scan, it is recommended that you contact your Hewlett-Packard service representative immediately. The occurrence of parity errors may indicate a potential problem with the disk array hardware.

---

## Displaying Parity Scan Status

If a parity scan is in progress, you can monitor its progress. This will help you determine how much longer the parity scan will take.

To display parity scan progress, type:

```
ammgr -s <LUN> <ArrayID>
```

## Halting a Parity Scan

You can halt a parity scan currently in progress. This may be useful if the parity scan is impacting performance and you want to reschedule it to run when the host is less busy.

To halt a parity scan, type:

```
ammgr -H <LUN> <ArrayID>
```

## Managing Disk Array Logs

The disk array controller continuously gathers and stores information on all aspects of disk array operation. Array Manager 60 retrieves these controller log entries at regular intervals and stores them in a log directory on the host (`/var/opt/hparray/log`). The log files are valuable in monitoring disk array operation and isolating problems that may occur.

On HP08 and later, major event logging is available. MAjor event logging provides more detailed information on disk array operation. Many events that were not logged in the

---

previous firmware releases are logged in the major event log. Earlier versions of firmware (prior to HP08) use the standard log file format, also called Asynchronous Event Notification (AEN).

---

**Note** On firmware HP08 and later, major event logging is enabled by default. If major event logging has been disabled by disabling the UTM, only standard log entries will be available.

---

### Allocating Space for Disk Array Logs

The amount of space available for storing log files on the host is controlled by the host environment variable `AM60_MAX_LOG_SIZE_MB`. This variable defines the maximum size (in megabytes) allocated for storage of the disk array log files. The log files are stored as individual files in `/var/opt/hparray/log`. All log files will be maintained until the specified maximum value is exceeded. At that point, the oldest log file will be deleted to make room for the new file (FIFO).

The valid range for `AM60_MAX_LOG_SIZE_MB` is 100 to 4096 (100 MB to 4 GB). Any value outside of this range is ignored and the default value is used. The default value for allocated log storage area is 100 MB.

---

**Note**

- When changing the value for `AM60_MAX_LOG_SIZE_MB`, `AM60Srvr` must be stopped and restarted before the new value will take effect. To enforce the `AM60_MAX_LOG_SIZE_MB` value, it is necessary to stop and restart `AM60Srvr` every time the host is rebooted.
- When setting a large value for `AM60_MAX_LOG_SIZE_MB`, make sure the `/var` file system is large enough to accommodate the logs files. If it is not, the log files may overflow the `/var` file system.

---

---

## Viewing Disk Array Logs

To display the disk array controller log files, type:

```
amlog [-s <StartTime>] [-e <EndTime>] [-t <Recordtype[,Recordtype].....>]
[-c] [-d <LogDir>] [-a <ArrayID>]
```

- **StartTime** identifies the starting date and time. Log entries with an earlier date and time will not be displayed. The default is the time of the oldest log entry. The format for entering the date and time is MMddhhmm[yy], where:

MM = Month (01-12)

dd = Day (01-31)

hh = Hour (00-23)

mm = Minute (00-59)

yy = Year (00-99) [optional] Any number less than 90 will be evaluated as 2000 + yy

- **EndTime** identifies the ending date and time. Log entries with a later date and time will not be printed. The default is the time of the last log entry. Uses the same format as StartTime.
- **Recordtype** specifies the log record types to be displayed. The following record types are valid:
  - ctrlr - controller log sense. This is the default record type if none is specified.
  - mel - major event log. These entries can be filtered using the -c option. Major event logging is available in firmware versions HP08 and later.
- **-c** limits the major event log entries displayed to only those with critical status. Note that using this option will limit the amount of information displayed, including non-critical log entries that may be useful for isolating problems.
- **LogDir** identifies the path name of the log directory. The default is /var/opt/hparray/log. The log entries are organized into a series of files with the naming convention AML000nn. When a log files reaches a certain size, a new file is created using the next sequential number. An associated catalog file (AMLCATLG) provides the information for navigating the binary entries in the main log files.
- **ArrayID** identifies a specific disk array for which log entries should be displayed. The default is to print entries for all disk arrays connected to and recognized by the host. The

---

actual ArrayID must be used here. An alias cannot be used because alias names are not recorded in the log file.

### Command Example

The following example displays the major event log entries for disk array rack\_1. The log entries displayed are limited to only critical entries, and entries made after 0900 on 15 May 2000.

```
amlog -s 05150000 -t mel -c rack_1
```

### Sample Log Entries

The following is a sample of a major event log entry.

```
Major event log for Subsystem 001100A0B8060166 at Wed Nov 29
17:38:27 2000
  Sequence Number      = 14
  Time Stamp           = Wed Nov 29 17:38:19 2000
  Event Type           = Global Hot Spare Not Large Enough (224F)
  Event Category       = Notification
  Component Type       = Disk
  Enclosure            = 5
  Slot (0-based)       = 3
  Event Priority        = Informational
```

The following is a sample of an Asynchronous Event Log entry:

```
Controller log sense for Subsystem 001100A0B8060166 at Mon Dec 4
12:51:34 2000
  Controller Time Stamp = 032307 083318
  FRU Code              = 0x08
  FRU Code Qualifier    = 0x8142
  Sense Key             = 0x06
  Additional Sense Code = 0x3F
  Add Sense Code Qual   = 0xC7
Decoded Field Replaceable Unit Information:
  FRU Group             = Sub-Enclosure Group
  FRU ID Setting        = 1
  FRU Type              = Power Supply
```

---

```
FRU State                = Failed
Decoded SCSI Sense:
Non-media Component Failure
Reporting LUN            = 0
```



**This topic is discussed in greater detail in the Advanced User's Guide, which can be downloaded from <http://www.hp.com/support/FC60>.**

## Flushing Disk Array Log Contents

Array Manager 60 automatically retrieves the contents of the disk array controller log at regular intervals, typically 15 minutes. However, if necessary you can manually flush (retrieve) the contents of the disk array log to the host. This may be useful if you suspect a very recent event has not yet been retrieved.

To flush the contents of the disk array log, type:

```
amutil -l <ArrayID>
```

A message is displayed indicating that the log file has been flushed. The log contents can now be displayed using the `anlog` command as described above.

## Purging Controller Logs

The controller logs are retrieved at regular intervals and stored on the host. Over time, these logs may grow quite large. Major event logging in particular generates very large log files. To reduce the space consumed by the logs, the oldest log file can be purged from the log file directory.



---

To purge the oldest log file in the host directory, type:

**amutil -p**

**Note** Always use the `amutil -p` command to purge the controller logs. This command maintains the catalog pointers used to access the log files. Using a system command such as `rm` to remove the log files will cause log catalog errors.

Management of the log files can be automated by creating a script that purges the oldest log files at regular intervals using `amutil -p`. This technique can be used to ensure that the log files don't grow to a size that may cause the `?var` file system to overflow.

---

## Resetting Battery Age

The battery age should be reset to zero when the battery is replaced. This ensures an accurate indication of the age of the battery

To reset the battery age to zero, type:

**ammgr -b <ArrayID>**

---

## Managing the Disk Array Using STM

STM is an online diagnostic tool, but it can be used to perform some of the common tasks involved in managing the disk array. The tasks described here are available to all users and do not require the purchase of a license. See "Support Tools Manager" on page 199 for more information on using this tool.

### Checking Disk Array Status Information

The STM Information Tool displays disk array status information. See "Using the STM Information Tool" on page 204 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Information Tool > Run

### Binding a LUN

The STM Expert Tool can be used to bind a LUN. See "Using the STM Expert Tool" on page 207 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Expert Tool > Run Select: Utilities > Bind LUN

---

## Unbinding a LUN

The STM Expert Tool can be used to unbind a LUN. See "Using the STM Expert Tool" on page 207 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Expert Tool > Run Select: Utilities > Unbind LUN

## Adding a Global Hot Spare

The STM Expert Tool can be used to add a global hot spare. See "Using the STM Expert Tool" on page 207 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Expert Tool > Run Select: Utilities > Hot Spares > Create

## Removing a Global Hot Spare

The STM Expert Tool can be used to remove a global hot spare. See "Using the STM Expert Tool" on page 207 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Expert Tool > Run Select: Utilities > Hot Spares > Delete

---

## Locating Disk Modules

The STM Expert Tool can be used to locate disk modules, to aid in identification. The LEDs on the disk array components are flashed to aid in identification. See "Using the STM Expert Tool" on page 207 for more information on running and using this tool.

STM Tool	Action
xstm, mstm	Select: Tools > Expert Tool > Run Select: Utilities > Flash Component Lights > Drive or Utilities > Flash Component Lights > LUN or Utilities > Flash Component Lights > Disk Enclosure or Utilities > Flash Component Lights > Array

---

# 5

# HP-UX DIAGNOSTIC TOOLS

---

Overview ..... 198

---

Support Tools Manager ..... 199

---

---

## Overview

STM (Support Tools Manager) is the primary diagnostic tool available for the Disk Array FC60. For diagnosing problems, STM provides the capability to gather and display detailed status information about the disk array. STM can also be used to perform common management tasks.

---

## Support Tools Manager

The STM host-based utility provides capability for managing the Disk Array FC60. STM comes with HP-UX instant ignition and support media.

The Support Tools Manager (STM) host-based utility is the primary online diagnostic tool available for the HP SureStore E Disk Array FC60. STM provides the capability for testing, configuring, and evaluating the operational condition of the disk array. STM comes with HP-UX instant ignition and support media.

---

**Note** Using the information in this chapter requires some understanding of the basic operation of STM. If you are not familiar with STM, go to the Systems Hardware, Diagnostics, and Monitoring section of HP's Online Documentation web site ([www.docs.hp.com](http://www.docs.hp.com)), which contains information on using STM, including a tutorial to get you started. You can also download a copy of the Diagnostic/IPR Media User's Guide, which contains complete information on using STM.

---

### STM User Interfaces

STM presents three different user three interfaces: an X-Windows interface (xstm), a menu-based interface (mstm), and a command-based interface (cstm). Each STM interface has an extensive online help facility as well as man pages.

---

**Note** Each Disk Array FC60 is displayed twice in the STM interface — once for each active disk array controller. Each controller is identified by a unique hardware path. STM operations can be performed by selecting either disk array controller. The effect on the disk array will be the same regardless of which controller is selected.

STM identifies the Disk Array FC60 by its product number of A5277A. This is the identifier you will look for when selecting a disk array in STM.

---

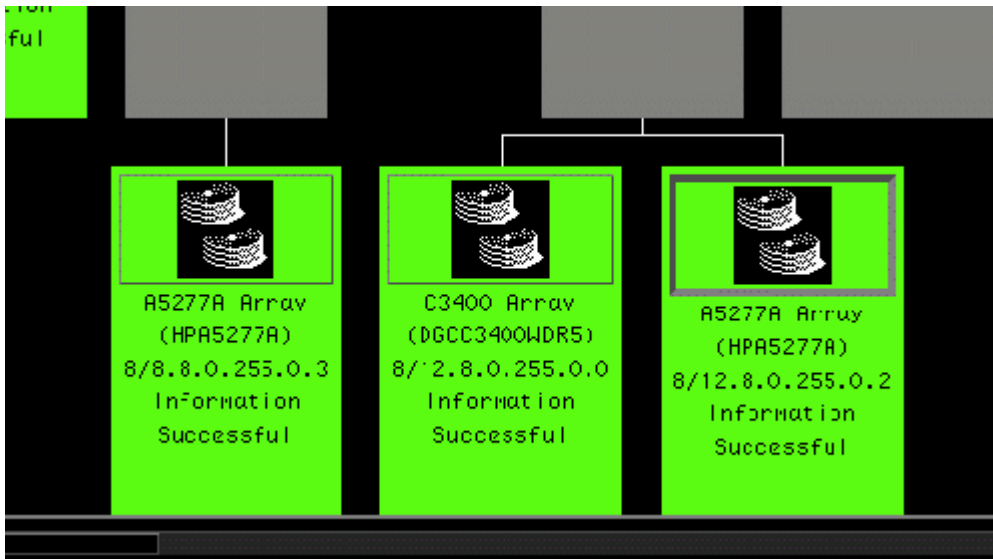
---

## xstm — the X Windows Interface

xstm is the X-Windows screen-based STM interface. Because it is the easiest to use, xstm is the recommended interface for systems that support graphical displays.

The main xstm window displays a map representing system resources. The STM system map represents each Disk Array FC60 as two icons labeled “A5277A Array”. See figure Figure 45. Each icon represents one of the disk array controllers, which are identified by their hardware paths. Click on the icon for the disk array you will be testing.

Navigation through xstm is done using menus and submenus. For more information on navigating in xstm, download a copy of the STM Quick Reference Guide from the Systems Hardware, Diagnostics, and Monitoring section of HP’s Online Documentation website ([www.docs.hp.com](http://www.docs.hp.com)).



**Figure 45** xstm Interface Main Window Disk Array Icons



---

## mstm — the Menu-based Interface

mstm is the menu-based STM interface. It serves as an alternate interface for systems that do not support graphical displays.

The main mstm window displays a list of system resources. The Disk Array FC60 is identified as product type “A5277A Array”. See Figure 46. Each entry in the list represents one of the disk array controllers, which are identified by their hardware paths. Select the entry for the disk array you will be testing.

In mstm, you traverse screens and menus, and select commands from pulldown menus, which are similar to those found in xstm.

- **Navigating between screens.** To navigate between one screen and another, use the Tab key.
- **Navigating within screens.** To navigate from one portion to another portion of the same screen, use the Prev and Next keys (or, alternatively, the Cursor Up and Cursor Down keys).
- **Navigating the menu bar.** The following methods are used to navigate the menu bar:
  - Position the cursor on the menu bar — Use the Tab key (or the MenuBar on/off function key).
  - Move to a particular pulldown menu — Use the cursor arrow keys.
  - Expand a menu/sub-menu — Use the Return key.
  - Highlight a command or sub-menu — Use the cursor keys.
  - Perform a command — Use the Return key.
  - Invoke a menu directly — Use Alt function key; then press nonunderlined character in menu title.
  - Activate a menu command/expand pulldown sub-menu — Press the letter that is underlined in command or submenu.

For more information on navigating in mstm, download a copy of the STM Quick Reference Guide from the Systems Hardware, Diagnostics, and Monitoring section of HP’s Online Documentation website ([www.docs.hp.com/hpux/systems](http://www.docs.hp.com/hpux/systems)).

The screenshot shows a terminal window titled 'xterm' with a blue header bar. The terminal displays the output of a command, likely 'mstm', showing a list of hardware devices. The output is organized into columns: File System, Device, Tools, and Options. The 'Tools' column contains 'Last Active Tool' and 'Last Op Status'. The 'Options' column contains the user's email address, 'morpheus.rose.hp.com'.

