IBM TotalStorage FAStT700 and Copy Services

Explore the new features of the FAStT700 hardware

Learn about FlashCopy, its components, and how to use them

Find out how Remote Volume Mirroring can protect data

Barry Mellish
Christian Demmer
Graham McGeoch
Ian Wallace

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM TotalStorage FAStT700 and Copy Services

July 2002

SG24-6808-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (July 2002)

This edition applies to the FAStT200, FAStT500, FAStT700, and Storage Manager version 8.2 that was available 28 June 2002.

Contents

Notices ix
Trademarks x

Preface xi
The team that wrote this redbook. xi
Become a published author xiii
Comments welcome. xiii

Chapter 1. Introduction to IBM TotalStorage FAStT700 1
1.1 Models. 2
1.1.1 The FAStT200 Storage Server 2
1.1.2 FAStT500 Storage Server. 3
1.1.3 FAStT700 Storage Server. 4
1.2 Managing FAStT700 5
1.2.1 New terminology 5
1.2.2 IBM FAStT Storage Manager 8.2 6
1.2.3 New features 7
1.2.4 Storage subsystem management methods. 8
1.2.5 Arrays and logical drives 8
1.2.6 Storage partitioning and heterogeneous hosts 9
1.2.7 Building blocks and sample configurations 10
1.2.8 Advanced copy services 11

Chapter 2. Hardware details 13
2.1 IBM TotalStorage FAStT700 Storage Server 14
2.1.1 Front view 15
2.1.2 Rear view 15
2.1.3 FAStT700 mini-hubs 16
2.1.4 Host-side connections. 17
2.1.5 Drive-side connections 17
2.1.6 FAStT700 diagnostic LEDs. 18
2.2 FAStT EXP700 expansion enclosure 21
2.2.1 Front view 21
2.2.2 Rear view 22
2.2.3 EXP700 ESM boards 22
2.2.4 EXP700 diagnostic LEDs 24
2.2.5 FAStT700 and EXP700 cabling 26
2.3 General cabling and connectivity issues 27
2.4 FAStT FC-2 Host Bus Adapter 29
2.5 TotalStorage SAN Switch F16 and F08 30

Chapter 3. Managing the FAStT Storage Server. 33
3.1 The FAStT Storage Manager software 34
3.1.1 Supported host systems 34
3.1.2 What's new 34
3.1.3 Storage subsystem management methods. 36
3.1.4 The FAStT Storage Manager Client 39
3.1.5 Event Monitor 42
3.2 Basic management 43

3.2.1 Arrays and logical drives 43
3.2.2 Storage partitioning 45
3.2.3 FlashCopy 48
3.2.4 Remote Volume Mirroring (RVM) 49
3.3 Advanced management 51
3.3.1 Data protection 51
3.3.2 Tuning and performance 57
3.3.3 FAStT Utilities 58
3.3.4 Script Editor and command line interface 58
3.3.5 Diagnostics 64

Chapter 4. Step-by-step procedures for the FAStT Storage Server 65
4.1 Driver and host software installation 66
4.1.1 Microsoft Windows NT 4.0 and Windows 2000 installation 66
4.1.2 Linux (RedHat 7.1 and 7.2, SuSE 7.3, and TurboLinux 7.0) 70
4.1.3 Novell NetWare 5.1 and 6.0 76
4.1.4 AIX 4.3.3 and 5.1 79
4.1.5 HP-UX v11 79
4.1.6 Solaris (SPARC) 2.6, 2.7, and 2.8 83
4.1.7 Network attached management 85
4.2 Configuring the FAStT Storage Server 85
4.2.1 Planning the configuration 85
4.2.2 Starting the FAStT Storage Manager Client 86
4.2.3 Updating the controller microcode 88
4.2.4 Updating the drive and ESM board microcode 94
4.2.5 Initial configuration steps 96
4.2.6 Defining hot spare drives 98
4.3 Creating arrays and logical drives 100
4.4 Configuring storage partitioning 105
4.5 Finishing the initial configuration 111
4.5.1 Monitoring and alerting 111
4.5.2 Saving the subsystem profile 113
4.6 Maintenance, tuning, diagnostics, and troubleshooting 114
4.6.1 Component properties 114
4.6.2 Modification priority 116
4.6.3 Cache settings 117
4.6.4 Media scanning 119
4.6.5 Loading and saving the configuration 121
4.6.6 The major event log 121
4.6.7 Changing ownership of a logical drive 122
4.6.8 Maintaining arrays 123
4.6.9 Maintaining logical drives 125
4.6.10 Handling premium features 127
4.6.11 Network setup of the controllers 129
4.6.12 Resetting the controllers 131
4.7 The FAStT Management Suite Java 133
4.7.1 Overview of the FAStT MSJ 133
4.7.2 Installing the FAStT MSJ 134
4.7.3 Using the FAStT MSJ GUI 135

Chapter 5. FlashCopy and the FAStT Storage Server 139
5.1 FlashCopy: How it works 140
5.1.1 Creating a FlashCopy logical drive 140

5.1.2 Disabling and recreating a FlashCopy logical drive 141
5.1.3 FlashCopy parameters 142
5.1.4 Estimating FlashCopy repository logical drive capacity 145
5.1.5 Estimating FlashCopy repository life. 146
5.1.6 Increasing the capacity of a FlashCopy repository logical drive 147
5.2 Step-by-step guide to FlashCopy 148
5.2.1 Checking the status of the Flash Copy premium feature 148
5.2.2 Enabling the FlashCopy premium feature. 150
5.2.3 Creating a FlashCopy drive. 150
5.2.4 Mapping a FlashCopy drive to a host 161
5.2.5 Viewing the FlashCopy drive status 164
5.2.6 Disabling a FlashCopy logical drive 169
5.2.7 Re-creating a FlashCopy logical drive 171
5.2.8 Resizing a FlashCopy repository drive 172
5.2.9 Deleting a FlashCopy drive. 174
5.3 Scenarios using scripts 176
5.3.1 Windows 2000 basic disks 177
5.3.2 Windows 2000 dynamic disks. 183
5.3.3 Linux disks 188