File System	Device	Tools	Options
			morpheus.rose.hp.com
Path	Product	Last Active Tool	Last Op Status
8	Bus Adapter (580)		
10	Bus Adapter (580)		
10/0	SCSI Interface (15)		
10/0.5.0	SCSI Disk (SEAGATEST19171)		
10/0.6.0	SCSI Disk (SEAGATEST19171)		
10/4	Bus Adapter (500)		
10/4/0	Access Port Multiplexor (		
10/8	Unknown ( )		
10/8.8	Fibre Channel Driver (Mas		
10/12	Core I/O Adapter (44)		
10/12/0	Centronics Interface (44)		
10/12/5	SCSI Interface (44)		
10/12/5.2.0	SCSI Disk (TOSHIBACD-ROM)		
10/12/6	HSC LAN Interface (44)		
10/12/7	Built-in Keyboard/Mouse (		
10/16	Bus Adapter (500)		
12	Bus Adapter (580)		
14	Bus Adapter (580)		
14/8	Unknown ( )		
14/8.8	Fibre Channel Driver (Mas		
14/8.8.0.255.0.0	A5277A Array (HPA5277A)	Expert Tool Successful	
14/8.8.0.255.0.0.1	Unknown ( )		
14/8.8.0.255.0.0.2	Unknown ( )		
14/8.8.0.255.0.0.3	Unknown ( )		
14/12	Unknown ( )		
14/12.8	Fibre Channel Driver (Mas		
14/12.8.0.255.0.1	A5277A Array (HPA5277A)	Expert Tool Incomplete	
14/12.8.0.255.0.1.1	Unknown ( )		

**Figure 46** mstm Interface Main Window

---

## STM Tools

The STM tools available for use with the HP SureStore E Disk Array FC60 are listed in Table 24.

**Table 24** Available Support Tools

Tool type	Description
Information	Provides detailed configuration and status information for all disk arrays components.
Expert	Provides capability to perform common disk array management tasks.

## STM Device Logs

When you run an STM support tool, the results are indicated by the status of the device icon on the system map (xstm) or the device entry on the device list (mstm, cstm). If the results are other than successful, you can view the appropriate device log for more information.

- **Failure Log** - If the results indicate Failure, view the device Failure Log for information identifying the likely cause of the failure and recommended action for correcting the problem.
- **Test Activity Log** - If the results are other than Successful or Failure, view the Test Activity Log for an explanation of what occurred.

---

## Using the STM Information Tool

The STM Information Tool gathers status and configuration information about the selected disk array and stores this information in three logs: the information log, the activity log, and the failure log.

### Running Information Tool in X Windows

1. At the system prompt:
  - Type **xstm &**
2. Click on the desired A5277A Array device icon.
3. To run the Information Tool and view the Information log:
  - From the menu bar, select Tools
  - Select Information
  - Select Run. The Information Tool builds and displays the Information log.
  - Select Done when done viewing the log.
4. To view the Activity log:
  - From the menu bar, select Tools
  - Select Information
  - Select Activity Log.
  - Select Done when done viewing the log.
5. To exit STM:
  - From the menu bar, select File
  - Select Exit

---

## Running Information Tool in Menu Mode

1. At the system prompt:
  - Type **mstm**
  - Select Ok
2. To select the desired disk array:
  - Scroll down using the arrow key, select the A5277A Array
  - Press <Enter>.
3. To run the Information Tool and display the Information log:
  - From the Menubar, select Tools
  - Select Information
  - Select Run. The Information Tool builds and displays the Information log.
  - Select Done when done viewing the log
4. To view the Activity log:
  - From the Menubar, select Tools
  - Select Information
  - Select Activity Log
  - Select Done when done viewing the log
5. To exit STM:
  - From the Menubar, select File
  - Select Exit
  - Select OK

---

## Interpreting the Information Tool Information Log

The Information Log contains status and configuration information for all disk array components. The log is separated into the following sections:

- Controller Enclosure – information for the components in the disk array controller enclosure
- Disk Enclosure(s) – information about the components in the disk enclosure(s), excluding the disks
- Map – a graphic layout showing the configuration of the disks enclosures and disks
- Configuration – information on LUN configuration
- Disks – information about each disk

Much of the information is self-explanatory, but if you need assistance interpreting any fields, a complete online description is available by performing the following steps:

1. Run the Information tool on the desired disk array as described in the preceding section.
2. Return to the main STM window and select the Tools menu.
3. Select Information.
4. Select Info. The Information Log Description window is displayed describing each field.
5. Select Done to close the description window.

## Interpreting the Information Tool Activity Log

The Information Tool Activity Log allows you to review details of the most recent execution of the Information tool against a specific device. You can view, save, and print the contents of the activity log, as desired.

---

## Using the STM Expert Tool

The Expert Tool provides the capability to manage the HP SureStore E Disk Array FC60. Before using the Expert Tool for the first time you are encouraged to read through the Expert Tool help topics. The Step-by-Step instructions in particular provide useful tips on using the Expert Tool.

As you perform tasks using the Expert Tool, the status of each operation is displayed in the main window. The Expert Tool window also provides a status of the activities performed on the selected device.

### Running Expert Tool in X Windows

1. At the system prompt:
  - Type **xstm &**
2. Click on the A5277A ARRAY device icon
3. To run the Expert tool:
  - From the menu bar, select Tools
  - Select Expert
  - Select Run
4. Perform the desired operation using the menus in the Expert Tool window. The Expert Tool menu options are listed in Table 25.
5. To exit the Expert Tool:
  - From the menu bar, Select File
  - Select Exit

---

## Running Expert Tool in Menu Mode

1. At the system prompt:
  - Type **mstm**
  - Select Ok
2. To select the disk array:
  - Scroll down using the arrow key, select A5277A ARRAY.
  - Press <Enter>
3. To run the Expert tool:
  - Select Menubar on or use arrow keys to get to Menubar
  - Select Tools
  - Select Expert Tool
  - Select Run
4. Perform the desired operation using the Expert Tool menus. The Expert Tool menu options are listed in Table 25.
5. To exit Expert Tool:
  - Select Menubar on or use arrow keys to get to Menubar
  - Select File
  - Select Exit
  - Select Yes

### Expert Tool Menu Options (xstm and mstm)

When using xstm or mstm, all Expert Tool functions are accessed from a series of menus. The menu options are described in Table 25.



**Table 25** Expert Tool Menus and Descriptions

<b>Menu</b>	<b>Option</b>	<b>Property</b>	<b>Description</b>	
Logs	View Event Log	NA	Displays selected event log entries	
Tests	Parity Scan	NA	Perform a parity scan on a LUN.	
Utilities	Bind LUN	NA	Bind selected disk modules into a LUN with a specified RAID level.	
	Unbind LUNs	NA	Unbind a LUN.	
	Replace LUN Zero	NA	Unbind and rebind LUN 0	
	Hot Spares	Create		Create hot spares.
		Delete		Delete hot spares.
	Flash Component Lights	Drive		Flash Fault LED on selected disk.
		LUN		Flash Fault LEDs on all disks in the selected LUN.
		Disk Enclosure		Flash Fault LEDs on all disks in the selected disk enclosure.
Array			Flash Fault LEDs on all disks in the disk array.	



---

# 6

# TROUBLESHOOTING

---

<b>Introduction</b> .....	<b>212</b>
<b>Disk Array Installation/Troubleshooting Checklist</b> .....	<b>217</b>
<b>Power-Up Troubleshooting</b> .....	<b>218</b>
<b>Controller Enclosure Troubleshooting</b> .....	<b>219</b>
<b>SureStore E Disk System SC10 Troubleshooting</b> .....	<b>222</b>

---

---

## Introduction

The modular design of the Disk Array FC60 simplifies the isolation and replacement of failed hardware components. Most disk array components are hot-swappable Field Replaceable Units (FRUs), which can be replaced while the disk array is operating. Some of the FRUs are customer replaceable. Other array components can be replaced in the field, but only by a trained service representative. A complete list of product and part numbers are included in "Replaceable Parts" on page 264.

---

**CAUTION** *To ensure proper operation, disk array controller modules and disk modules should not be removed or replaced when disk array power is off. Also, disk modules should not be moved to different slot locations within the disk enclosure.*

---

## About Field Replaceable Units (FRUs)

The Disk Array FC60 consists of a Controller Enclosure and one or more SureStore E Disk System SC10 enclosures. Table 26 identifies the disk array FRUs and whether they are customer replaceable. See "Removal and Replacement" on page 229 for more information.

**Table 26** Field Replaceable Units

Field Replaceable Units	Customer Replaceable Unit	Hot Swappable Modules
<b>Disk Enclosure</b>		
Disk Modules	Yes	Yes
BCC (Bus Controller Card) Modules – Note - disk enclosure should be powered down when replacing the BCC.	No	
Power Supply Modules	Yes	Yes
Disk Fan Modules	Yes	Yes
<b>Controller Enclosure</b>		
Controller Modules	No	Yes
Media Interphase Adapter (MIA)	Yes	Yes
Battery Backup Module	Yes	Yes
Controller Fan Module	Yes	Yes
SCSI Cables	Yes	No
Fibre Channel Cables	Yes	Yes
Power Supply Modules	Yes	Yes
Power Supply Fan Module	Yes	Yes

---

## HP-UX Troubleshooting Tools

There are several tools available for troubleshooting the disk array on an HP-UX host. This includes monitoring the operation of the disk array and gathering information that will help identify and solve the problem.

- Array Manager 60 - primarily used to manage the disk array, Array Manager 60 can also be used to check the status of the disk array and to retrieve log information.
- Support Tools Manager (STM) - the Information Tool of STM is one the most valuable sources of information for identifying problems with the disk array.

## Windows NT and Windows 2000 Troubleshooting Tools

The *HP Storage Manager 60* software includes a set of tools for troubleshooting the disk array. See the *HP Storage Manager 60-NT User's Guide* for information on using the *HP Storage Manager 60* software.

## EMS Hardware Event Monitoring (HP-UX Only)

EMS hardware monitoring is one the most important troubleshooting tools available for the Disk Array FC60. With EMS monitoring, you can be alerted to problems as they occur, allowing you to respond quickly to correct a problem before it impacts disk array operation. All components and aspects of disk array operation are monitored. And EMS gives you the flexibility to deliver event notification using a variety of different methods.

More information on EMS Hardware Monitoring can be downloaded from the Systems Hardware, Diagnostics, and Monitoring section of HP's Online Documentation web site at [www.docs.hp.com/hpux/systems](http://www.docs.hp.com/hpux/systems).

---

## EMS Monitor Event Severity Levels

Each event detected and reported by the EMS monitor is assigned a severity level, which indicates the impact the event may have on disk array operation. The following severity levels are used for all events:

Critical	An event that causes host system downtime, or other loss of service. Host system operation will be affected if the disk system continues to be used without correction of the problem. Immediate action is required.
Serious	An event that may cause, host system downtime, or other loss of service if left uncorrected. Host system and hardware operation may be adversely affected. The problem needs repair as soon as possible.
Major Warning	An event that could escalate to a serious condition if not corrected. Host system operation should not be affected and normal use of the disk system can continue. Repair is needed but at a convenient time.
Minor Warning	An event that will not likely escalate to a severe condition if left uncorrected. Host system operation will not be interrupted and normal use of the disk system can continue. The problem can be repaired at a convenient time.
Information	An event that is expected as part of the normal operation of the hardware. No action is required.

## Sample EMS Event Message

A sample EMS hardware monitor event message is shown in Figure 47. A message typically includes the following information:

- Message Data – Date and time the message was sent, the source and destination of the message, and the severity level
- Event Data – Date and time of the event, the host, event ID, name of the monitor, event number, event class, severity level, hardware path, associated OS error log entry ID
- Error Description – Information indicating the component that experienced the event and the nature of the event

- 
- Probable Cause/Recommended Action – The cause of the event and suggested steps toward a solution. This information should be the first step in troubleshooting.

```
Notification Time: Thu Aug 6 15:18:03 1998
yourserver sent Event Monitor notification information:
/peripherals/events/mass_storage/LVD/enclosure/10_12.8.0.255.0.10.0 is
!=1.
Its current value is CRITICAL(5)
Event data from monitor:
Event Time: Thu Aug 6 15:18:03 1998
Hostname: yourserver.rose.hp.com      IP Address: 15.43.212.175
Event ID: 0x0035ca2b9b00000002      Monitor: dm_ses_enclosure
Event # : 101                        Event Class: I/O
Severity: Critical

Enclosure at hardware path 10/12.8.0.255.0.10.0: Hardware failure
Associated OS error log entry id(s): None

Description of Error:

    The disk device is indicating a fault.

Probable Cause/Recommended Action:

    The disk drive in slot 1 has failed. Check the connection, reseal,
    or replace the disk drive.
```

**Figure 47** Sample EMS Hardware Event Notification Message



---

## Disk Array Installation/Troubleshooting Checklist

The following checklist is intended to help isolate and solve problems that may occur when installing the disk array.

- Check Fibre Optic and SCSI Cables and SCSI Terminators:
  - No damaged fibre optic cables
  - No damaged or loose screws on connectors
  - All cables tightly secured to the connectors on the Fibre Channel Interconnect PCA
  - Shortest possible fibre optic cable lengths between disk arrays and host adaptors
- Check Disk Array (Pre-Power-Up):
  - All modules properly seated in the disk array enclosures
  - Disk modules in their proper slots
- Check Disk Array Functionality and Configuration (Following power up)
  - Date and time synchronized with the host.
  - No fault indicators on any disk array components
  - All disk modules spun up
  - LUN 0 replaced on non-factory built disk array (A5277A).
  - Host ioscan displays proper disk array configuration
  - Fibre Channel drivers installed
  - Different capacity disk modules not mixed in same LUN
  - Disk modules with different part/product numbers not mixed in same LUN

---

## Power-Up Troubleshooting

When the disk array is powered up, each component perform an internal self-test, to ensure it is operating properly.

Visual indications of power-up are:

- The green Power LED on the controller enclosure is on
- The green Power LED on each disk enclosures is on
- All fans are operating
- No Fault LEDs are on. See Figure 48 on page 220 and Figure 49 on page 223.

The following LEDs will also be on:

- Controller enclosure:
  - The green Full Charge LEDs on the controller batteries, if the batteries are fully charged. If the batteries are not fully charged, the Full Charge LEDs will flash until charging is complete, and then remain on.
  - The Power LEDs on the controller enclosure
- Disk enclosure
  - The green Power Supply LEDs
  - Fan LEDs

If the power-up sequence fails, the appropriate amber Fault LED will be on, identifying the failing component. If the power-up tests fail, all modules in the disk array must be examined for Fault LEDs that are on. See Table 27 on page 221 and Table 28 on page 223).

---

**Note**      Inserting a module while the power-up tests are running can cause the tests to fail. Insert all modules and then cycle power to rerun the power-up tests.

---

---

**Note** If no LEDs are ON and the fans are not running, it indicates that no AC power is being supplied to the disk array power supply modules. Check the input AC power to the disk array.

---

See "Applying Power to the Disk Array" on page 109 for information on powering up the disk array.


## Controller Enclosure Troubleshooting

### Introduction

This section contains information on identifying problems with the disk array controller enclosure.

## Controller Enclosure LEDs

Figure 48 shows the locations of the status LEDs for the controller enclosure. Table 27 summarizes the operating LED states for all components within the controller enclosure.



**Figure 48** Controller Enclosure LEDs

Table 27 Normal LED Status for Controller Enclosure

Module	LED	Normal State
Controller Enclosure	Power On	On (green)
	Power Fault	Off
	Fan Fault	Off
	Controller Fault	Off
	Fast Write Cache	On (green) while data is in cache
Controller	Controller Power	On (green)
	Controller Fault	Off
	Heartbeat	Blink (green)
	Status	Green There are 8 status LEDs. The number and pattern of these LEDs depend on how your system is configured.
Controller Battery	Fault - B	Off
	Full Charge - B	On (green) <sup>1</sup>
	Fault - A	Off
	Full Charge - A	On (green) <sup>1</sup>
Controller Power Assembly	Power 1	On (green)
	Power 2	On (green)
Controller Fan Assembly	Fan Power	On (green)
	Fan Fault	Off

1. Both Full Charge A and Full Charge B LEDs are ON after batteries are fully charged. The LEDs flash while charging is in progress, and remain on when charging is complete.