Chapter 6. Remote Volume Mirroring 191
6.1 Introduction to the Remote Volume Mirror option 192
6.2 Primary and secondary logical drives 192
6.3 Mirror repository 193
6.4 Mirror relationships 193
6.5 Logical drive limits. 195
6.6 Fabric configuration. 196
6.6.1 Distance limits. 196
6.6.2 Switch zoning 197
6.7 Data replication 197
6.8 Performance considerations 198
6.9 Loss of connection to secondary drive 199
6.10 Logical drive ownership. 200
6.11 FlashCopy and Remote Volume Mirroring 201
6.12 Supported hardware and operating systems 201
6.13 Supported configurations 201
6.14 Using Remote Volume Mirroring. 202
6.14.1 Enabling the Remote Volume Mirror premium feature option. 202
6.14.2 Remote Volume Mirror status 203
6.14.3 Enabling and activating the Remote Volume Mirror option. 204
6.14.4 Creating Remote Volume Mirror relationships 208
6.14.5 Viewing mirror relationships 212
6.14.6 Changing the mirror synchronization priority. 218
6.14.7 Removing mirror relationships 218
6.14.8 Mapping secondary mirror drives 219
6.15 Disaster recovery 219
6.15.1 Disaster plan 219
6.15.2 Key file backups and availability 219
6.15.3 Recovery from a switch failure 220
6.15.4 Recovery from a storage subsystem failure 220
6.15.5 Entire site failover 220
6.15.6 Recreating a mirror relationship 221
6.16 Reversing the roles of the primary and secondary drives 221

6.16.1 Changing the secondary to the primary logical drive 222
6.16.2 Changing the primary to the secondary logical drive 226
6.17 Remote Volume Mirroring solution design 226
6.17.1 Solution 1: Simple department with minimum redundancy 227
6.17.2 Solution 2: Intersite with redundant fabric 228
6.17.3 Solution 3: Intersite with FlashCopy drives and tape backup 229

Chapter 7. Migrating from an earlier version of FAStT Storage Manager 231
7.1 Prerequisites for migration 232
7.2 Updating the FAStT host software 232
7.2.1 Updating from a previous version of Storage Manager. 232
7.2.2 Updating Storage Manager in a Microsoft Cluster (MSCS) environment 233
7.2.3 Upgrading from Windows NT 4.0 to Windows 2000 234
7.2.4 Upgrading from Windows NT4 to Windows 2000 in a cluster setup 234
7.3 Updating the FAStT controller microcode 236

Chapter 8. FAStT700 maintenance 249
8.1 Performance monitoring and tuning 250
8.1.1 Performance Monitor 250
8.1.2 Cache parameters 252
8.1.3 Controller ownership 254
8.1.4 Segment size 254
8.1.5 RAID levels 254
8.1.6 Logical drive modification priority 255
8.1.7 Remote Volume Mirroring 255
8.2 Error reporting 256
8.2.1 Problem notification 256
8.2.2 Event Viewer 257
8.2.3 Recovery Guru 258
8.2.4 Missing logical drives 259
8.3 Diagnostics 260
8.3.1 Controller runtime diagnostics. 260
8.3.2 Read Link Status diagnostics 261
8.4 Common problems 263
8.4.1 Out of compliance error message in Recovery Guru 263
8.4.2 FAStT Client cannot connect to the FAStT Storage Server 263
8.4.3 Logical drives not on the preferred path 265
8.4.4 Storage subsystem is unresponsive 265
8.4.5 Obtaining the world-wide name of the host bus adapter. 265

Appendix A. Additional host-specific instructions for FlashCopy logical drives . . . 269
Operating system resources for additional instructions 270
Windows: Basic/regular disks 271
Process overview 271
Additional instructions for Windows 2000 basic disks. 272
Additional instructions for Windows NT regular disks 276
Windows: Dynamic disks 280
Process overview 280
Additional instructions for Windows 2000 dynamic disks 281
UNIX: Regular disks. 286
Process overview 286
Additional instructions for UNIX using regular disks 287
UNIX: Logical Drive Manager logical drives. 290
Process overview 290

Additional instructions for AIX: LVM Logical Logical Drives 291
Additional instructions for Solaris: Veritas Logical Drive Manager 295
Additional instructions for HP-UX: Logical Logical Drive Manager 299
NetWare 303
Process overview 303
Additional instructions for NetWare. 304

Appendix B. Critical event descriptions 307

Appendix C. Command line reference 317
Command line interface parameters 318
Commands and parameters. 321

Appendix D. FlashCopy example scripts 327
Windows 2000 basic disks 328
Initial flashcopy script 328
Disable flashcopy script 330
Recreate flashcopy script 333
Windows 2000 dynamic disks 335
Initial flashcopy script 335
Disable flashcopy script 337
Linux disks 342
Disable flashcopy script 342
Recreate flashcopy script 343

Related publications 345
IBM Redbooks 345
Other resources 345
Referenced Web sites 345
How to get IBM Redbooks 346
IBM Redbooks collections. 346

Index 347

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks(logo)™	Netfinity®	SP™
AIX®	Perform™	SP2®
Balance®	pSeries™	Tivoli®
FlashCopy®	Redbooks™	TotalStorage™
IBM®	SANergy™	Wave®
IBM eServer®	Sequent®	xSeries™

The following terms are trademarks of International Business Machines Corporation and Lotus Development Corporation in the United States, other countries, or both:

Lotus®	Notes	Word Pro®

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM Redbook introduces the IBM TotalStorage FAStT700 Storage Server and the new Copy Services functionality provided in the latest release of the Storage Manager software. It explains FAStT700 and discusses the hardware in detail. It examines FAStT Storage Manager, with an emphasis on the release of version 8.2, which introduces Remote Volume Mirroring. It offers a step-by-step guide to using the Storage Manager to create arrays, logical drives, and other basic management tasks.

This redbook includes details on:

- FlashCopy: How it is implemented and how to use it
- Remote Volume Mirroring: What it is and how to set up and control it
- Migrating from Storage Manager 7.10 to Storage Manager 8
- Performance monitoring and error reporting

Plus you'll find practical information on:

- Host-specific instructions for FlashCopy
- Critical event descriptions
- Command Line Reference
- FlashCopy example scripts

This book is intended for IBM technical professionals, Business Partners, and Customers who want to learn more about the capabilities of the advanced functions introduced with FAStT Storage Manager Software v8.2. It also targets those who have a FAStT storage subsystem and need detailed advice on how to configure it.

You can learn more about the FAStT storage subsystem in the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246. It introduces you to the FAStT range of storage products with particular emphasis on FAStT200.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Barry Mellish is a certified I/T Specialist and project leader at the International Technical Support Organization, San Jose Center. He has coauthored ten redbooks and has taught many classes on storage subsystems. He joined IBM UK 18 years ago. Prior to joining the ITSO, he worked as a Senior Storage Specialist on the Disk Expert team in EMEA.

Christian Demmer is an IT Specialist and Pre-sales Consultant in the IBM Server Group for IBM Germany. He provides technical support to IBM professionals, Business Partners, and Customers in IBM @server xSeries solutions. He holds a degree in mathematics from the University of Dortmund. His areas of expertise include Linux, Microsoft Windows NT and 2000, and high availability solutions. Christian also coauthored the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246.

Graham McGeoch is a Technical Consultant for Anix Computers, an IBM Business Partner in the United Kingdom. Graham spent 17 years in Control Engineering before moving into IT five years ago. He is a Microsoft Certified Systems Engineer (MCSE) in Windows NT4 and Windows 2000. He is also a Citrix Certified Administrator, IBM @server xSeries Certified

Systems Expert, and IBM Certified Systems Specialist in SAN implementation and TotalStorage networking. Anix Computers is a TotalStorage Solutions Centre and specializes in Enterprise Class solutions based on IBM @server pSeries and xSeries servers.

Ian Wallace is a Senior IT Specialist for IBM Australia. He provides pre-sales support for storage products to the IBM Enterprise Server Group, Business Partners, and customers. He has 13 years of experience in storage technologies, working in a variety of positions, such Customer Service Engineer, International Support, and Technical Consultant. He has an associate diploma in applied science from the Canberra Institute of Technology. His areas of expertise include storage networking, network data protection, and UNIX administration.

The redbook team (left to right): Ian Wallace, Christian Demmer, Barry Mellish, and Graham McGeoch

Thanks to the following people for their contributions to this project:

Michael Quillen
IBM Beaverton

John Murtagh
IBM Dallas

Melanie Jarvis
Christopher McCann
Bob Reiland
Mic Watkins
IBM Raleigh

Tai Chang
Funda Eceral
IBM San Jose

Dave Binning
Jannia Blackmon
Shawn Bramblett
Veronika Bronnikova
Obinna Eneh
Ernesto Eustaquio Jr.
Dave Goss
George Kattampallil
Shane Monaghan
Mario Rodriguez
Antoine Tounou
IBM Tucson

Dave Worley
LSI Logic

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. QXXE Building 80-E2
  650 Harry Road
  San Jose, California 95120-6099

1

Introduction to IBM TotalStorage FAStT700

This chapter introduces the IBM TotalStorage FAStT700 and where it fits into the FAStT family of storage products. This chapter also discusses the latest version of the controlling software – Storage Manager Version 8.2 – and the new advanced Copy Service functions – FlashCopy and Remote Volume Mirroring (RVM).

For further information on FAStT200 and FAStT500, read the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246.

> **Attention:** Some of the operating systems that are discussed in this section, at the time this redbook was written, are not formally supported. You must check the FAStT supported servers Web site prior to connecting hosts to the FAStT Storage Server or using the advanced functions. You can locate the FAStT700 supported servers Web site at:
>
> `http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm`

1.1 Models

The FAStT Storage Server is a RAID storage subsystem that contains Fibre Channel (FC) interfaces to connect both the host systems and the disk drive enclosures. The storage server provides high system availability through use of hot-swappable and redundant components. Although we focus on FAStT700, we give a brief overview of the other products. Figure 1-1 shows the FAStT evolution, which includes:

- **FAStT200 Storage Server**

  This solution is suitable for entry and medium level environments.

- **FAStT500 Storage Server**

  This storage server can support medium to high-end configurations with massive storage capacities and a large number of heterogeneous host systems. It offers a higher level of availability, performance, and expandability than FAStT200.

- **FAStT700 Storage Server**

  This storage server supports high-end configurations with up to 16 terabytes of storage and 64 heterogeneous host systems. It offers the highest level of performance and expandability in the range.

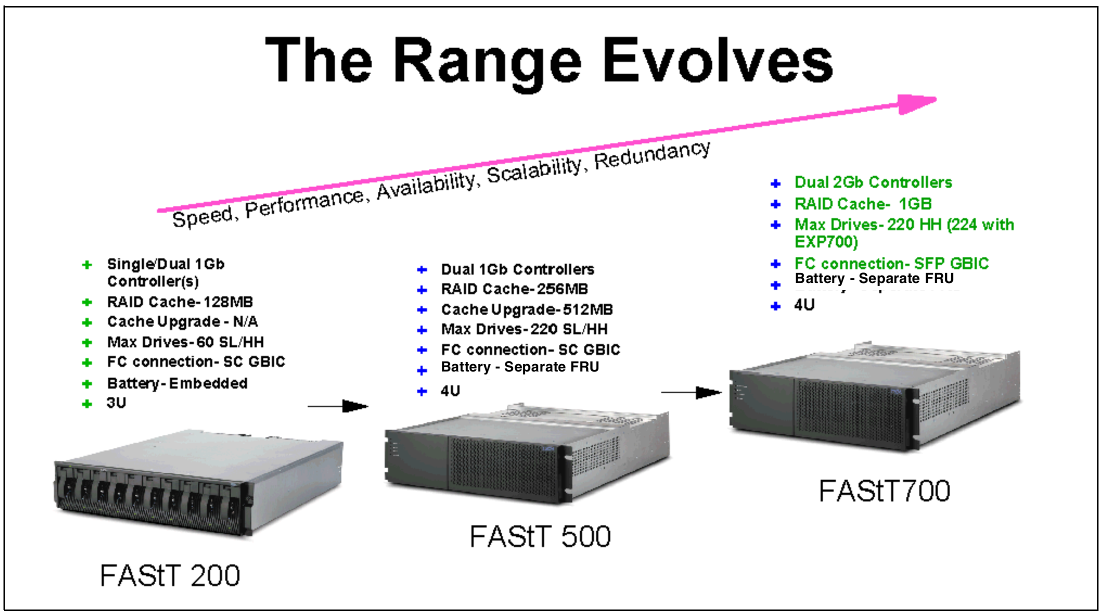

Figure 1-1 The FAStT evolution

1.1.1 The FAStT200 Storage Server

The FAStT200 Storage Server is a 3U rack-mountable Fibre Channel RAID storage system and disk drive enclosure. It targets the entry and midrange segment of the FC storage market. A typical use of the FAStT200 is in a 2-node cluster environment with up to 30 Fibre Channel disk drives attached to the storage server, as shown in Figure 1-2.

Figure 1-2 FAStT200

Two models are available:

- The FAStT200 Storage Server, with a single RAID controller
- The FAStT200 High Availability (HA) Storage Server, which contains two RAID controllers and can therefore provide higher availability

Both models feature hot-swap and redundant power supplies and fans and you can install up to 10 slim-line or half-high FC disk drives. If you need to connect more than 10 disks, you can use the EXP500 FC storage expansion enclosures. Each EXP500 can accommodate 10 additional disk drives and up to five EXP500s are supported on FAStT200. This means that the maximum supported number of disk drives is 60.