---

## SureStore E Disk System SC10 Troubleshooting

This section contains information on identifying and isolating problems with the Disk System SC10 disk enclosure.

### Disk Enclosure LEDs

Figure 49 shows the locations of the disk enclosure status LEDs. Table 28 summarizes the operating LED states for all components within the disk enclosure.


### Losing LUN 0

If LUN 0 becomes unavailable because of multiple disk failures, Array Manager 60 may not be able to communicate with the disk array. If this occurs, replace the failed disks as quickly as possible. When the disk are replaced it may be necessary to replace LUN 0, recreating it in its original configuration.

---

**CAUTION** *If LUN 0 is lost, do not power off the disk array. Without LUN 0, the disk array will not operate normally when it is powered back on.*

---



**Figure 49** Disk Enclosure LEDs

**Table 28** Disk Enclosure LED Functions

LED	State	Indication
System Power	Green	Power is on. Normal operation.
	OFF	Power is off
System Fault	Amber	Self-test <sup>1</sup> / Problem <sup>2</sup>
	OFF	Normal operation
	Flashing	BCC A & B DIP switch settings do not match

**Table 28** Disk Enclosure LED Functions (cont'd)

LED	State	Indication
BCC Fault	Amber	Self-test <sup>1</sup> / Fault
	OFF	Normal operation
	Flashing	Peer BCC DIP switch settings do not match
LVD	Green	Bus operating in LVD mode
	OFF	Bus operating in single-ended mode
Term. Pwr.	Green	Termination power is available from the host. Normal operation.
	OFF	There is no termination power.
Full Bus	Green	Disks are on a single bus of ten addresses.
	OFF	Disks are split between two buses, five addresses each
Bus Active	Green	Bus is available for use.
	OFF	Isolator chip is disabled and bus is not available.
Fan	Green	Normal operation
	Amber	Startup <sup>1</sup> / Fault
	OFF	Power is off
Power Supply	Green	Normal operation
	Amber	Startup <sup>1</sup> / Fault
	OFF	Power is off
Disk Fault	OFF	Disk is operating normally.
	Amber	A fault has occurred with the disk.
	Flashing	Used to identify a disk using the software management tools.
Disk Activity <sup>3</sup>	Flashing	Indicates that the disk is being accessed by the disk array.

1 Startup and self-tests occur briefly when the unit is powered ON.

2 A component has failed; temperature or voltage is out of normal range.  
See "Isolating Causes" on page 226

3 When a disk is installed with power on, its activity LED stays ON until the disk has spun up. When the disk is ready, the LED turns off. Thereafter, it flashes when there is I/O to the disk.



---

**Note** It is normal for the amber Fault LED on a component to go on briefly when the component initially starts up. However, if the Fault LED remains on for more than a few seconds, a fault has been detected.

---

## Interpreting Component Status Values (HP-UX Only)

Common status terms have specific indications for various disk enclosure components. The component status condition terminology used by Array Manager 60 (AM60) may differ from that used by STM. Both terms are identified in Table 29.

**Table 29** Status Indications by Component

Component/Status	Disks	BCCs	Fans, Power Supplies	Sensors
<b>AM60:</b> Good <b>STM:</b> Optimal	Component is installed and no error conditions are known.			
<b>AM60:</b> Critical <b>STM:</b> Fault - 2	N/A	Hardware has failed.		Voltage/temperature exceeds critical limit.
<b>AM60:</b> Non-critical <b>STM:</b> Fault - 3	N/A	BCC A and B have different firmware versions.	N/A	Voltage/temperature exceeds warning limit.
<b>AM60:</b> Not Installed <b>STM:</b> Not Present	Component is not installed.			
<b>AM60:</b> Unknown <b>STM:</b> Fault - 0	N/A	N/A	N/A	Sensor has failed or status is not available.
Not Available	Component is installed without known errors, but has not been turned on or set into operation; or the controller that owns reporting has failed.			

## Isolating Causes

Table 30 lists the probable causes and solutions for problems you may detect on the disk enclosure. When more than one problem applies to your situation, investigate the first description that applies. The table lists the most basic problems first and excludes them from subsequent problem descriptions.

**Table 30** Disk Enclosure Troubleshooting Table

Problem Description	HW Event Category	LED State	Status	Probable Cause/Solution
The disk enclosure does not power on	none	System Power LED off	none	<ul style="list-style-type: none"> <li>– Neither power cord is plugged in.</li> <li>– The power switch is not pressed.</li> <li>– AC breaker is tripped.</li> <li>– AC power source has failed.</li> <li>– The PDU/PDRU is defective.</li> <li>– Power switch is defective.</li> <li>– A faulty component is causing power supplies to turn off.</li> </ul> Remove all components and reinsert one at a time until the faulty component is isolated.
System fault LED is on	none	Power Supply LED off	Power supply Not Available	<ul style="list-style-type: none"> <li>– The power supply is not plugged in.</li> <li>– The PDU/PDRU or primary power source has failed.</li> </ul>
	Critical	Part Fault LED on	Critical	A component has failed. See problem descriptions below.
Buzzer sounds when BCC installed	Critical	BCC Fault flashing	Critical	BCC DIP switch settings do not match peer BCC switch settings. Reset switches.
BCC Fault LED is on	Critical	BCC Fault on	Critical	BCC hardware is faulty. Replace the BCC.
Fan LED is amber	Critical	Amber	Critical	Fan has slowed or stopped. Replace the fan.

**Table 30** Disk Enclosure Troubleshooting Table (cont'd)

Problem Description	HW Event Category	LED State	Status	Probable Cause/Solution
Power Supply LED is amber	Critical	Amber	Critical	<ul style="list-style-type: none"> <li>– An incompatible or defective component caused a temporary fault.</li> <li>– Power supply hardware is faulty. Unplug the power cord and wait for the LED to turn off. Reinsert the power cord. If fault persists, replace the power supply.</li> </ul>
Temperature is over limit	Critical	none	Critical Temp is >54.5° C	<ul style="list-style-type: none"> <li>– A fan is faulty. Check status and correct.</li> <li>– Airflow is obstructed; vents are blocked.</li> <li>– One or more slots are empty.</li> <li>– Power supply is faulty. Check status and correct.</li> <li>– Room temperature is too high. If ambient temperature cannot be reduced in a reasonable time, turn OFF product to prevent shortened life.</li> <li>– Temperature sensor is faulty. Compare temperature reported by peer BCC.</li> </ul>
	Major Warning	none	Non-critical Temp is >36° C	
<p><b>Temperature sensors are on the BCC and are independent of power supplies. Investigate temperature warnings immediately, before power supplies sense critical temperature and turn off.</b></p>				
Temperature is under limit	Critical	none	Critical Temp is <9.5° C	<ul style="list-style-type: none"> <li>– Room temperature is too low.</li> <li>– Temperature sensor is faulty. Compare temperature reported by peer BCC.</li> </ul>
	Major Warning	none	Non-critical Temp is <15.5° C	
Voltage is over limit	Critical	none	Critical	Power supply is faulty. Check status and correct.
	Major Warning	none	Non-critical	
Voltage is under limit	Critical	none	Critical	Either power supply is faulty. Check status and correct.
	Major Warning	none	Non-critical	

---

**Table 30** Disk Enclosure Troubleshooting Table (cont'd)

<b>Problem Description</b>	<b>HW Event Category</b>	<b>LED State</b>	<b>Status</b>	<b>Probable Cause/Solution</b>
Peer BCC status, temperature and voltage are Not Available	Major Warning	none	Both BCCs: Non-critical	Firmware on BCC A and BCC B are different versions.
		none	Not Available	Internal bus is faulty. Contact HP technical support to replace midplane.

---

**7****REMOVAL AND REPLACEMENT**

---

<b>Overview</b> .....	<b>230</b>
<b>Disk Enclosure Modules</b> .....	<b>232</b>
<b>Controller Enclosure Modules</b> .....	<b>242</b>

---

---

## Overview

This chapter describes removal and replacement procedures for the disk array hot-swappable modules that are customer replaceable. Hot-swappable modules can be replaced without impacting host interaction with the disk array.

Procedures for replacing the following modules are included in this chapter:

- Disk Enclosure
  - Disk Modules
  - Power Supply Modules
  - Fan Modules
- Disk Array Controller Enclosure
  - Battery Backup Module
  - Controller Fan Module
  - Power Supply Modules
  - Power Supply Fan Module

All modules other than those listed above must be replaced by a trained service representative.

---

### Note

#### **What is a “hot-swappable” component?**

The term hot-swappable identifies a component that can be removed and replaced without disrupting host access to the disk array. Host I/Os continue without interruption when replacing a hot-swappable component. Disk array controller modules, disk modules, power supply modules, and fan modules are all hot swappable components.

---

---

**Note****Is the HP SureStore E Disk Array FC60 customer repairable?**

Although the modular design of the Disk Array FC60 makes it easy to replace failed components, it is recommended that repair of the product be done by a trained service representative. This includes troubleshooting, and removal and replacement of hot-swappable components. Customer repair should be performed only in situations where a representative is not available, or system operation may be impacted if repair is not performed immediately.

The HP SureStore E Disk Array FC60 is fully covered by a warranty from Hewlett-Packard. Additional support services may have also been purchased for the disk array. During the warranty period, or if the product is covered by a service contract, it is recommended that you contact your service representative for all service and support issues.

---

---

## Disk Enclosure Modules

This section describes the procedures for replacing the hot swappable modules in the disk enclosure.

### Disk Module or Filler Module

**Hot Swappable Component!**

This procedure describes how to add or replace disk modules and disk slot filler modules. When adding or replacing disk filler modules use the same procedure, ignoring any steps or information that applies to disk modules only.

---

**CAUTION** - To prevent damage from static electricity, follow standard ESD procedures and avoid touching exposed circuitry.



- Do not remove a disk module or filler module until you have a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.

- To avoid removing the wrong disk, it is recommended that you identify the disk by flashing its amber Fault LED. Identifying a disk can be done using the management software tools. Removing the wrong disk can cause data to become unavailable if the disk is part of the same LUN as the failed disk.

- Never move disks from one array to another or within an array with power on. Moving disks with power on may result in data loss. Always power down the disk array before moving disks. Moving disk should not be confused with replacing a failed disk, which can be done with power on.

---



---

**Note** When a disk module is replaced, the new disk inherits the group properties of the original disk. For example, if you replace a disk that was part of LUN 1, the replacement will also become part of LUN 1. If the disk is a replacement for a global hot spare or an unassigned disk, the replacement will become a global hot spare or an unassigned disk.

A special feature called drive lockout prevents unsupported disk drives from being used in the disk array. If an unsupported disk drive is installed in the disk array, the drive will be failed. Locked out drives will be indicated with a special drive status (0x1B).

---

## Removing a Disk Module or Filler Module

1. Unlock and open the disk enclosure door.
2. If removing a disk module, insert the plug end of your ESD wrist strap in the disk enclosure ESD socket (A in Figure 50).
3. Release the module by squeezing the latch tab (B in Figure 50) and sliding the module part way out of the enclosure.

---

**CAUTION** *While the disk is spinning, it is vulnerable to damage. Wait for the disk to stop spinning (about 15 seconds) before fully removing the module from the enclosure.*

---

---

**WARNING** **To avoid electric shock hazard, do not touch the backplane or adjacent drive electronics when removing and inserting disks.**


---

4. Slide the module out of the slot until you can grasp the handle (C). Support the disk module with your other hand around the enclosed side.

---

**Note** When removing adjacent disk modules, removing them from right to left improves access to successive modules.

---



**Figure 50** Disk Module Removal

## Installing a Disk Module or Filler Module

**CAUTION** *Touching the disk circuit board can cause high energy discharge and damage the disk.*



*Disks modules are fragile and should be handled carefully.*




---

**Note** If the disk module you are installing has been removed from another Disk Array FC60, you should ensure that the module has a status of Unassigned. This is done by unbinding the LUN the disk module was a part of in the original disk array. See "Moving Disks from One Disk Array to Another" on page 136.

---

1. Remove the replacement disk from its ESD bag, being careful to grasp the disk by its handle (A in Figure 52).
2. Pull the cam latch (B) away from the disk module.
3. Mark the slot number that the disk is being installed into on the Slot Location Map on the top of the disk module. See Figure 51. This facilitates reinstalling the disk module in the correct slot.



**Figure 51** Disk Slot Identification Label

---

**Note** Use both hands to hold the disk — one on the handle and the other on the carrier frame.

---

4. Slide the disk module into the slot with capacity label up (C in Figure 52).

- 
5. Close the cam latch to seat the module firmly into the backplane. An audible click indicates the latch is closed properly.
  6. Check the LEDs (D in Figure 52) above the disk module for the following behavior:
    - Both LEDs should turn on briefly.
    - The amber Fault LED should turn off.
    - The green disk Activity LED should blink for a few seconds and then go out. If the host begins to access the disk, the Activity LED will flash.

If this behavior does not occur, refer to chapter 6, Troubleshooting, for information on identifying and solving the problem.

7. Unplug the ESD strap, and close and lock the disk enclosure door.
8. Disable WCE on the new disk module using the following command:

**amutil -w <ArrayID>**

See "Disabling Disk Module Write Cache Enable (WCE)" on page 183 for more information.




Figure 52 Disk Module Replacement

---

## Disk Enclosure Fan Module

### Hot Swappable Component!

A failed fan module should be replaced as soon as possible. There are two fan modules in the enclosure. If a fan fails, the remaining fan module will maintain proper cooling. However, if the remaining fan module fails before the defective fan is replaced, the disk enclosure must be shut down to prevent heat damage.


---

**CAUTION** *Do not remove a disk fan module until you a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

---

### Removing the Fan Module

1. Loosen the two module locking screws (A in Figure 53).
2. Slide the fan module out of the chassis by the pull tab (B).



**Figure 53** Disk Enclosure Fan Module Removal and Replacement

### Installing the Fan Module

1. Slide the replacement fan module into the empty slot (C in Figure 53).
2. Tighten the locking screws (A).
3. Check the fan module LED for the following behavior:
  - The Fan Fault LED should flash amber, and then turn green after a few seconds. If the LED does not turn green, refer to "Troubleshooting" on page 211.

---

## Disk Enclosure Power Supply Module

### Hot Swappable Component!

A failed power supply module should be replaced as soon as possible. When one power supply fails, the remaining power supply will maintain the proper operating voltage for the disk enclosure. However, if the remaining power supply fails before the first power supply is replaced, all power will be lost to the disk enclosure.

---


**CAUTION** *Do not remove a power supply module until you have a replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

---

### Removing the Power Supply Module

1. Disconnect the power cord to the power supply.
2. Loosen the screw (B in Figure 54) from the power supply handle (A).
3. Pull the handle down to disengage the power supply from the backplane.
4. Slide the power supply out of the chassis. Support the power supply with your free hand as it clears the chassis.