Note: Up to 126 FC devices may be configured in an FC loop. But arbitrated loop physical addresses (AL_PA) are also needed for controller ports and expansion enclosure monitoring. Therefore, you could connect up to 10 EXP500 with 110 disks total (10 disks in FAStT200, 100 disks in additional EXP500 units). This is not recommended from a performance perspective. Performance of the FAStT200 Storage Server is best with (up to) 30 disk drives

Each RAID controller unit contains 128 MB of battery-backup cache. Hardware components of FAStT200 and FAStT200 HA are explained in detail in the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246.

1.1.2 FAStT500 Storage Server

The FAStT500 Storage Server (Figure 1-3) is a 4U rack-mountable Fibre Channel RAID storage system. It provides the levels of performance, availability, and expandability needed to satisfy mid-range to high-end storage requirements. You typically use the FAStT500 Storage Server in an advanced cluster environment and with heterogeneous operating systems running on the host systems.

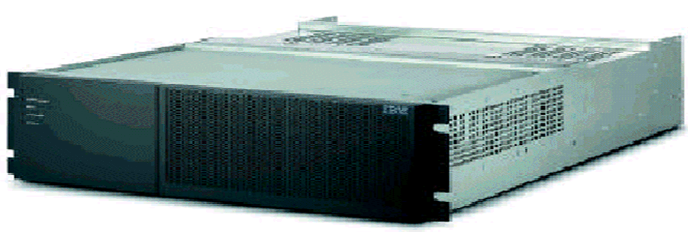

Figure 1-3 FAStT500

The storage server features two RAID controller units, redundant power supplies, and fans. All these components are hot-swappable, which assures excellent system availability. You use the EXP500 external storage expansion enclosures to install the FC disk drives, and you can connect up to 22 EXP500 enclosures to FAStT500. This means a total of up to 220 disk drives. The enclosures can be connected in a fully redundant manner, which provides a very high level of availability. On the host side FC connections, you can use up to four mini-hubs. This allows you to establish up to eight host connections without needing an external hub or switch.

For performance and availability, each RAID controller unit contains 256 MB of battery-backed cache. You can expand this amount further to a maximum of 512 MB per controller, giving 1 GB of cache within FAStT500.

The hardware components of FAStT500 are explained in detail in the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246.

1.1.3 FAStT700 Storage Server

The FAStT700 Storage Server, shown rack mounted in Figure 1-4, is the newest addition to the range. It is the same physical size as FAStT500 with new higher performance controllers. These new controllers are 2 Gbps and connect via new mini-hubs to the new FAStT FC-2 Host Bus Adapter (HBA) and the new 2109 F16 Fibre Channel switch to give a full 2 Gps fabric.

Like FAStT500, it attaches to up to 220 FC disks via 22 EXP500 expansion units or up to 224 FC disks via 16 EXP700 expansion units to provide scalability for easy growth (18 GB up to 16 TB). To avoid single points of failure, it also sports high availability features such as dual hot-swappable RAID controllers, two dual redundant FC disk loops, write cache mirroring, redundant hot-swappable power supplies, fans, and dual AC line cords.

Using the new Storage Manager Version 8.2, it supports FlashCopy, Dynamic Volume Expansion, and Remote Volume Mirroring with controller based support for up to 64 storage partitions. RAID levels 0, 1, 3, 5, and 10 are supported, and for performance, it includes a 2 GB battery-backed cache (1 GB per controller).

Figure 1-4 FAStT700

1.2 Managing FAStT700

To manage the FAStT700 Storage Server, use the IBM FAStT Storage Manager software. This software allows you to:

- Configure arrays and logical drives
- Assign your logical drives into storage partitions
- Replace and rebuild failed disk drives
- Expand the size of arrays and logical volumes
- Convert from one RAID level to another
- Perform troubleshooting and management tasks, like checking the status of FAStT Storage Server components, update the firmware of RAID controllers, and similar actions
- Configure and manage FlashCopy Volumes and Remote Volume Mirroring (FlashCopy and RVM are Premium Features that must be purchased)

At the time this redbook was written, the current publicly available version of Storage Manager was 8.2 with support on FAStT700 for Microsoft Windows, Linux, and Novell NetWare. Refer to the following Web site for the current support matrix, which changes constantly as new support is added:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

It is also possible to use the serial interface and a terminal emulation utility. However, this is only meant for advanced troubleshooting and management. It should only be used when other management methods fail and must be done under the supervision of IBM level 2 support.

1.2.1 New terminology

If you are used to using a previous version of the storage-management software, some of the terms you are familiar with have changed. It is important that you familiarize yourself with the new terminology. Table 1-1 provides a list of some of the changed terms.

Table 1-1 New terminology

Term used in previous versions	New term
RAID module	Storage subsystem
Drive group	Array
Volume	Logical drive
Auto Volume Transfer (AVT)	Auto Logical Drive Transfer (ADT)

It is important to understand the distinction between the two terms explained in the following sections when reading this book.

Management station

A management station is a system that is used to manage the storage subsystem. It can be attached to the storage subsystem in the following ways:

- Through a TCP/IP Ethernet connection to the controllers in the storage subsystem
- Through a TCP/IP connection to the host-agent software that is installed on a host computer, which in turn, is directly attached to the storage subsystem through the Fibre Channel input/output (I/O) path

Host computer

A host computer is a system that is directly attached to the storage subsystem through a Fibre Channel I/O path. This system is used to:

- Serve data (typically in the form of files) from the storage subsystem.
- Function as a connection point to the storage subsystem for a remote-management station.

Notes:

- The terms host and host computer are used interchangeably throughout this book.
- A host computer can also function as a management station.
- The term logical unit number (LUN) is sometimes used to describe a logical drive.

1.2.2 IBM FAStT Storage Manager 8.2

The Storage Manager V8.2 software consists of a suite of components that includes:

- Microsoft Virtual Machine (Windows NT4 only)
- Storage Manager 8.2 Client
- Redundant Disk Array Controller (RDAC)
- Storage Manager 8.2 Agent
- Storage Manager 8.2 Utilities

Microsoft Virtual Machine

For systems running Windows NT 4.0, the Microsoft Virtual Machine is required to support the Storage Manager 8.2 Agent and the Event Monitor option.

Storage Manager 8.2 Client

The Storage Manager 8.2 Client (SMclient) component provides the graphical user interface (GUI) for managing storage subsystems. The SMclient contains:

- **Enterprise Management**: You can use this component to add, remove, and monitor storage subsystems within the management domain.
- **Subsystem Management**: You can use the Subsystem Management component to manage the components of an individual storage subsystem.