**Figure 54** Disk Enclosure Power Supply Module Removal and Replacement

### Installing the Power Supply Module

1. With the handle down, slide the replacement power supply into the empty slot (D in Figure 54). The supply begins to engage the backplane with 3/8 inch (8 mm) still exposed.
2. Swing the handle upward to seat the power supply into the backplane. The power supply should be flush with the chassis.
3. Tighten the screw (B) in the power supply handle (A).
4. Plug the power cord into the power supply and electrical source.
5. Check the Power Supply LED. It should then turn green. If the LED remains off or is amber, see chapter 6, Troubleshooting for information on identifying and solving the problem.

---

## Controller Enclosure Modules

This section provides removal and replacement procedures for the controller enclosure modules, plus the controller enclosure front cover. Most controller modules are hot swappable, however certain restrictions need to be observed for some modules, as identified in these descriptions.

The controller modules, the controller fan module, and the BBU are accessed from the front of the controller enclosure. Access to these modules requires that the front cover be removed. The power supply modules and the power supply fan module are accessed from the rear of the controller enclosure.

---

**CAUTION** *Do not remove a module from the controller enclosure until you are ready to install the replacement. An empty slot will disrupt the cooling airflow within the enclosure and may cause overheating.*

---

---


## Front Cover Removal/Replacement

### Hot Swappable Component!

To gain access to the front of the controller module, the controller fan module, or the battery backup unit (BBU), the front cover must be removed.

### Removing the Front Cover

1. Pull the bottom of the cover out about one inch to release the pins. See Figure 55.
2. Slide the cover down one inch and pull it away from the controller enclosure.



**Figure 55** Controller Enclosure Front Cover Removal and Replacement

---

## Installing the Front Cover

1. Slide the top edge of the cover up under the lip of the chassis.
2. Push the cover up as far as it will go, then push the bottom in until the pins snap into the mounting holes.

## Controller Fan Module

**Hot Swappable Component!**


---

**CAUTION** *Do not operate the controller enclosure without adequate ventilation and cooling to the controller modules. Operating without proper cooling to the controller modules may damage them. You must shut down the controller enclosure to avoid overheating if you anticipate needing more than 15 minutes to replace the fan module.*

---

## Removing the Controller Fan Module:

1. Remove the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 243.
2. Loosen the screw on the fan module. See Figure 56.
3. Slide the controller fan module out of the enclosure.



**Figure 56** Controller Fan Module Removal and Replacement

### Installing the Controller Fan Module.

1. Slide the new module into the slot and tighten the screw. The captive screw is spring-loaded and will not tighten unless it is inserted all the way into the chassis. If the screw keeps spinning without tightening, push on the front of the module until it snaps into place, and then tighten the screw.
2. Check the system LEDs on the controller fan module for the following behavior:
  - The green Power LED should be on and the Fan Fault LED should be off. If the Fan Fault LED turns on or the Power LED remains off, make sure that the controller enclosure fan module is inserted all the way into the chassis, and the thumbscrew is tight. If the Fan Fault LED remains on, it may indicate a problem with the new module.
3. Install the front cover. See "Front Cover Removal/Replacement" on page 243.

---

## Battery Backup Unit (BBU) Removal/Replacement

### Hot Swappable Component!

---

**Note** If the Fast Write Cache LED is on when the BBU is removed from the enclosure (or if the BBU fails), write caching will be disabled and the write cache data will be written to disk. However, if a power outage occurs prior to completing the cache write to disk, data may be lost. Therefore, make sure the Fast Write Cache LED is off before replacing the BBU.

---

### Removing the BBU


1. Remove the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 243.
2. Loosen the four screws (two on each side) on the BBU. See Figure 57.

---

**CAUTION** *The BBU weighs approximately 24 pounds so be prepared to support its weight when you pull it out of the chassis. The BBU is not as deep as other modules and disengages the enclosure sooner than you may expect.*

---

3. Grasp the two handles and slide the module out about two inches.
4. Grasp both sides of the module with both hands and slide it out of the enclosure.




**Figure 57** BBU Removal and Replacement

---

## Installing the BBU

1. Unpack the new BBU. Save the shipping material for transporting the used BBU to the disposal facility.
2. Fill in the following information on the “Battery Support Information” label on the front of the battery. See Figure 58.
  - a. Record the current date on the blank line next to “Date of Installation.”
  - b. Record the expiration date (two years from the current date) on the line next to “Replacement Date.”
3. Slide the new BBU into the slot and tighten the screws. See Figure 57.
4. Install the controller enclosure front cover. See "Front Cover Removal/Replacement" on page 243.
5. Allow the system to run up to 7 hours to properly charge the BBU. The Full Charge LEDs will flash while the BBU is charging. When properly charged, both Full Charge LEDs on the front of the BBU will remain on.



**Figure 58** Battery Information Label



- 
- Dispose of the old BBU.

**Note** Dispose of the used BBU according to local and federal regulations, which may include hazardous material handling procedures.

---

## Power Supply Fan Module Removal/Replacement

**Hot Swappable Component!**


---

**CAUTION** *Do not operate the enclosure without adequate ventilation and cooling to the power supplies. Operating the power supplies without proper cooling may damage their circuitry. You must shut down the controller enclosure to avoid overheating if you take more than 15 minutes to remove the Controller Fan Module and insert the new one.*

---

### Removing the Power Supply Fan Module

- Gain access to the rear of the disk array controller enclosure.
- Lift up on the ring/latch on the module to release the catch. See Figure 59.
- Slide the power supply fan module from the enclosure.



**Figure 59** Power Supply Fan Module Removal and Replacement

## Installing the Power Supply Fan Module

1. Slide the power supply fan module into the enclosure. The latch will snap down when the module is seated properly. If the latch remains up, lift up on the ring/latch and push in on the module until it snaps into place.
2. Check the module LEDs for the following behavior:
  - The green Fan Power LED should be on and the amber Fan Fault LEDs should be off. If the Power LED remains off or the Fault LED turns on, make sure the power supply fan module is inserted all the way into the chassis and locked in place. If the Fault LED remains on, the new power supply fan module may be malfunctioning, refer to Chapter 6, Troubleshooting.


## Power Supply Module Removal/Replacement

### ☑ Hot Swappable Component!


A power supply should be replaced as quickly as possible to avoid the possibility of the remaining supply failing and shutting down the disk array.

### Removing the Power Supply Module

1. Turn off the power switch and unplug the power cord from the failed power supply module. See Figure 60.
2. Lift up on the pull ring to release the latch. See Figure 61.
3. Slide the supply out of the enclosure.



**Figure 60** Removing Power from the Controller Enclosure



**Figure 61** Power Supply Module Removal and Replacement

### **Installing the Power Supply Module**

1. Slide the supply into the slot until it is fully seated and the latch snaps into place.
2. Plug in the power cord and turn on the power. See Figure 60.
3. Check the power supply module LED for the following behavior:
  - The Power LED should go on. Once the power supply is installed and operating, there may be a delay of up to several minutes before the Power Fault LED goes off.

---

## SCSI Cables

Replacing SCSI cables requires that the disk enclosure be shut down. Shutting down the enclosure will degrade the performance of the array during the replacement. When the replacement is completed and the disk enclosure is powered up, the array will perform a rebuild (since I/O has occurred to the array while the disk enclosure was powered off). Array performance will be reduced until the rebuild is complete.

To replace a SCSI cable, complete the following steps:

1. Open the disk enclosure front door.
2. Power the disk enclosure down by setting the power switch to off.

---

**CAUTION** *When replacing a SCSI cable, it must be disconnected from both the controller enclosure and the BCC module. Disconnect the cable from the controller enclosure first, then from the failed BCC. DO NOT disk connect the SCSI cable from the BCC first or errors may occur when the enclosure is powered on.*

---

3. Remove the SCSI cable from the controller enclosure connector.
4. Remove the SCSI cable from the BCC
5. Connect the SCSI cable to the BCC connector.
6. Connect the other end of the SCSI cable to the controller enclosure SCSI connector.
7. Disconnect all disk modules from the backplane connectors by releasing the locking levers and pulling module out, about one inch. This step is required to ensure the disk modules return to full operation when the enclosure is powered back up.
8. Power up the disk enclosure.
9. Reseat the disk modules one at a time:
  - a. Push the disk module in slot 0 into slot until it snaps into place.
  - b. Observe the Disk Fault LED, it should come on. When the Disk Fault LED goes out, insert the next disk module.
  - c. Repeat the above steps until all disk modules are re-seated.

---

Once the disk enclosure is powered up, check the status of the disk modules using one of the software management tools. Initially the disk modules status will be either “write failed” or “no\_response.” Eventually, all the disk modules should return to “replaced” status. Once this occurs, the disk array will perform a rebuild (a result of I/O occurring during the period the enclosure was powered off). The disk array will operate at reduced performance until the rebuild is completed.

---

# 8

# REFERENCE / LEGAL / REGULATORY

---

Models and Options . . . . . 258

---

PDU/PDRU Products . . . . . 263

---

Replaceable Parts. . . . . 264

---

A5277A/AZ Controller Enclosure Specifications . . . . . 266

---

A5294A/AZ Disk Enclosure Specifications . . . . . 270

---

Warranty and License Information . . . . . 274

---

Regulatory Compliance . . . . . 280

---

FCC Statements (USA Only) . . . . . 281

---

Reference / Legal /  
Regulatory

---

# System Requirements

## Host Systems

### HP-UX

**Table 31 Supported HP-UX Host Platform Information**

Supported Host	Boot Support on HP-UX 11.x?	Fibre Channel I/O Adapter
K-class	Yes	A3404A
V-class	Yes	A5158A, A3740A
L-class	Yes	A5158A, A3740A
D-class	Yes	A3591B
N-class	Yes	A5158A, A3740A
R-class	Yes	A3591B
T-class	No	A3636A
C-class	No	A5158A on HP-UX 11.x A3740A on HP-UX 10.20
A4xx-A5xx class	Yes	A5158A

### Windows NT 4.0 and Windows 2000

Any host running Windows NT 4.0 or Windows 2000.

## Supported Operating Systems

- HP-UX 11.0
- HP-UX 10.20
- Windows NT 4.0 (with Service Pack 4, 5, or 6)



- 
- Windows 2000

## **Fibre Channel Host Adapters**

### **HP-UX**

- K-class: A3404A (assy number J2389-60001), 1063 Mbps, short-wave, non-OFC
- D- & R-class: A3591A (number A3395-60001), 1063 Mbps, short-wave, non-OFC
- T-600 class: A3636A (number A3329-60107), 1063 Mbps, short-wave, non-OFC
- V-class: A3740A (number A3740-60001), 1063 Mbps, short-wave, non-OFC

### **Windows NT 4.0 and /2000**

See the *HP Storage Manager 60 User's Guide* for a list of supported host adapters.

---

## Models and Options

The HP SureStore E Disk Array FC60 consists of two products: the A5277A/AZ controller enclosure and the A5294A/AZ SureStore E Disk System SC10, or disk enclosure. Each of these products have their own options as indicated below.

### **A5277A/AZ Controller Enclosure Models and Options**

- **A5277A** is a field racked controller enclosure integrated by qualified service -trained personnel. This model can be ordered with up to six A5294A. disk enclosures.
- **A5277AZ** is a factory-racked controller enclosure integrated into the specified rack. This model must be ordered with from one to six A5294AZ disk enclosures.

Both models include the following components (except as indicated):

- Rackmount enclosure
- Two controller modules with 256 Mbytes of cache (single controller available)
- Two power supplies
- Two fan modules
- One battery backup unit (BBU)
- Two power cords (w/ ferrite bead)
- Mounting rail kits for HP cabinets
- Terminators for unused SCSI ports
- Two 1/2 EIA Filler Panels (A5277A only)
- User documentation

**Table 32** A5277A/AZ Product Options

Option	Description
<b>Controller Options (must order one option)</b>	
203	Single controller module with 256 Mbyte cache, one Media Interface Adaptor, and one controller slot filler module. Configured with HP-UX firmware.
204	Two controller modules with 256 Mbyte cache and two Media Interface Adaptors. Configured with HP-UX firmware.
205	Two controller modules with 256 Mbyte cache and two Media Interface Adaptors. Configured with Windows NT/2000 firmware and Windows NT NVSRAM. Includes Storage
304	Two controller modules with 512 Mbyte cache and two Media Interface Adaptors. Configured with HP-UX firmware.
305	Two controller modules with 512 Mbyte cache and two Media Interface Adaptors. Configured with Windows NT firmware.
<b>Host Connect Cable Options</b>	
0Z4	2-meter Fibre Channel cable
AFY	16-meter Fibre Channel cable
0Z5	50-meter Fibre Channel cable
0Z6	100-meter Fibre Channel cable
701	Replace 1.5M VHDCI SCSI cable with 5M VHDCI SCSI cable
<b>Operating System Support Option (No longer required)</b>	
UM4	Support for HP-UX 11.0

---

## A5294A/AZ Disk Enclosure SC10 Models and Options

Order the following product and options as required. Enter the following product and options as sub-items to the A5277A and A5277AZ products above.

- **A5294A disk enclosure** is a field racked Sure Store E Disk System SC10 integrated by a service-trained engineer. This product may be ordered in conjunction with A5277A controller enclosure. To order a disk system SC10 without integration into an array, order A5272A.
- **A5294AZ** Factory-racked Sure Store E Disk System SC10 (disk enclosure) integrated into the specified rack with the A5227AZ controller. This product may only be ordered in conjunction with A5277AZ controller enclosure. To order a disk enclosure SC10 without integration into an array, order A5272AZ

Both models include the following components (except as indicated):

- SC10 Rack-mount enclosure (accommodates 10, 1.6” or 1,” disk modules)
- Two power supplies
- Two fan modules
- Two power cords
- Two bus controller cards (BCC; with enclosure monitoring)
- Mounting rails for HP racks
- Two SCSI bus terminators
- User documentation
- 1/2 EIA filler panel (A5294AZ as required for proper rack appearance)
- One or two 2.0 meter SCSI cables - The number of cables included depends on the number of A5294A/AZ disk enclosures ordered:

<u>Number of A5294A/AZ per A5277A/AZ</u>	<u>Number of SCSI cables per A5294A/AZ</u>
1,2, or 3	2
4, 5, or 6	1

**Table 33** A5294A Custom Cabling Option

Option	Description
701	Delete one 2m cable included in A5294A product and add one 5m VHDCI SCSI cable for connection of A5277A to A5294A in a different rack

**Table 34** A5294A/AZ Storage Capacity Options

Option	Description
Note: All disk enclosures ordered with a single A5277A/AZ must have identical Storage Capacity Options.	
104	4, 9-Gbyte 10K rpm disk drives
108	8, 9-Gbyte 10K rpm disk drives
110	10, 9-Gbyte 10K rpm disk drives
204	4, 18-Gbyte 10K rpm disk drives
208	8, 18-Gbyte 10K rpm disk drives
210	10, 18-Gbyte 10K rpm disk drives
304	4, 36-Gbyte 10K rpm disk drives
308	8, 36-Gbyte 10K rpm disk drives
310	10, 36-Gbyte 10K rpm disk drives
504	4, 18-Gbyte 15K rpm disk drives
508	8, 18-Gbyte 15K rpm disk drives
510	10, 18-Gbyte 15K rpm disk drives

**Table 35** Supporting Software (CD ROMS)

Product	Description
B6191A	IPR Distribution CD-ROM
A5628A	<i>HP Storage Manager 60</i>