The Event Monitor is a separate program that is bundled with the SMclient. If installed, it monitors storage subsystems whenever the Enterprise Management window is closed. It runs continuously in the background and can send alert notifications in the event of a critical problem.

Redundant Disk Array Controller

The Redundant Disk Array Controller (RDAC) contains the multi-path device driver that is necessary to provide controller failover support when a component on the Fibre Channel I/O path fails.

Storage Manager 8.2 Agent

The Storage Manager Agent (SMagent) provides a management conduit for the Storage Manager Client to configure and monitor the subsystem through the Fibre Channel I/O path.

Storage Manager 8.2 Utility

The Storage Manager 8.2 Utility (SMutil) dynamically defines and maps new logical drives to the operating system. This software package contains the following components:

- **hot_add utility:** The hot_add utility enables you to register newly created logical drives with the operating system.
- **SMdevices utility:** The SMdevices utility enables you to associate storage subsystem logical drives with operating-system device names.
- **SMflashcopyassist utility:** The SMflashcopyassist utility enables you to flush cached data before creating a FlashCopy logical drive and resolve duplicate logical drive signatures between the base logical drives and the FlashCopy logical drives (Windows NT only).

IBM Storage Manager V8.2 adds support for such new features as FlashCopy, Dynamic logical volume expansion (DVE), Remote Volume Mirroring (RVM), and a large increase in number of hosts and logical volumes. FlashCopy and RVM are separately chargeable options.

IBM Storage Manager V8.2 is available as a chargeable upgrade for FAStT200 and FAStT500. Note that on FAStT200, RVM is not available.

For an update on supported operating systems, servers, firmware and drivers, see the IBM FAStT Storage Servers - List of Supported Servers Web site at:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

1.2.3 New features

IBM FAStT Storage Manager Version 8.2 software supports two premium features that can be enabled by purchasing a premium feature key, as well as several new standard features. These features include:

- **FlashCopy**: A premium feature that supports the creation and management of FlashCopy logical drives. A FlashCopy logical drive is a logical point-in-time image of another logical drive, called a *base logical drive*, in the storage subsystem.
- **Remote Volume Mirroring**: A premium feature that is used for online, real-time replication of data between storage subsystems over a remote distance. When the RVM is enabled, the maximum number of logical drives per storage subsystem is reduced.
- **Dynamic logical-drive expansion**: Enables you to increase the capacity of an existing logical drive.
- **2048 logical drive support**: Enables you to increase the number of defined logical drives up to 2048 for each storage subsystem.
- **Storage partitioning**: Supports up to 64 storage partitions.
- **Script Engine and command line interface (CLI) language**: Expands previous version support.
- **Read Link Status Diagnostic (RLS)**: Enhances controller diagnostics so you can isolate the source of link problems on a Fibre Channel loop.
- **Redesigned user interface**: Supports a more integrated process of creating logical drives and then mapping them to hosts or host groups and logical unit numbers to create storage partitions.
- **Controller default IP address**: Controllers A and B are automatically assigned the following IP addresses if no BOOTP/DHCP server is found:
  - Controller A (192.168.128.101)
  - Controller B (192.168.128.102)

  The subnet mask for the default IP address is 255.255.255.0.

You can learn more about these new features in 3.1.2, “What’s new” on page 34.

1.2.4 Storage subsystem management methods

The storage-management software provides two methods for managing storage subsystems:

- Host-agent management method
- Direct management method

Depending on your specific storage-subsystem configurations and host systems, you can use either or both methods. The management methods you select determine where you need to install the software components. You can learn more about these methods in 3.1.3, “Storage subsystem management methods” on page 36.

1.2.5 Arrays and logical drives

FAStT Storage Server is basically a RAID controller. It supports RAID levels 0, 1, 3, and 5 (RAID level 1 is actually implemented as RAID 10). You should always select the RAID level that is most appropriate for your particular data access pattern and that suits your availability requirements. Various RAID levels differ in performance, disk space efficiency, and level of fault-tolerance. The RAID levels are discussed in detail in 3.3.1, “Data protection” on page 51.

To configure the storage server, use the FAStT Storage Manager to group your physical disks into arrays and then create one or more logical drives inside those arrays. Logical drives are the entities that appear to the operating system as physical disk drives. The RAID level is specified per each array. This means all logical drives inside an array use the same RAID level. You can configure up to 2048 logical drives per FAStT700 Storage Server.

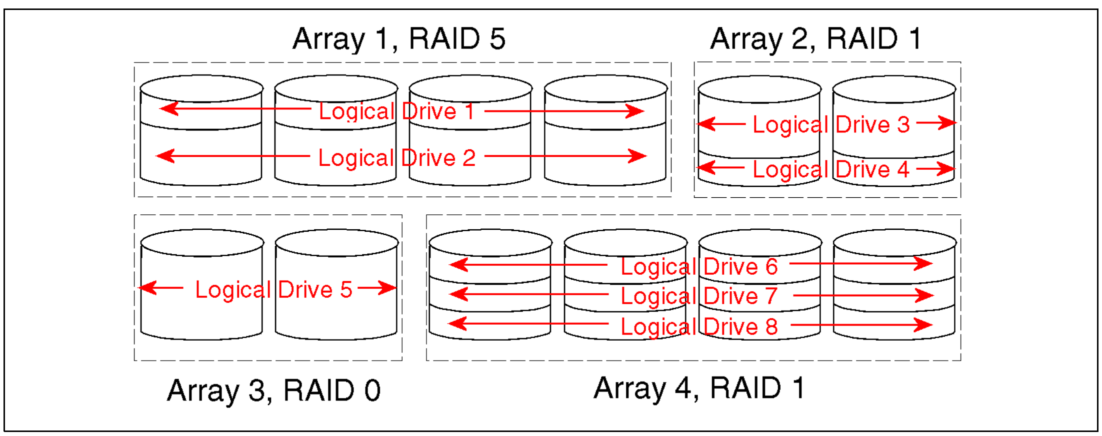

Figure 1-5 Arrays and logical drives

Figure 1-5 shows a sample configuration that consists of 12 physical disk drives. We divided the physical disk drives into four arrays and created a total of eight logical drives in these arrays. The operating system sees these logical drives as physical disks.

Several array and logical drive operations can be done dynamically using the FAStT Storage Manager:

- Adding new disk drives to arrays
- Changing the RAID level for an array
- Changing the segment size of a logical drive
- Defragmentation
- Dynamic Volume Expansion

You can find detailed information on working with arrays and logical drives in 4.3, “Creating arrays and logical drives” on page 100.

1.2.6 Storage partitioning and heterogeneous hosts

Storage partitioning allows you to connect multiple host systems to the same storage server. It is a way of assigning logical drives to specific host systems or groups of hosts. This is known as *LUN masking*.

Logical drives in a storage partition are only visible and accessible by their assigned group or individual hosts. Heterogeneous host support means that the host systems can run different operating systems. However, be aware that all the host systems within a particular storage partition must run the same operating system because they have unlimited access to all logical drives in this partition. Therefore, the file systems on these logical drives must be compatible with host systems.

The best way to ensure this is to run the same operating system on all hosts within the same partition. Some operating systems (Linux for example) may be able to mount foreign file systems. In addition, keep in mind that FAStT Storage Manager does not provide any control of concurrent access to the same disk blocks. Access control must be provided by cluster or file sharing software.

Storage partitioning gives you a high level of flexibility when attaching multiple hosts to the same storage server. For example, if you attach multiple non-clustered hosts, your storage space can be divided into storage partitions and each host can only access its own logical drives. Another example is to use the storage space for a mixture of several clusters and individual hosts. For this to work, each cluster and each host must have access to its own logical drives and must not gain access to other logical drives. With storage partitioning, you create several storage partitions, typically one per each cluster or host.

You can assign up to 32 logical drives into each storage partition, and each logical drive can belong to exactly one storage partition.

Figure 1-6 shows an example of storage partitioning. Five logical drives are divided into two storage partitions. Logical drives 1 and 5 are assigned to Host A. Logical drives 2, 3, and 4 are assigned to Cluster B.

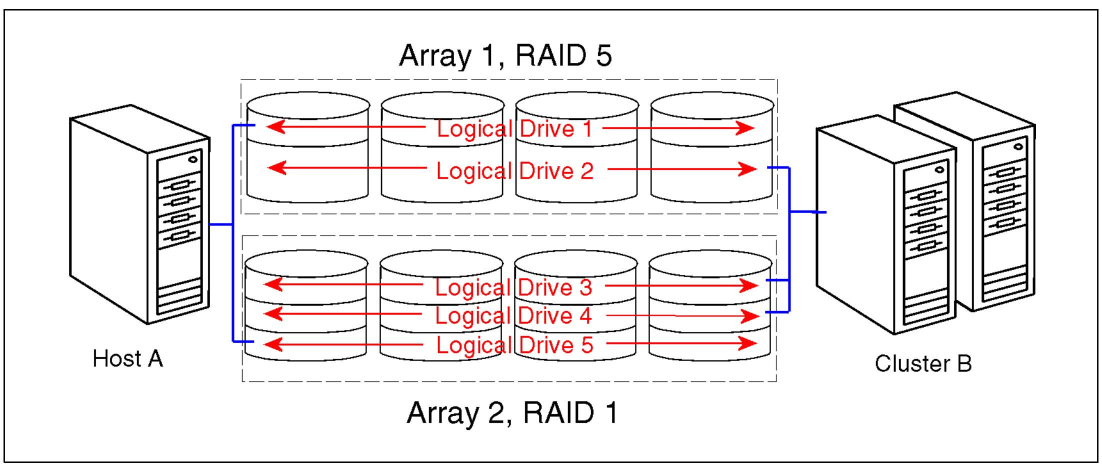

Figure 1-6 Storage partitioning

You can find detailed information on storage partitioning in 4.4, “Configuring storage partitioning” on page 105.

1.2.7 Building blocks and sample configurations

The fibre components that are required to build a storage server-based solution include:

- FAStT700 Storage Server
- FAStT FC-2 Host Adapters
- EXP500 or EXP700 Fibre Storage Expansion Enclosures
- Fibre Channel managed hubs and switches

The best way to explain the role of each component is through sample configurations. Consider this example in Figure 1-7. It shows a minimum FAStT700 configuration that consists of a FAStT700 Storage Server, up to eight EXP700 disk expansion drawers, and a host system with two FAStT FC-2 host bus adapters. The storage server is directly attached to the host system. It also shows the maximum supported number of direct attached hosts, which is four.

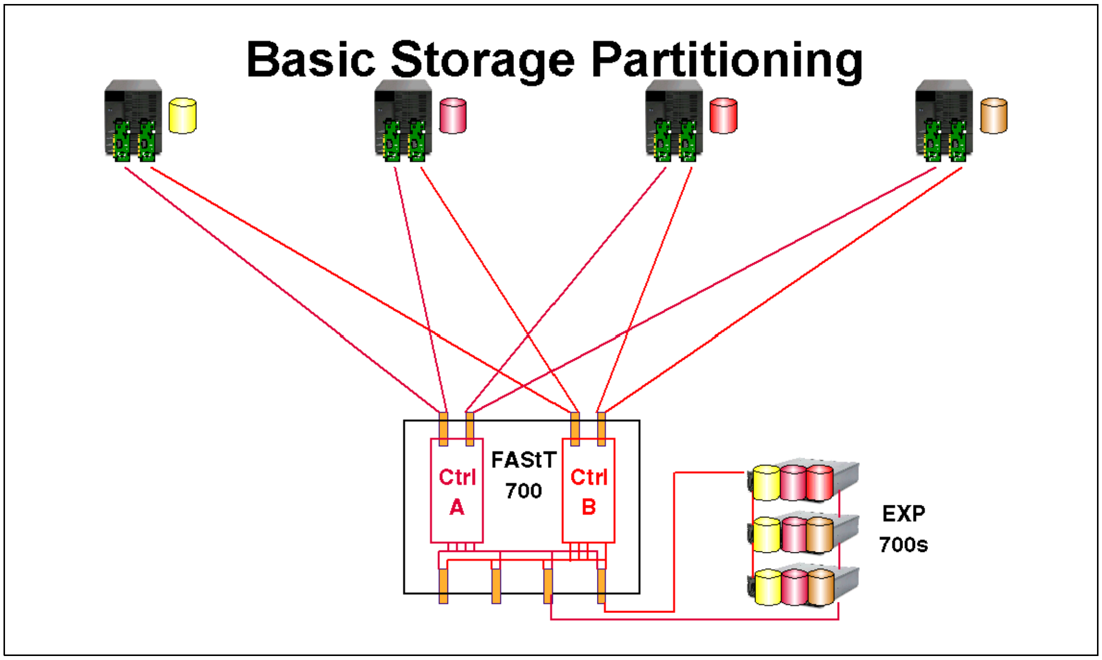

Figure 1-7 Basic FAStT700 configuration

A more complicated example is shown in Figure 1-8. It shows multiple hosts and a cluster accessing FAStT700. We must now move to using a switched fabric. For redundancy, the example uses two switches. For performance, the disks are split across two drive loops and we have dual paths to each switch. Currently only the Windows RDAC supports the dual paths to the switch. The dual paths must be to separate mini-hubs.