---

## Disk Array FC60 Upgrade and Add-On Products

Order the following parts to expand or reconfigure your original purchase:

**Table 36** Upgrade Products

Order No.	Description
A5276A	9.1-Gbyte disk drive module 10K rpm Ultra 2 LVD
A5282A	18.2-Gbyte disk drive module 10K rpm Ultra 2 LVD
A5633A	18.2-Gbyte disk drive module 15K rpm Ultra 2 LVD
A5595A	36.4-Gbyte disk drive module 10K rpm Ultra 2 LVD
A5622A	73.4-Gbyte disk drive module 10K rpm Ultra 2 LVD
A5278A	Add-on controller, includes Media Interface Adaptor (MIA) Must also order Option 002 for 256 Mbytes of cache
002	256 Mbyte cache option for A5278A
A5279A	Controller cache 512-Mbyte upgrade kit. Includes two 256-Mbyte DIMMs
A5306A	2 m, VHDCI to VHDCI SCSI Cable, M/M
A5307A	5 m, VHDCI to VHDCI SCSI Cable, M/M
A5308A	10 m, VHDCI to VHDCI SCSI Cable, M/M
A5250A	Rail kit for HP legacy cabinets: C2785A, C2786A, C2787A, A1896A, and A1897A
A5251A	Rail kit for HP Rack System/E: HP A4900A, A4901A, A4902A, J1500A, J1501A, J1502A
A5672A	Rail kit for Rittal 9000 racks
A3583A	2-meter Fibre Channel cable
A3531A	16-meter Fibre Channel cable
A3735A	50-meter Fibre Channel cable
A3736A	100-meter Fibre Channel cable
A5296A	VHDCI SCSI terminator (LVD/SE)
A5649A	Storage Partition Support (up to 8 partitions)

## PDU/PDRU Products

Hewlett-Packard offers the following PDUs and PDRUs, with US and international power options, for meeting electrical requirements:

**Table 37** PDU/PDRU Products

Order No.	Description
Supported on Original Racks	
E7676A	19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 10 C20 outlets
E7671A	19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 2 C19 & 6 C13 outlets
E7674A	19 inch, 100-240 V, 16 Amp, 1 C20 inlet, 1 C19 & 7 C13 outlets
E7679A	19 inch, 100-127 V, 16 Amp, 2 C20 inlets, 2 C19 outlets, switch accessory
E7680A	19 inch, 200-240 V, 16 Amp, 2 C20 inlets, 2 C19 outlets, switch accessory
E7681A	19 inch, 200-240 V, 30 Amp, L6-30P, 2 C19 & 8 C13 outlets, switch accessory
E7682A	19 inch, 200-240 V, 30 Amp, IEC-309, 2 C19 & 2 C13 outlets, switch accessory
E4452A	36 inch, 220 V, 16 Amp, L6-20P, 6 C-13 outlets
E4453A	36 inch, 220 V, 16 Amp, no plug, 6 C-13 outlets
E5933A	36-inch, 110-220 V, 16 Amp, UPS, IEC-320, 10 C-13 outlets
E4456A/B	60 inch, 220 V, 16 Amp, L6-20P, 10 C-13 outlets
E4457A/B	60 inch, 220 V, 16 Amp, no plug, 10 C-13 outlets
E5930A	60 inch, 110-220 V, 16 Amp, UPS, IEC-320, 10 C-13 outlets
E5931A	60 inch, 220 V, 16 Amp, UPS, LP-30P, 10 C-13 outlets
E5932A	60 inch, 220 V, 16 Amp, UPS, no plug, 10 C-13 outlets
E7677A	Switch panel accessory for PRU
E7678A	Switch control jumper cord for PRU

---

## Replaceable Parts

### A5277A/AZ Controller Enclosure Replaceable Parts

Table 38 Controller Enclosure Replaceable Parts

Part Number	Field Replaceable Units	Exchange Part #
A5278-60001	HP-UX <sup>1</sup> Controller Module (5v model <sup>2</sup> ) w/32 MB SIMM (no cache DIMMs) This part has been replaced by the A5278-60006.	A5278-69001
A5278-60006	HP-UX <sup>1</sup> Controller Module (3.5v model <sup>2</sup> ) w/32 MB SIMM (no cache DIMMs)	A5278-69006
A5635-60002	Windows NT <sup>1</sup> Controller Module w/32 MB SIMM (no cache DIMMs) Includes Windows NT NVSRAM settings	A5635-69002
A5278-60002	128 MB DIMM	n/a
5065-5201	256 MB DiMM	n/a
A5278-60004	16MB SIMM Module	n/a
A5277-60009	Battery Backup Module	A5277-69009
A5277-60003	Controller Fan Module	n/a

<sup>1</sup>Controller modules have operating-specific firmware installed and are not interchangeable. The HP-UX 3.5v controller (A5278-60006) can be upgraded to a Windows NT/2000 controller (A5635-60002) using the HP Storage Manager 60-NT software (A5628A). The HP-UX 5v controller (A5278-60001) cannot be upgraded for use on Windows NT/2000.

<sup>2</sup>The 5v HP-UX controller and the 3.5v HP-UX controller are interchangeable. The 5v model will be available until supplies are exhausted and then replaced by the 3.5v model.

---



Table 38 Controller Enclosure Replaceable Parts (cont'd)

Part Number	Field Replaceable Units	Exchange Part #
A5277-60004	Power Supply Modules	n/a
A5277-60002	Power Supply Fan Module	n/a
A5277-60001	Front Door Assembly	n/a
5021-1121	Terminator, SCSI, 68 pin, LVD	n/a
5064-2464	Media Interface Adapter (MIA)	n/a

## A5294A/AZ Disk Enclosure Replaceable Parts

Table 39 Disk Enclosure Replaceable Parts

Replacement Part Order No.	Part Description	Exchange Part Order No.
8120-6514	Power cord	n/a
A5236-60003	Fan	n/a
A5236-60023	Power supply	A5236-69023
A5272-67004	Bus Control Card (BCC)	A5272-69004
A5236-60009	Disk filler assembly	n/a
A5276-67001	9.1-Gbyte disk drive module 10K rpm Ultra 2 LVD	A5276-69001
A5282-67001	18.2-Gbyte disk drive module 10K rpm Ultra 2 LVD	A5282-69001
A5633-67001	18.2-Gbyte disk drive module 15K rpm Ultra 2 LVD	A5633-69001
A5595-67001	36.4-Gbyte disk drive module 10K rpm Ultra 2 LVD	A5595-69001
A5622-67001	73.4-Gbyte disk drive module 10K rpm Ultra 2 LVD	A5622-69001

---

## A5277A/AZ Controller Enclosure Specifications

### Dimensions:

---

Height	Width	Depth
6.75 inches (17.1 cm)	17.5 inches (44.5 cm)	24 inches (61 cm)

---

### Weight:

---

Component	Weight of Each (lbs)	Quantity	Subtotal (lbs)
Controller modules	6.6	2	13.2
Controller Fan	1.9	1	1.9
Battery	21.4	1	21.4
Power Supply	3.3	2	6.6
Power Supply Fan	1.5	1	1.5
Front Cover	2	1	2
Chassis	31.6	1	31.6
Controller Enclosure - Total			78 lbs

---

---

## AC Power:

### AC Voltage and Frequency:

- 120 VAC (100 - 127 VAC), 50 to 60 Hz single phase
- 230 VAC (220 - 240 VAC), 50 to 60 Hz single phase
- Auto-ranging

### Current:

Voltage	Typical Operating Current	Maximum Operating Current	In-Rush Current
100 - 127 VAC	1.5 A	2.3 A	21.7 A
220 - 240 VAC	0.8 A	1.2 A	42.9 A

### Power Consumption:

Incoming Voltage AC RMS	Typical Power Consumption
100- 127 VAC	180 watts
200 - 240 VAC	180 watts

### Heat Output:

- 615 BTU/hr.

---

## Environmental Specifications

---

**Note** The HP SureStore E Disk Array FC60 has been tested for proper operation in supported Hewlett-Packard cabinets. If the disk array is installed in an untested rack configuration, care must be taken to ensure that all necessary environmental requirements are met. This includes power, airflow, temperature, and humidity. Failure to meet the required operating specifications may result in product failure.

---

### Operating Environment

For continuous, trouble-free operation, the high availability disk array should not be operated at its maximum environmental limits for extended periods of time. The recommended operating range provides a less stressful operating environment. Operating within the recommended operating range ensures maximum reliability for the disk array.

The following specifications were type-tested under controlled conditions. Hewlett-Packard maintains an active program of auditing production products to ensure these specifications remain true when products are again tested under the same conditions. The limits of these specifications do not represent the optimum for long, trouble-free operation and are specifically not recommended for maximum customer satisfaction. The recommended conditions are stated separately where appropriate.

- Operating temperature: 5° C to 40° C (41° F to 104° F)  
**Recommended:** 20° C to 25.5° C (68° F to 78° F)
- Maximum gradient: 20° C per hour (36° F per hour)
- Relative humidity: 10% to 80% RH at 28° C (wet bulb)  
**Recommended:** 30% to 50%
- Altitude: 3048 m (0 - 10,000 ft)

---

**Note** For continuous, trouble-free operation, the product should NOT be operated at its maximum environmental limits for extended periods of time. Operating within the recommended operating range, a less stressful operating environment, ensures maximum reliability.

---

---

### **Non-operating Environmental (shipping and storage):**

- Temperature: -40° C to 70° C (-40° F to 158° F)
- Maximum gradient: 20° C per hour (68° F per hour)
- Relative humidity: 10% to 90% RH @ 28° C (wet bulb)
- Altitude: 4572 m (0 - 15,000 ft)

### **Acoustics**

- Meets or exceeds all known international acoustics specifications for computing environments.

### **Battery Back-Up Module Disposal**

The BBU contains batteries that may be harmful to the environment if not disposed of properly. Please follow the local, state or country regulations regarding the proper disposal of these batteries.

---

## A5294A/AZ Disk Enclosure Specifications

### Dimensions:

Height	Width	Depth
5.91 in. (15.0 cm)	18.9 in. (48.0 cm)	27.2 in. (69.1 cm)

### Weight:

Component	Weight of Each (lbs)	Quantity	Subtotal (lbs)
Disk Drive (HH)	2.8	10	28
Fan	3.3	2	7
Power Supply	10.6	2	22
BCC	4.5	2	9
Midplane-Mezzanine	6	1	6
Door	2	1	2
Chassis	35	1	35
Total, Approx.			110 lbs. (50 kg.)

---

## AC Power:

### AC Voltage and Frequency:

- 100 - 127 VAC, 50 to 60 Hz single phase
- 220 - 240 VAC, 50 to 60 Hz single phase:

### Current:

Voltage	Typical Current	Maximum Current
100 - 127 VAC	4.8 a	6.5 a
220 - 240 VAC	2.4 a	3.2 a
50-60 Hz , single phase, power factor corrected		

### Power Consumption:

Incoming Voltage AC RMS	Typical Power Consumption
100- 127 VAC	347 watts
200 - 240 VAC	347 watts

### DC Power Output:

- Disk: +5 V and +12 V from power supply
- BCC: +5 V and +12 V from power supply, 3.3 V self-generated
- Fan: +12 V from power supply

### Heat Output:

- 2200 BTU/hr.

---

## Environmental Specifications

---

**Note** The HP SureStore E Disk Array FC60 has been tested for proper operation in supported Hewlett-Packard cabinets. If the disk array is installed in an untested rack configuration, care must be taken to ensure that all necessary environmental requirements are met. This includes power, airflow, temperature, and humidity. Failure to meet the required operating specifications may result in product failure.

---

### Operating Environment

For continuous, trouble-free operation, the high availability disk array should not be operated at its maximum environmental limits for extended periods of time. The recommended operating range provides a less stressful operating environment. Operating within the recommended operating range ensures maximum reliability for the disk array.

The following specifications were type-tested under controlled conditions. Hewlett-Packard maintains an active program of auditing production products to ensure these specifications remain true when products are again tested under the same conditions. The limits of these specifications do not represent the optimum for long, trouble-free operation and are specifically not recommended for maximum customer satisfaction. The recommended conditions are stated separately where appropriate.

### Temperature and Altitude:

- Operating temperature: 5° C to 40° C (50° F to 104° F)  
**Recommended:** 20° C to 25.5° C (68° F to 78° F)
- Maximum gradient: 20° C per hour (36° F per hour)
- Relative humidity: 20% to 80% noncondensing, max. wetbulb at 26° C  
**Recommended:** 30% to 50% noncondensing
- Altitude: 3000 m (10,000 ft)



---

**Note**

For continuous, trouble-free operation, the disk enclosure should NOT be operated at its maximum environmental limits for extended periods of time. Operating within the recommended operating range, a less stressful operating environment, ensures maximum reliability.

---

The environmental limits in a nonoperating state (shipping and storage) are wider:

- Temperature: -40° C to 70° C (-40° F to 158° F)
- Maximum gradient: 24° C per hour (43.2° F per hour)
- Relative humidity: 15% to 90% noncondensing
- Altitude: 4600 m (15,000 ft)

**Acoustics**

- Sound power: 6.4 Bels
- Sound pressure at operator's position: 56.2 dB(A)

---

## Warranty and License Information

### Hewlett-Packard Hardware Limited Warranty

HP warrants to you, the end-user Customer, that HP SureStore E Disk Array FC60 hardware components and supplies will be free from defects in material and workmanship under normal use after the date of purchase for *three years*. If HP or Authorized Reseller receives notice of such defects during the warranty period, HP or Authorized Reseller will, at its option, either repair or replace products that prove to be defective. Replacement parts may be new or equivalent in performance to new.

Should HP or Authorized Reseller be unable to repair or replace the hardware or accessory within a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price upon return of the HP SureStore E Disk Array FC60.

#### Replacement Parts Warranty

HP replacement parts assume the remaining warranty of the parts they replace. Warranty life of a part is not extended by means of replacement.

#### Items Not Covered

Your HP SureStore E Disk Array FC60 warranty does not cover the following:

- Products purchased from anyone other than HP or an authorized HP reseller
- Non-HP products installed by unauthorized entities
- Consumables, such as batteries
- Customer-installed third-party software
- Routine cleaning, or normal cosmetic and mechanical wear
- Damage caused by misuse, abuse, or neglect
- Damage caused by parts that were not manufactured or sold by HP
- Damage caused when warranted parts were repaired or replaced by an organization other than HP or by a service provider not authorized by HP

---

## Software Product Limited Warranty

The HP Software Product Limited Warranty will apply to all Software that is provided to you by HP as part of the HP SureStore E Disk Array FC60 for the NINETY (90) day period specified below. This HP Software Product Limited Warranty will supersede any non-HP software warranty terms that may be found in any documentation or other materials contained in the computer product packaging with respect to covered Software.

**Ninety-Day Limited Software Warranty.** HP warrants for a period of NINETY (90) DAYS from the date of the purchase that the Software will execute its programming instructions when all files are properly installed. HP does not warrant that the software will be uninterrupted or error free. In the event that this software product fails to execute its programming instructions during the warranty period, Customer's remedy will be a refund or replacement. Should HP be unable to replace the media within a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price (license fee) upon return of the product and all copies.

**Removable Media (if supplied).** HP warrants the removable media, if supplied, upon which this product is recorded to be free from defects in materials and workmanship under normal use for a period of NINETY (90) DAYS from the date of purchase. In the event the media proves to be defective during the warranty period, Customer's remedy will be to return the media to HP for replacement. Should HP be unable to replace the media with a reasonable amount of time, Customer's alternate remedy will be a refund of the purchase price upon return of the product and destruction of all other non-removable media copies of the software product.

Note that removable media does not include "hot-swap" hard drives, which are warranted under the HP Hardware Limited Warranty.

**Notice of Warranty Claims.** All warranty claims must be made during the applicable warranty period or within thirty (30) days after the expiration of the applicable warranty period. Any claim made after that time will not be eligible for warranty service but may be service under a separate HP support contract covering the product.

## Limitation of Warranty

HP does not warrant uninterrupted or error-free operation of an HP SureStore E Disk Array FC60.

---

This warranty extends only to the original owner in the original country of purchase and is not transferable. Consumables, such as batteries, have no warranty.

The above warranties will not apply to products from which serial numbers have been removed or to defects resulting from misuse (including operation of HP SureStore E Disk Array FC60 without covers and incorrect input voltage), unauthorized modification, operation or storage outside the environmental specifications for the product, in-transit damage, improper maintenance, or defects resulting from use of third-party software, accessories, media, supplies, consumables, or such items not designed for use with the product.

The HP warranty does not cover errors, malfunctions, or problems caused by or related to third-party products that are external to your HP SureStore E Disk Array FC60. HP reserves the right to charge for services where the cause of the error is found to be a non-certified, third-party product. In addition, HP reserves the right to charge for resolving problems caused by unauthorized personnel who have serviced your computer equipment.

**Limitation of liability and Remedies.** In no event will HP or its affiliates, subcontractors, or suppliers be liable for any of the following:

1. Damages for loss of data, or software restoration;
2. Damages relating to Customer's procurement of substitute products or services (i.e., "cost of cover");
3. Incidental, special, or consequential damages (including lost profits, loss of use, Customer downtime, cost of data recovery/re-creation), even if HP is informed of their possibility;
4. Third-party claims against Customer for losses or damages.

The HP SureStore E Disk Array FC60 is not specifically designed, manufactured or intended for sale as parts, components or assemblies for the planning, construction, maintenance, or direct operation of a nuclear facility. Customer is solely liable if Products or Support purchased by Customer are used for these applications.

HP's liability for damage to tangible property per incident is limited to the greater of \$300,000 or the actual amount paid for the product that is the subject of the claim, and for damages for bodily injury or death, to the extent that all such damages are determined by a court of competent jurisdiction to have been directly caused by a defective HP product.

**TO THE EXTENT ALLOWED BY LOCAL LAW, THE REMEDIES IN THIS WAR-**

---

**RANTY STATEMENT ARE CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES. EXCEPT AS INDICATED ABOVE, IN NO EVENT WILL HP OR ITS SUPPLIERS BE LIABLE FOR LOSS OF DATA OR FOR DIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL (INCLUDING LOST PROFIT OR DATA), OR OTHER DAMAGE, WHETHER BASED IN CONTRACT, TORT, OR OTHERWISE.**

Some Countries, states or provinces do not allow limitations on the duration of an implied warranty, so the above limitation or exclusion might not apply to you. The warranty gives you specific legal rights and you might also have other rights that vary from country to country, state to state or province to province.

**THE WARRANTY TERMS CONTAINED HERE, EXCEPT TO THE EXTENT LAWFULLY PERMITTED, DO NOT EXCLUDE, RESTRICT OR MODIFY AND ARE IN ADDITION TO THE MANDATORY STATUTORY RIGHTS APPLICABLE TO THE SALE OF THIS PRODUCT TO YOU.**

## **Hewlett-Packard Software License Terms**

The disk array hardware described herein uses Licensed Internal Code ("LIC"). The LIC, including any updates or replacements, any LIC utility software, and Supplier's software are collectively referred to as "Software". HP or its Suppliers own the copyrights in the Software and all copies of the Software. HP distributes the Software under license. The following Terms and Conditions shall govern the Software and Customer's use of the Software.

### **1. DEFINITIONS**

- a. "Software" is one or more programs capable of operating on a computer, processor, or controller which is either listed separately as a Software Product on the Price List, included with another Product on the Price List, or fixed in hardware and not removable in normal operation.
- b. "Use" means storing, loading, installing, executing, or displaying Software on a computer, processor, or controller, or making a copy of Software for archival or backup purposes only.
- c. "Products" include hardware, Software, options, documentation, accessories, supplies, spare parts, and upgrades on HP's Price List on the date HP receives Customer's order.

**License Grant.** HP grants you a license to Use one copy of the Software. "Use" means storing, loading, installing, executing or displaying the Software. You may not modify the

---

Software or disable any licensing or control features of the Software. If the Software is licensed for "concurrent use", you may not allow more than the maximum number of authorized users to Use the Software concurrently. You may not allow the Software to be used by any other party or for the benefit of any other party.

**Ownership.** The Software is owned and copyrighted by HP or its third party suppliers. Your license confers no title or ownership and is not a sale of any rights in the Software, its documentation or the media on which they are recorded or printed. Third party suppliers may protect their rights in the Software in the event of any infringement.

**Copies.** You may only make one copy of the Software for archival purposes or when copying is an essential step in the authorized Use of the Software on a backup product, provided that copies are used in no other manner and provided further that Use on the backup product is discontinued when the original or replacement product becomes operable. You must reproduce all copyright notices in the original Software on all copies. You may not copy the Software onto any public or distributed network.

**No Disassembly or Decryption.** You may not disassemble, decompile or reverse engineer the Software without HP's prior written consent. Where you have other rights under statute, you will provide HP with reasonably detailed information regarding any intended disassembly or decompilation. You may not decrypt the Software unless necessary for the legitimate use of the Software.

**Transfer.** You may not sell, sublicense or otherwise transfer the Software. Your license will automatically terminate upon any attempt to transfer of the Software. If you transfer, sell or otherwise dispose of the product upon which this Software is fixed, you must erase the Software before any such transfer or disposal. You will make reasonable efforts to protect the confidentiality of the Software.

**Termination.** HP may terminate your license upon notice for failure to comply with any of these License Terms. Upon termination or when you cease to use the Software, you must immediately destroy the Software, together with all copies and merged portions in any form.

---

## Restricted Rights Legend

1.) Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013.

Hewlett-Packard Company  
3000 Hanover Street  
Palo Alto, Ca 94304 U.S.A.

Copyright © 1997, 1998 Hewlett-Packard Company. All Rights Reserved.

2.) Customer further agree that Software is delivered and Licensed as "Commercial Computer Software" as defined in DFARS 252.227-7013, or "restricted computer software" as defined in FAR 52.227-19 c(1,2) if used, respectively, in the performance of a Department of Defense on non-Department of Defense U.S. Government contract.

c. Any term of this Agreement which is held to be invalid will be deleted, but the remainder of the Agreement will not be affected.

d. Neither party may assign any rights or obligations without prior written consent of the other party.

e. Customer who exports Products assumes responsibility for complying with applicable laws and regulations and for obtaining required export and import authorization. Customer will not export or re-export the Software, any copy, or any technical data in violation of any applicable laws or regulations.

f. Any disputes arising in connection with this agreement will be governed by the laws of California.

g. These HP Software License Terms supersede any previous communications, representations or agreements between the parties, whether oral or written, regarding transactions hereunder.

---

## Regulatory Compliance

### Safety Certifications:

- UL listed
- CUL certified
- TUV certified with GS mark
- Gost Certified
- CE-Mark

### EMC Compliance

- US FCC, Class A
- CSA, Class A
- VCC1, Class A
- BCIQ, Class A
- CE-Mark
- C-Tick Mark



---

## FCC Statements (USA Only)

The Federal Communications Commission (in 47 CFR 15.105) has specified that the following notice be brought to the attention of the users of this product.

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 when the equipment is operated 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. The end user of this product should be aware that any changes or modifications made to this equipment without the approval of the manufacturer could result in the product not meeting the Class A limits, in which case the FCC could void the user's authority to operate the equipment.

## IEC Statement (Worldwide)

This is a CISPR 22 Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

## CSA Statement (For Canada Only)

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement.

---

## VCCI Statement (Japan)

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

---

This equipment is in the Class A category information technology equipment based on the rules of Voluntary Control Council For Interference by Information Technology Equipment (VCCI). When used in a residential area, radio interference may be caused. In this case, user may be required to take appropriate

## Harmonics Conformance (Japan)

高調波ガイドライン適合品

## Class A Warning Statement (Taiwan)

警告使用者：這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

---

## Spécification ATI Classe A (France Seulement)

DECLARATION D'INSTALLATION ET DE MISE EN EXPLOITATION d'un matériel de traitement de l'information (ATI), classé A en fonction des niveaux de perturbations radioélectriques émis, définis dans la norme européenne EN 55022 concernant la Compatibilité Electromagnétique.

Cher Client,

Conformément à la Réglementation Française en vigueur l'installation ou le transfert d'installation, et l'exploitation de cet appareil de classe A, doivent faire l'objet d'une déclaration (en deux exemplaires) simultanément auprès des services suivants:

- Comité de Coordination des Télécommunications 20, avenue de Ségur-75700 PARIS
- Préfecture du département du lieu d'exploitation

Le formulaire à utiliser est disponible auprès des préfectures.

La déclaration doit être faite dans les 30 jours suivant la mise en exploitation.

Le non respect de cette obligation peut être sanctionné par les peines prévues au code des Postes et Télécommunications et celles indiquées dans la loi du 31 mai 1993 susvisée.

Arrêté du 27 Mars 1993, publié au J.O. du 28 Mars-ATI

## Product Noise Declaration (For Germany Only)

### Acoustic Noise Emissions:

- LpA: 45.0 dB (seeking)
- At bystander position per ISO 7779.
- All data are the results from type tests of the product configuration having the highest acoustic emissions: 12 disk modules.
- All other configurations have lower emission levels.

---

## **Geräuschemission (For Germany Only)**

- LpA: 45.0 dB (suchend)
- Am fiktiven Arbeitsplatz nach DIN 45635 T. 19.
- Die Daten sind die Ergebnisse von Typprüfungen an Gerätekonfigurationen mit den höchsten Geräuschemissionen: 12 Plattenlaufwerke.
- Alle andere Konfigurationen haben geringere Geräuschpegel.
- Für weitere Angaben siehe unter Umgebungsbedingungen.

---

## Declaration of Conformity

according to ISO / IEC Guide 22 and EN 45014

**Manufacturer Name:** Hewlett-Packard Company  
**Manufacturer Address:** Enterprise Storage Business Unit  
P.O.Box 15  
Boise, Idaho U.S.A. 83707

**Declares, that the product**

**Product Name:** SureStore E Disk Array FC60  
**Product Numbers:** A5277A, A5277AZ, A5635A  
**Product Options:** All

conforms to the following Product Specifications:

**Safety:** IEC 950: 1991 +A1+ A2 +A3 + A4 / EN 60950 (1992) + A1 + A2 + A3 + A4 + A11  
IEC 825-1: 1993 / EN 60825-1: 1994 Class 1 (Laser/LED) GB4943 (1995)

**EMC:** CISPR 22: 1993 / EN 55022 (1994) Class A  
EN 50082-1: 1992 - Generic Immunity, including:  
IEC 801-2: 1991: / prEN 55024-2 (1991): 4 kV CD, 8 kV AD  
IEC 801-3: 1984: / prEN 55024-3 (1991): 3 V/m  
IEC 801-4: 1988: / prEN 55024-4 (1992): 1 kV Power Lines, 0.5 kV Data Line  
IEC 801-5: 1992: / IEC 61000-4-5 (1995): 2 kV OC, 1 kV SC (Common Mode),  
1 kV Peak (Differential Mode)  
IEC 61000-3-2: 1997 / EN 61000-3-2 (1995): Harmonic Current  
IEC 61000-3-3: 1994 / EN 61000-3-3 (1995): Voltage Fluctuations and Flicker  
IEC 61000-4-6: 1996 / EN 61000-4-6 (1996): Conducted Immunity  
GB9254 (1988)

**Supplementary Information**

The product herewith complies with the requirements of the Low Voltage Directive 73/23/EEC and the EMC Directive 89/336/EEC and carries the CE Marking accordingly.

1.)This product was tested with Hewlett-Packard Unix server host computer system.



Boise Idaho U.S.A., 03/22/99

---

Dan T. Michauld / QA Manager

**European Contact: Your local Hewlett-Packard Sales and Service office or Hewlett-Packard GmbH,  
Department HQ-TRE, Herrenberger Straße 130, D-71034 Boblingen (FAX +49-7031-14-3143)**

Reference / Legal /  
Regulatory



---

## GLOSSARY

### **adapter**

A printed circuit assembly that transmits user data (I/Os) between the host system's internal bus and the external Fibre Channel link and vice versa. Also called an I/O adapter, FC adapter, or host bus adapter (HBA).

### **ArrayID**

The value used to identify a disk array when using Array Manager 60. The ArrayID can be either the disk array S/N, or an alias assigned to the disk array.

### **auto failover**

The process of transferring the I/O data path from a failed controller to the remaining operational controller in the disk array. Auto failover allows uninterrupted access to all user data with no downtime. The failed controller can be removed, then replaced by a new array controller, without any downtime, loss of data, or interruption to the host computer system.

### **battery backup unit (BBU)**

A hardware component providing up to 5 days (120 Hours) of backup power to the disk array should power be lost. The BBU power maintains the integrity of the contents of write cache, which is written to the disk media when power to the disk array is restored. A functioning BBU is required to enable disk array cache.

### **BCC module**

The BCC (bus controller card) module manages the disk modules in the disk systems. Each disk system contains two BCC controllers.

---

**bind**

The process of configuring unassigned disks into a LUN disk group. Disks can be bound into one of the following LUN disk groups: RAID 5, RAID 1 (single mirrored pair), RAID 0/1 (multiple mirrored pairs).

**bootware**

This controller firmware comprises the bring-up or boot code, the kernel or executive under which the firmware executes, the firmware to run hardware diagnostics, initialize the hardware and to upload other controller firmware/software from Flash memory, and the XMODEM download functionality.

**bus**

A physical connection between printed-circuit boards in a CPU or a disk array used to communicate data. See also *SCSI-2 bus*.

**caching**

The technique of temporarily storing data in RAM on the disk array controller to improve performance during an I/O. By using higher-speed RAM whenever possible, I/Os can be executed faster than accessing the disk.

**cascaded FC-AL hubs**

One FC-AL hub connected to another FC-AL hub to increase arbitrated loop distances. Cascaded hubs allow distances up to 500 meters between hubs, or between a hub and a device.

**channel**

One of the six SCSI busses connecting the disk array controller enclosure to the disk enclosures. Each channel is an independent LVD SCSI bus.

**channel:ID**

The two-part value used to identify a disk module within the disk array when using the array management tools. The channel value indicates the SCSI channel that the disk enclosure containing the disk is connected to. The ID value indicates the SCSI ID assigned to the slot in which the disk is installed.



---

**Class of Service**

The types of services provided by the Fibre Channel topology and used by the communicating port.

**controller**

A removable unit that contains an array controller.

**dacstore**

A region on each disk used to store configuration information. During the Start Of Day process, this information is used to configure controller NVSRAM and to establish other operating parameters, such as the current LUN configuration.