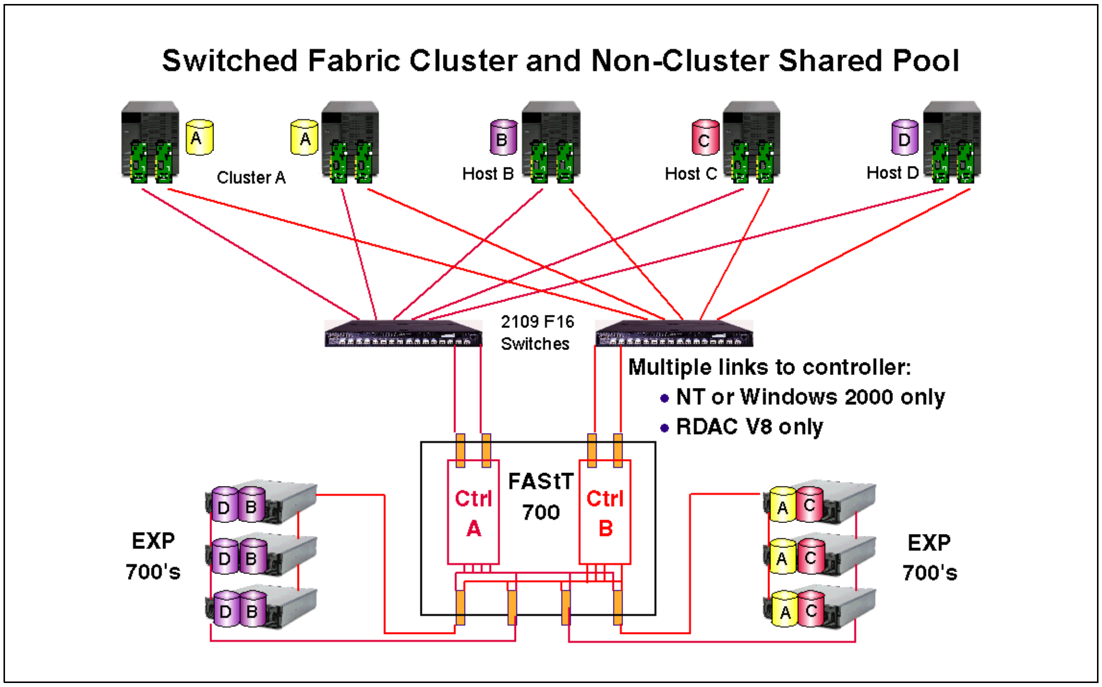

Figure 1-8 FAStT700 switched fabric

1.2.8 Advanced copy services

Two new advanced copy services are available in the IBM FAStT Storage Manager 8.2 software:

- FlashCopy
- Remote Volume Mirroring

Both services are premium features and are enabled by purchasing a premium feature key.

FlashCopy supports the creation and management of *FlashCopy logical drives*. You can learn more about FlashCopy in 3.1.2, “What’s new” on page 34.

Remote Volume Mirroring (RVM) is used for online, real-time replication of data between storage subsystems over a remote distance. You can learn more about Remove Volume Mirroring in 3.1.2, “What’s new” on page 34.

2

Hardware details

This chapter describes the hardware components of a Fibre Channel storage solution. First it covers the IBM TotalStorage FAStT700 Storage Servers. Then it discusses additional hardware components that are essential for a complete storage solution. These include the EXP500 and EXP700 external expansion enclosure, the TotalStorage FAStT-2 Host Adapter, and the FC switches 2109-F08 and 2109-F16.

You can find a description of the other members of the FAStT product family, such as FAStT200, FAStT500, and EXP500, in the IBM Redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246.

2.1 IBM TotalStorage FAStT700 Storage Server

The FAStT700 Storage Server (Figure 2-1) provides a higher level of performance, availability, and expendability than FAStT200 or FAStT500. It is a 4U rack-mountable device that is suitable for high-end storage needs.

Figure 2-1 IBM TotalStorage FAStT700

High performance

FAStT700 uses an Intel processor in each RAID controller unit. It supports 1 GB of cache memory in each controller. These factors contribute to a very high throughput. This is especially true when you compare the performance of FAStT700 to the other members of the FAStT product family, such as FAStT200 and FAStT500. In addition, FAStT uses the new 2 Gb/sec. Fibre Channel standard on the host side as well as on the drive side. This guarantees a high throughput even if a large number of hosts and disk drives are attached to FAStT700.

High availability

The following high availability features are standard on the FAStT700 Storage Server:

- It contains two hot-swappable RAID controller units.
- Each RAID controller unit can be connected to two host loops through two mini-hubs.
- Each RAID controller unit contains 1 GB of battery backed cache for a total of 2 GB of cache.
- The drive-side FC loops can be connected in a fully redundant way.
- It contains hot-swappable and redundant power supplies and fan units.
- The cache backup battery is hot-swappable.
- With proper FC cabling, you can avoid any single point of failure.

High scalability

You can connect up to eight EXP700 expansion enclosures or 11 EXP500 expansion enclosures to each redundant FC loop on the drive side. That is 112 FC disk drives use the EXP700 or 110 FC disk drives use the EXP500. Two such fully redundant loops can be used. Therefore, you can connect up to 16 EXP700 expansion enclosures or 22 EXP500 expansion enclosures. This means you can attach up to 224 disk drives by using the EXP700 or 220 disk drives using the EXP500 to FAStT700 without a single point of failure.

2.1.1 Front view

After you remove the front bezel, you can access the following components as shown in Figure 2-2:

- **Two RAID controllers**: Each RAID controller or blade (this terminology is used within the online help) contains diagnostic LEDs and a reset switch.
- **Controller fan module** (which also contains status LEDs): Two fans are integrated in this module. Their role is to provide cooling to both RAID controller blades. FAStT700 diagnostic LEDs are also present on the module as well as an alarm switch.
- **Battery module**: It is possible to hot-replace the battery module. During this operation, the cache temporarily lacks battery protection. If write-back cache policy is used and FAStT500 experiences a fault, this may lead to data loss. However, the process of hot-swapping the battery does not take much time and such an event is highly unlikely.