**data redundancy**

The technique of storing extra data along with the original data to protect against a disk failure. This extra, or redundant, data is used to rebuild the original data if a disk fails or becomes inaccessible. Data redundancy is achieved using one of two techniques: mirroring or parity.

**data transfer rate**

Data transfer rate is the speed at which data is transferred between a host computer system and a peripheral.

**degraded mode**

The operating mode a LUN enters when if a disk fails or becomes inaccessible. In degraded mode the LUN must recreate the data on the missing disk using either mirrored data or parity information. I/O performance typically suffers while the LUN is in degraded mode. In addition, the LUN is susceptible to a second disk failure while in degraded mode. The LUN will remain in degraded mode until the content of the failed disk is rebuilt to a global hot spare or a replacement disk.

**DIMM**

Dual-inline memory module.

---

**disk array controller**

A printed-circuit board with memory modules that manages the overall operation of the disk array. The disk array controllers manage all aspects of disk array operation, including I/O transfers, data recovery in the event of a failure, and management of disk array capacity. There are two controllers (A and B) in the disk array enclosure. Both controllers are active, each assuming ownership of LUNs within the disk array.

**disk group**

A collection of individual disks that share a common role in disk array operation. All disks on the disk array will be a member of one of the following disk groups: a LUN, hot spare, or unassigned.

**disk module**

The individual disk drive assemblies used in the disk array. Each module houses a single disk. Up to ten disk modules can be installed in each disk system.

**disk striping**

A performance-enhancing technique in which data is distributed in uniformly sized segments across all the disks in a LUN. Striping allows multiple sets of read/write heads to simultaneously execute an I/O, resulting in improved performance. RAID 5 LUNs use disk striping.

**disk stripe size**

In a LUN using disk striping, the stripe size is calculated by multiplying the stripe segment size times the number of disks in the LUN.

**disk system**

The enclosure that contains the disk modules used by the disk array. Each disk system includes its own BCC controllers, power supplies, and cooling fans. Up to six disk systems can be connected to the disk array.

**Fibre Host ID**

Numbers that identify each controller to the host or other devices on the Fibre Channel network.

---

**EPROM**

Erasable Programmable Read-Only Memory.

**fabric**

A Fibre Channel term that describes a crosspoint switched network, which is one of three existing Fibre Channel topologies. A fabric consists of one or more fabric elements, which are switches responsible for frame routing. A fabric can interconnect a maximum of 244 devices. The fabric structure is transparent to the devices connected to it and relieves them from responsibility for station management.

**FC-AL**

See *Fibre Channel Arbitrated Loop (FC-AL)*.

**fibre**

A generic Fibre Channel term used to cover all transmission media specified in the Fibre Channel Physical Layer standard (FC-PH), including optical fibre, copper twisted pair, and copper coaxial cable.

**fibre optic cable**

An optical fibre cable made from thin strands of dielectric material such as glass through which data in the form of light pulses is transmitted by laser or LED. Fibre optic cable is used for high-speed transmission over medium to long distances.

**Fibre Channel**

Logically, a bidirectional, full-duplex, point-to-point, serial data channel structured for high performance capability. Physically, Fibre Channel interconnects devices such as host systems and servers, FC hubs, and disk arrays through ports, called N\_Ports, in one of three topologies: a point-to-point link, an arbitrated loop, or a crosspoint switched network, which is called a fabric. Fibre Channel can interconnect two devices in a point-to-point topology, from two to 126 devices in an arbitrated loop, and up to 224 devices in a fabric switched topology. The disk array default topology is arbitrated loop.

Fibre Channel is a generalized transport mechanism that has no protocol or native I/O command set, but can transport any existing protocol, such as SCSI in Fibre Channel frames. Fibre Channel is capable of operating at speeds of 100 MB/s (full speed), 50 MB/s (half speed),

---

25 MB/s (quarter speed), or 12.5 MB/s (eighth speed) over distances of up to 100 m over copper media, or up to 10 km over optical links. The disk array operates at full speed.

### **Fibre Channel Arbitrated Loop (FC-AL)**

One of three existing Fibre Channel topologies in which two to 126 ports are interconnected serially in a single loop circuit. Access to the FC-AL is controlled by an arbitration scheme. The FC-AL topology supports all classes of service and guarantees in-order delivery of Fibre Channel frames when the originator and responder are on the same FC-AL. The disk array's default topology is arbitrated loop.

### **Fibre Channel Arbitrated Loop Hub**

A full-duplex, 1.063 Gb/s intelligent hub used in an FC-AL topology to increase the loop's reliability, the number of loop connections, and the distances between the host system(s) and disk array(s). A maximum of ten devices can be connected to each FC-AL hub.

### **Fibre Channel Protocol for SCSI (FCP)**

FCP defines a high-level Fibre Channel mapping layer (FC-4) using lower-level Fibre Channel (FC-PH) services to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target across the Fibre Channel link using Fibre Channel frame and sequence formats.

### **flushing cache**

The process of writing the contents of write cache memory to a disk. Flushing occurs at regular intervals during normal operation, and when power to the disk array is interrupted.

### **foreign controller**

Any controller other than the last one to complete the Start-Of-Day process in a given slot. At the completion of the SOD process, the identification of the controller is stored. If a different controller is installed in the slot, it is considered foreign. See *resident controller*.

---

**frame**

The smallest indivisible unit of application-data transfer used by Fibre Channel. Frame size depends on the hardware implementation and is independent of the application software. Frames begin with a 4-byte Start of Frame (SOF), end with a 4-byte End of Frame (EOF), include a 24-byte frame header and 4-byte Cyclic Redundancy Check (CRC), and can carry a variable data payload from 0 to 2112 bytes, the first 64 of which can be used for optional headers.

**FRU (Field Replaceable Unit)**

A disk array hardware component that can be removed and replaced by a customer or Hewlett-Packard service representative.

**global hot spare**

A disk that is powered up and electrically connected to a disk array but not used until a disk failure occurs. A global hot spare is a dedicated, on-line, backup disk that cannot be used to store user data during normal disk array operation. If any disk in a LUN fails, the disk array automatically begins rebuilding the failed disk's content on an available global hot spare. When the rebuild process completes, the LUN functions normally, using the global hot spare as a replacement for the failed disk. Up to six disks (one for each channel) can be assigned as global hot spares.

**hardware path**

See *primary disk array path*.

**HBA**

See *adapter*.

**high availability**

The technique of designing systems that remain operational even if a hardware or other type of failure occurs. The disk array achieves high availability by using redundant data and global hot spare disks to maintain access to user data in the event of a disk failure. Redundant hardware assemblies also ensure high availability in the event of a failure.

---

**host**

A processor that runs an operating system using a disk array for data storage and retrieval.

**hot swappable**

Hot swappable components can be removed and replaced while the disk array is online without disrupting system operation. Disk array controller modules, disk modules, power supply modules, and fan modules are all hot swappable components.

**I/O operation**

An operation initiated by a host computer system during which data is either written to or read from a peripheral.

**image (disk image)**

See *mirroring*.

**JBOD**

An acronym for Just a Bunch Of Disks, it implies an enclosure that houses disk drives that by themselves do not provide data redundancy. The disk systems on the disk array are JBOD devices, but when connected to the disk array controllers, they provide data redundancy through the use of RAID technology.

**link**

In Fibre Channel, two unidirectional fibres transmitting in opposite directions and their associated transmitters and receivers that serve as the communication media between nodes in a topology. A link is comparable to a bus in the SCSI protocol.

**logical unit (LUN)**

An acronym for Logical Unit Number, the SCSI term for logical disk unit. A logical unit is an addressable part of a SCSI target. The terms “logical unit number” and “LUN” are interchangeable.

The LUN is the basic storage entity created on the disk array. Each LUN appears to the host as a single hard disk drive. Multiple LUNs, each using different RAID levels, can be

---

created on the same disk array. A numeric value is assigned to a LUN at the time it is created.

### **LVD-SCSI**

Low voltage differential implementation of SCSI. Also referred to as Ultra2 SCSI.

### **LVM (Logical Volume Manager)**

The default disk configuration strategy on HP-UX. In LVM one or more physical disk modules are configured into volume groups that are then configured into logical volumes.

### **loop address**

The unique ID of a node in Fibre Channel loop topology, sometimes referred to as a Loop ID.

### **Major Event Log**

A more detailed logging method implemented on HP08 and later firmware. Major event logging requires that the UTM be enabled.

### **MIA**

Media interface adapter, a device that allows you to connect fiber optic cable to a copper wire connector.

### **mirrored disks**

A mirrored disk is an exact copy of another disk. Mirrored disk are used when implementing the data redundancy technique of mirroring. See *mirroring*.

### **mirroring**

A technique which creates data redundancy by maintaining a duplicate copy of all data stored on the disk array. The mirror copy of data is available if the disk containing the original data fails or becomes inaccessible. RAID 1 and RAID 0/1 LUNs use mirroring.

---

## **NVSRAM**

The disk array controller stores operating configuration information in this non-volatile SRAM (referred to as NVSRAM). The contents of NVSRAM can only be accessed or changed using special diagnostic tools.

## **path**

See *primary disk array path* or *primary path*.

## **parity**

A data protection technique that provides data redundancy by creating extra data based on the original data. Parity is calculated on each write I/O by doing a serial binary exclusive OR (XOR) of the data segments in the stripe written to the data disks in the LUN. Parity is used by RAID 5 LUNs to reconstruct data from a failed disk.

## **peripheral device addressing (PDA)**

The addressing technique used by the host to address the disk array controllers. Because it is limited to addressing only eight LUNs, PDA has been superseded by VSA as the technique for addressing disk array LUNs.

## **point-to-point**

One of three existing Fibre Channel topologies, in which two devices are directly connected by a link with no fabric, loop, or switching elements present.

## **port**

The hardware entity that connects a device to a Fibre Channel topology. A device can contain one or more ports.

## **primary disk array path** or **primary path**

The main data path used for host I/Os to a LUN. The primary path can be set using LVM. Because the disk array has two controllers, one controller acts as the primary path, and the other acts as the alternate path. If a failure occurs in the primary data path, the alternate path is used, maintaining access to the LUN.



---

## **PROM (Programmable Read-Only Memory)**

SP-resident boot code that loads the SP microcode from one of the disk array's database drives when the disk array is powered up or when an SP is enabled.

## **RAID**

An acronym for "Redundant Array of Independent Disks." RAID was developed to provide data redundancy using independent disk drives. RAID is essentially a method for configuring multiple disks into a logical entity (LUN) that appears to the host system as a single, contiguous disk drive. RAID uses features such as data striping, disk mirroring, and parity to implement data redundancy. The RAID levels supported by the disk array include *RAID 1*, *RAID 3*, and *RAID 5*.

### **RAID 0**

A RAID level that provides improved performance through the use of data striping, but does not provide any data redundancy. RAID 0 should only be used for non-critical data storage.

### **RAID 1**

A RAID level in which the LUN uses a single mirrored pair of disks. One disk serves as the data disk, and the other serves as the mirror disk.

### **RAID 0/1**

A RAID configuration in which the LUN uses both mirroring for redundancy, and disk striping for performance. Half the disks serve as the data disks, and half serve as the mirror disks. Up to 30 disks can be included in a RAID 0/1 LUN.

### **RAID 5**

A RAID configuration in which the LUN uses data parity to for redundancy, and disk striping for performance. RAID 5 is efficient in its use of disk space, but suffers a performance penalty when performing write I/Os. Up to six disks (one per enclosure) can be used in a RAID 5 LUN.

### **rebuild**

The process of reconstructing the data that was on a failed disk onto another disk, typically a global hot spare. The disk array reconstructs the data using mirrored data or

---

parity information, depending on the RAID level of the LUN. Until a rebuild is complete, the disk array is operating in degraded mode and is vulnerable to a second disk failure.

### **reconstruction**

See *rebuild*.

### **resident controller**

The last controller to complete the Start-Of-Day process in a given slot. At the completion of the SOD process, the identification of the controller is stored. This controller remains the resident controller until another controller completes the SOD process in the same slot. See *foreign controller*.

### **SAM (System Administration Manager)**

A host-based system configuration and management utility.

### **SCSI**

An acronym for “Small Computer System Interface”, SCSI is an industry-standard protocol for connecting peripherals and hosts over a bus topology.

### **SCSI-2 bus**

A bus that complies with the SCSI standard. The six channels that connect the disk array controller system to the disk systems are SCSI 2 busses. Each disk system has two internal SCSI-2 busses that can be configured as independent busses that connect five disk slots to each BCC controller, or as a single internal bus.

### **SCSI ID**

A unique number assigned to each device connected to a SCSI bus. This number is used by the *HBA* to address each device on the bus. Each disk in a disk system is assigned a SCSI ID of 0 - 9.

### **secondary disk array path or secondary path**

See *primary disk array path* or *primary path*.

---

**SIMM (Single In-line Memory Module)**

A memory module that provides the local storage (cache) for an SP. An SP must have at least two 4-MB memory modules to support the storage system cache.

**Start Of Day (SOD)**

The initialization process used by the disk array controllers to configure itself and establish various operating parameters. Each controller runs its own SOD process. The SOD process occurs following a power on reset, or following the insertion of a controller.

**status light**

Lights on the front panel of each disk module, power module, and fan used to indicate the operating status of the hardware.

**STM (Support Tools Manager)**

A host-based tool used for disk array configuration, management, and diagnostics.

**stripe boundary crossing**

In disk striping, if the stripe segment size is too small for a single I/O operation, the operation requires access to two stripes. Called a stripe boundary crossing, this event reduces I/O performance.

**stripe segment size**

The amount of information simultaneously read from or written to each disk in a LUN using disk striping. The default stripe segment size is set to the same value used for cache page size (4K or 16K). The stripe segment size is configurable and can affect I/O performance. If the segment size for the LUN is set too small, it may result in many stripe boundary crossings.

**striping (disk striping)**

See *disk striping*.

**terminator**

An electrical connection at each end of the SCSI bus composed of a set of resistors (or possibly other components). Its function is to provide a pull-up for open-collector

---

drivers on the bus, and also impedance matching to prevent signal reflections at the ends of the cable. The SCSI bus requires termination at both ends of the bus. One end of the SCSI bus is terminated by the adapter's internal termination. The other end should have a terminator placed on the 68-pin high density SCSI connector on the last SCSI peripheral. If this device is not terminated, data errors may occur.

### **topology**

The physical layout of devices on a network. The three Fibre Channel topologies are fabric, arbitrated loop, and point-to-point. The disk array's default topology is arbitrated loop.

### **unbind**

Unbinding reverses the LUN binding process, deleting the LUN and returning the disk in the LUN to the unassigned disk group. Unbinding destroys all data on the LUN.

### **Uninterruptible Power Supply (UPS)**

An Uninterruptible Power Supply is a power supply that is capable of maintaining power even if the input ac mains supply loses its source of power.

### **Universal Transport Mechanism (UTM)**

A special LUN reserved for communication between the host and the disk array. SCSI commands are passed through the UTM. By default LUN 31 is used for the UTM. If the UTM is disabled, LUN 0 is used for the SCSI command path.

### **VHDCI**

Very high density cable interface.

### **volume set addressing (VSA)**

An enhanced technique for addressing disk array LUNs. VSA overcomes the eight LUN limit imposed by PDA, allowing all 32 LUNs on the disk array to be addressed by the host.

## A

A5628A, Array Manager 60 18

AC power specifications 271

disk enclosure 267

acoustics

controller enclosure 269

disk enclosure 273

adapter

Fibre Channel host 257

adding a disk module from another array 136

adding a global hot spare

using Array Manager 60 177

using SAM 154

using STM 195

adding capacity 135

with disk enclosures 137

with disk modules 135

with higher capacity disk modules 137

adding disk enclosures 137

adding disk modules 135

address

Fibre Channel host ID 93

addressing

full-bus example 105

split-bus example 105

alias, disk array

assigning using Array Manager 60 178

assigning using SAM 145

AM60Srvr

starting 122

amcfg

binding a LUN 170

changing LUN ownership 174

unbinding a LUN 173

amdsp

checking rebuild progress 186

listing disk arrays 173

amlog

viewing logs 190

ammgr

adding hot spare 177

assigning an alias 178

displaying parity scan status 188

halting parity scan 188

performing a parity scan 187

removing hot spare 177

resetting battery age 193

setting cache flush limit 182

setting cache flush threshold 181

setting cache page size 181

setting controller date and time 178

amutil

changing rebuild priority settings 186

flushing disk array log 192

locating disk modules 185

managing the UTM 179

Array Manager 60 18

adding a global hot spare 177

assigning an alias 178

binding a LUN 170

calculating LUN capacity 173

changing LUN ownership 174

changing rebuild priority settings 186

checking disk array status 163

checking rebuild progress 186

command summary 159

described 120

displaying parity scan status 188

flushing disk array log 192

halting parity scan 188

---

- identifying disk modules 173
- installing 122
- listing all disk arrays 169
- locating disk modules 185
- managing disk array logs 190
- managing log files 188
- performing a parity scan 187
- removing a global hot spare 177
- replacing a LUN 175
- rescanning for disk arrays 169
- resetting battery age 193
- running 122
- setting cache flush limit 182
- setting cache flush threshold 181
- setting cache page size 181
- synchronizing controller date and time 178
- unbinding a LUN 174

array *See disk array*

ArrayID 161

- listing for all arrays 169

assigning LUN ownership 128

**B**

backplane controller card. *See BCC module*

batteries *See battery backup module*

battery age

- resetting 193

battery backup module

- cache protection 44
- described 43
- life expectancy 44
- removal and replacement 246

battery charger *See battery backup module*

BCC module

- described 27
- troubleshooting 226

binding a LUN

- using Array Manager 60 170
- using SAM 148
- using STM 194

**C**

cabinets

- supported 16

cable installation

- Fibre Channel 107
- SCSI 98

cache

- upgrading to 512 MB 139

cache flush limit

- impact on performance 133
- setting 182

cache flush threshold

- impact on performance 132
- setting 181

cache page size

- impact on performance 133
- setting 181

calculating LUN capacity 173

capacity

- adding 135
- managing 123
- maximum 57

changing LUN ownership 174

channel number

- disk module 125

channel:ID

- described 161

checking disk array status

- using Array Manager 60 163
- using SAM 141
- using STM 194

cntrlrID

- described 161

command summary

- Array Manager 60 159

command-based interface *See cstm*

configuration

- switches 89

configuring LUNs 123

controller date

- synchronizing with host 178

controller enclosure

- 
- acoustics 269
  - dimensions 266
  - environmental 268
  - front cover 35
  - host ID 37
  - LED locations 115
  - models 258
  - power down 41
  - power supply fan module
    - described 41
  - power supply modules
    - described 40
  - specifications 266
  - voltage 267
  - weight 266
  - controller enclosure modules
    - described 32
  - controller fan module
    - described 38
    - removal and replacement 244
  - controller memory modules
    - DIMM 38
    - SIMM 38
  - controller module
    - described 36
    - interface connectors 37
    - LEDs, described 37
  - controller time
    - synchronizing with host 178
  - cstm
    - described 203
  - current
    - disk enclosure 271
  - D**
  - device logs
    - STM 203
  - dimensions
    - controller enclosure 266
    - disk enclosure 270
  - DIMM 38
  - disabling disk WCE 183
  - disk array
    - capacity management features 50
    - configuration planning 53
    - described 16
    - high availability features 45
    - high availability planning 55
    - options 260
    - power input specifications 271
    - power-down sequence 116
    - power-up sequence 109
    - upgrade and add-on products 262
    - ventilation 249
  - disk array capacity
    - maximum 57
  - disk array configurations
    - five disk enclosure, high availability and performance 74
    - five disk enclosure, maximum capacity 76
    - four disk enclosure, high availability and performance 70
    - four disk enclosure, maximum capacity 72
    - one disk enclosure, non-high availability 60
    - recommended 59
    - six disk enclosure, high availability and performance 78
    - six disk enclosure, maximum capacity 80
    - three disk enclosure, high availability and performance 66
    - three disk enclosure, maximum capacity 68
    - two disk enclosure, high availability and performance 62
    - two disk enclosure, maximum capacity 64
  - disk array logs
    - managing 190
    - viewing 190
  - disk array management software 18
  - disk array status
    - checking using Array Manager 60 163
    - checking using SAM 141
    - checking using STM 194
  - disk drives *See disk modules*
  - disk enclosure

---

- acoustics 273
- adding 137
- current 271
- described 21
- dimensions 270
- heat output 271
- hot swappable modules 25
- operating with empty slot 26
- operation features 23
- options 260
- power 271
- power input specifications 267
- power on switch 110
- specifications 270
- status indicators 23
- status indicators, described 23
- voltage 271
- weight 270

- disk enclosure modules
  - described 25
- disk filler
  - described 25
- disk module
  - adding 135, 232
  - adding or replacing 232
  - channel number 125
  - described 25
  - enclosure number 126
  - identifying in array 125
  - identifying using Array Manager 60 173
  - moving 136
  - removal and replacement 232
  - replacing with higher capacity modules 137
  - SCSI ID 125
  - slot number 126
- disk *See disk module*
- disk system *See disk enclosure*
- disk WCE
  - disabling 183
  - enabling 184
- drive lockout 233

## E

- EMC compliance 280
- EMS hardware event monitoring 19, 214
- enabling disk WCE 184
- enclosure number 126
- environmental specifications 268, 272
- evaluating performance 131
- event messages, contents of 215
- event notification
  - on HP-UX 215
- expanding storage capacity
  - overview 57
- Expert Tool 207
  - menu options 208
  - using menu mode 208
  - using X Windows 207

## F

- factory default configurations *See recommended array configurations*
- fan
  - troubleshooting 227
- fan module
  - troubleshooting 226
- fan module, disk enclosure
  - described 29
  - fault detection 29
  - removal and replacement 238
- fast write cache LED 38
- Fibre Channel
  - controller modules 37
  - host connections 37
  - hub connection 53
  - throughput 53
- Fibre Channel cable
  - installation 107
- Fibre Channel connection 53
- Fibre Channel host ID
  - addressing 93
- filler module
  - adding or replacing 232
  - removal



---

Firmware Update Tool 209

firmware, BCC

troubleshooting 228

flushing disk array log 192

front cover, controller enclosure 35

removal and replacement 243

full-bus mode

setting 92

switch 92

## G

global hot spare

adding using Array Manager 60 177

adding using SAM 154

adding using STM 195

removing using Array Manager 60 177

removing using SAM 156

removing using STM 195

## H

hardware event monitoring *See EMS*

*hardware event monitoring*

heat output

controller enclosure 267

disk enclosure 271

high availability

features 19, 45

planning 55

host

connections 37

host ID

controller modules 37

hot spare *See global hot spare*

hot swappable modules 16

disk enclosure 25

## I

I/Os per second

performance 56

identifying disk modules 125

using Array Manager 60 173

increasing capacity

of a LUN 126

of the disk array 135

information log

STM 206

Information Tool 204

interpreting information log 206

using menu mode 205

using X windows 204

installing Array Manager 60 software 122

interpreting status indicators 145

## L

LEDs

controller enclosure normal operation 114,  
221

controller module, described 37

disk enclosure normal operation 112

disk enclosures, described 23

fast write cache 38

functions 223

power supply module 241

status 225

listing all disk arrays 169

locating disk modules

using Array Manager 60 185

using SAM 146

using STM 146, 196

log files

managing 188

logs

managing 190

loop ID *See Fibre Channel host ID*

losing LUN 0 222

LUN

assigning ownership 128

binding using Array Manager 60 170

binding using SAM 148

binding using STM 194

calculating capacity using Array Manager  
60 173

changing ownership using Array Manager  
60 174

---

- configuring 123
- replacing using Array Manager 60 175
- selecting disks for 124
- selecting RAID level for 128
- setting stripe segment size 130
- unbinding using Array Manager 60 174
- unbinding using SAM 152
- unbinding using STM 195

LUN 0, losing 222

## M

- management tools 18, 120
- managing disk array capacity 123
- managing LUNs
  - using Array Manager 60 170
  - using SAM 148
  - using STM 194
- menu-based interface *See mstm*
- MIA
  - installation 108
- missing LUN 0 222
- models and options
  - controller enclosure 258
- modules
  - controller enclosure 32
  - disk enclosure 25
- moving disk modules
  - from one array to another 136
  - within an array 136
- mstm
  - described 201

## N

- noise level 269

## O

- operating environment 272
  - specifications 268
- operating system
  - support 256
- operating system support 18
- overheating

- controller enclosure 38

## P

- parity scan 187
  - displaying status of 188
  - halting 188
- PDU/PDRU
  - order numbers 263
  - troubleshooting 226
- performance
  - array configuration 55
  - I/Os per second 56
  - impact of configuration settings 131
  - SCSI channels 54
  - split-bus configuration 56
- planning a disk array configuration 53
- power
  - AC input specifications 271
  - AC input, disk enclosure 267
  - DC specifications 268, 272
  - disk enclosure 271
- power cable
  - troubleshooting 226
- power consumption
  - controller enclosure 267
  - disk enclosure 271
- power cord
  - installation 97
- power cord, controller enclosure
  - with power supply 40
- power down
  - controller enclosure 41
- power supply
  - troubleshooting 227
- power supply fan module, controller enclosure
  - described 41
  - removal and replacement 249
- power supply module
  - LEDs 241
- power supply module, controller enclosure
  - described 40
  - removal and replacement 251

- 
- power supply module, disk enclosure
    - described 30
    - load sharing 30
    - removal and replacement 240
  - power switch
    - controller enclosure 111
    - disk enclosure 110
    - disk enclosure, location 24
  - power, controller enclosure
    - redundancy, controller enclosure 40
  - power-down sequence 116
  - powering up the disk array 109
  - power-up
    - controller 218
    - LEDs 218
    - self tests 218
  - products
    - add-on & support 262
  - proper ventilation 249
  - R**
  - rack
    - Rittal 16
  - racks
    - supported 16
  - RAID
    - supported levels 19
  - RAID 0/1
    - five disk enclosure array 74
    - four disk enclosure array 70
    - six disk enclosure array 78
  - RAID 1
    - three disk enclosure array 66
    - two disk enclosure array 62
  - RAID 5
    - five disk enclosure array 76
    - four disk enclosure array 72
    - six disk enclosure array 80
    - storage capacity 57
    - three disk enclosure array 68
    - two disk enclosure array 64
  - RAID level comparison
    - application and I/O pattern performance
      - characteristics 49
      - data redundancy characteristics 46
      - general performance characteristics 48
    - RAID levels
      - tips on selecting 128
    - rebuild
      - changing priority settings 186
      - checking progress of 186
    - recommended disk array configurations 59
    - regulatory compliance 280
    - removal and replacement
      - battery backup module 246
      - controller fan module 244
      - fan module, disk enclosure 238
      - front cover, controller enclosure 243
      - power supply fan module, controller enclosure 249
      - power supply module, controller enclosure 251
      - power supply module, disk enclosure 240
      - SCSI cables 253
    - removing a global hot spare
      - using Array Manager 60 177
      - using SAM 156
      - using STM 195
    - replacing a LUN
      - using Array Manager 60 175
      - using SAM 152
    - replacing disk modules with higher capacity
      - modules 137
    - rescanning for disk arrays 169
    - Rittal rack 16
    - running Array Manager 60 122
  - S**
  - safety compliance 280
  - SAM 141
    - adding a global hot spare 154
    - assigning an alias 145
    - binding a LUN 148
    - checking disk array status 141

---

- interpreting status indicators 145
- locating disk modules 146
- removing a global hot spare 156
- replacing a LUN 152
- unbinding a LUN 152
- scanning for parity errors 187
- SCSI
  - drive connections 37
- SCSI cables 98
- SCSI cabling
  - full-bus configurations 99
  - removal and replacement 253
  - split-bus configurations 102
- SCSI channels
  - number of connections 54
  - performance 54
  - throughput 54
- SCSI connections 92
- SCSI ID
  - disk module 125
- segment size
  - impact on performance 131
  - tips on setting 130
- selecting a RAID level 128
- selecting disks for a LUN 124
- setting
  - cache flush limit 182
  - cache flush threshold 181
  - cache page size 181
  - configuration switches 89
  - controller date and time 178
  - stripe segment size 130
- SF21 230
- SF88 230
- SIMMs 38
- slot number 126
- software 19
  - management 18
- software tools 18
- specifications 270
  - controller enclosure 266
  - disk enclosure 270
- split-bus mode
  - performance 56
  - switch 92
- status
  - checking using Array Manager 60 163
  - checking using SAM 142
  - checking using STM 194
- status LEDs, controller enclosure
  - normal operation 114, 221
- status LEDs, disk enclosure
  - normal operation 112
- STM
  - adding a global hot spare 195
  - binding a LUN 194
  - checking disk array status 194
  - cstm 203
  - device logs 203
  - Expert Tool 207
  - Firmware Update Tool 209
  - Information Tool 204
  - locating disk modules 146, 196
  - managing with 194
  - mstm 201
  - removing a global hot spare 195
  - status indications 225
  - unbinding a LUN 195
  - users interfaces 199
  - xstm 200
- storage capacity
  - expansion 57
  - scalable 19
- storage capacity options 260
- stripe segment size
  - tips on setting 130
- supported operating systems 18
- SureStore E Disk System SC10 *See disk enclosure*
- switch configuration 89
  - disk enclosure 98
  - disk enclosure (tray) ID switch 89
  - Fibre Channel host ID address 93
  - full-bus 92

---

- split-bus 92
- tips on setting 89
- synchronizing
  - controller date and time 178
- system requirements
  - specifications 256

## T

- throughput
  - Fibre Channel 53
  - SCSI channels 54
- troubleshooting
  - checklist 217
  - event notification 215
  - isolating causes 226
  - power-up 218
  - table 226
  - Windows NT tools 214

## U

- unbinding a LUN
  - using Array Manager 60 174
  - using SAM 152
  - using STM 195
- universal transport mechanism 179
- upgrade and add-on products 262
- upgrading controller cache 139
- user interfaces for STM 199
- UTM 179

## V

- ventilation 249
- viewing logs 190
- voltage specifications
  - controller enclosure 267
  - disk enclosure 271

## W

- warranty information 274
- WCE
  - disabling on disks 183
  - disk 183, 184
  - enabling on disks 184

- weight
  - controller enclosure 266
  - disk enclosure 270
- Windows NT tools 18
- Windows NT troubleshooting tools 214
- write cache
  - battery backup 44

## X

- X Windows interface *See xstm*
- xstm
  - described 200