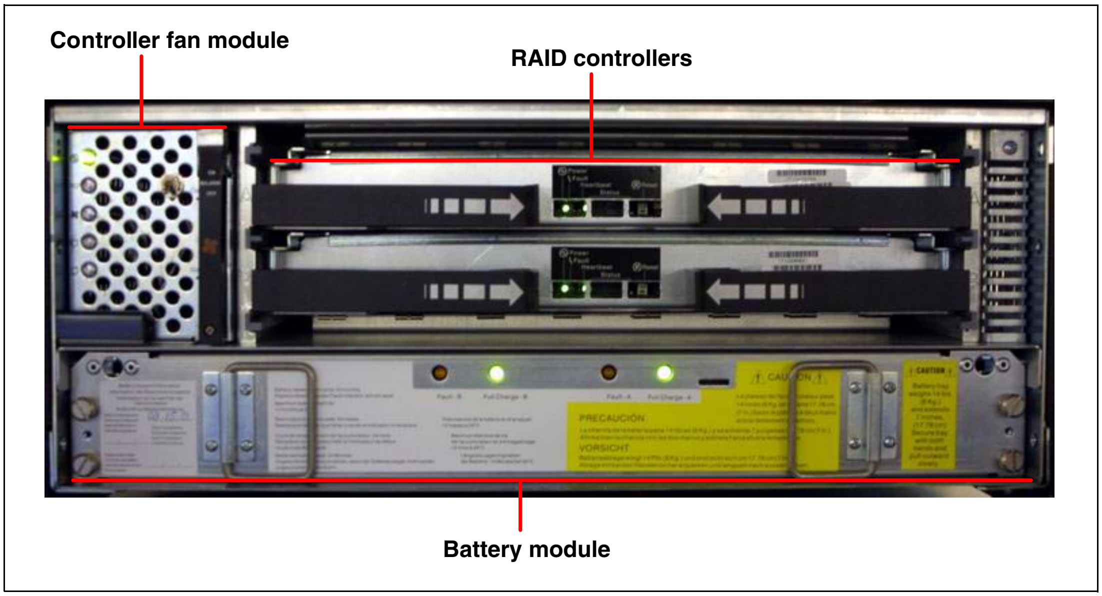

Figure 2-2 Front view of FAStT700 with a detached bezel

The upper controller is called controller A, and the lower one is called controller B.

2.1.2 Rear view

The rear view of FAStT700 is shown in Figure 2-3. The components that are accessible on this side of FAStT700 include:

- Mini-hubs with SFP module ports for host connectivity
- Mini-hubs with SFP module ports for FC expansion enclosure connectivity
- Fan and communications module with the management ports
- Two hot-pluggable and redundant power supplies

A SFP module, called *small-form pluggable module*, is used to convert electrical signals to optical signals that are required for Fibre Channel transmission to and from the controllers in the FAStT Storage Server. It is the successor of the GBIC modules used in FAStT200 and FAStT500.

The fan and communications module features a serial and Ethernet port per RAID controller unit for management and troubleshooting purposes. Besides management ports, this module has two fans that provide redundant cooling to both power supplies.

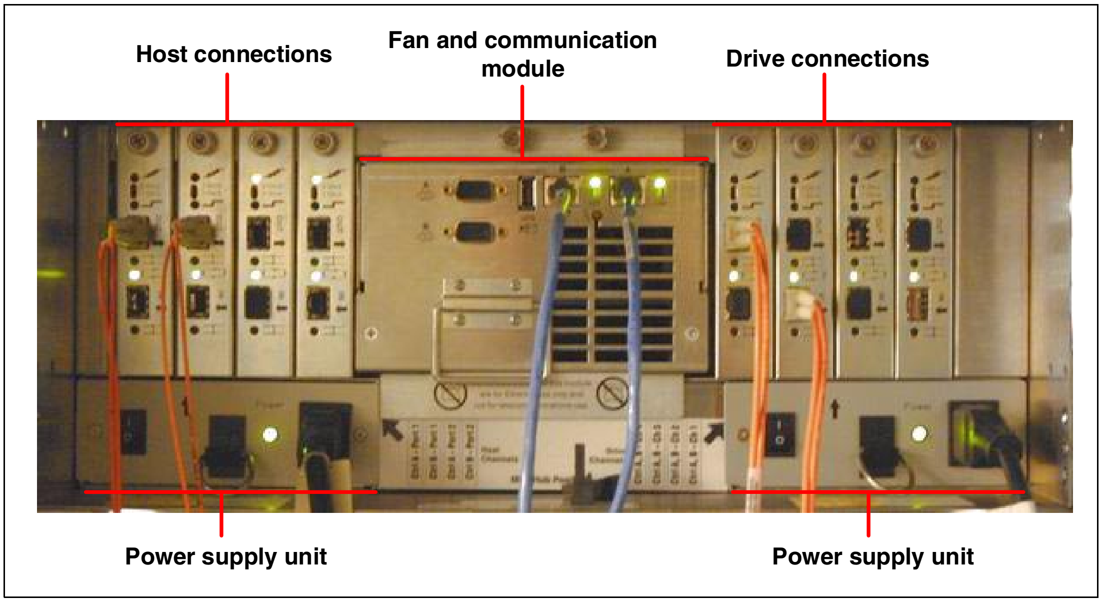

Figure 2-3 Rear view of FAStT700

2.1.3 FAStT700 mini-hubs

FAStT700 can take up to eight mini-hubs in the back – four for host connections and the other four for drive connections.

The mini-hubs of FAStT700 provide connectivity to the host systems, either direct or via a fibre switch or hub, and connect the drive enclosures. Each mini-hub provides two ports for SFP modules, which are available in short- or long-wave versions. The mini-hubs of the drive side are the same as those used on the host side. Figure 2-4 shows a mini-hub with a cable attached to a SFP module.

You shouldn’t use a SFP module in a mini-hub, if that SFP module isn’t used for any connection. If you install a second SFP module in a mini-hub, the mini-hub forces both connections provided through the two SFP modules to use the loop protocol. If only one SFP module is used, the connection uses the fabric protocol. But if two SFP modules are installed but only one is connected, then this configuration can cause problems.

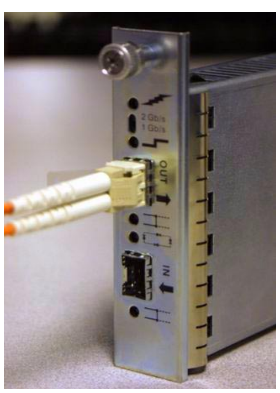

Figure 2-4 FAStT mini-hub

The mini-hubs on the host side have dedicated connections to the controllers. The mini-hubs in position 1 and 3 are connected to controller A, and the mini-hubs 2 and 4 are connected to controller B. The four mini-hubs on the drive side are connected to each of the two controllers to provide redundant attachment of the drives. The internal cabling is also shown in the diagram in Figure 2-5.

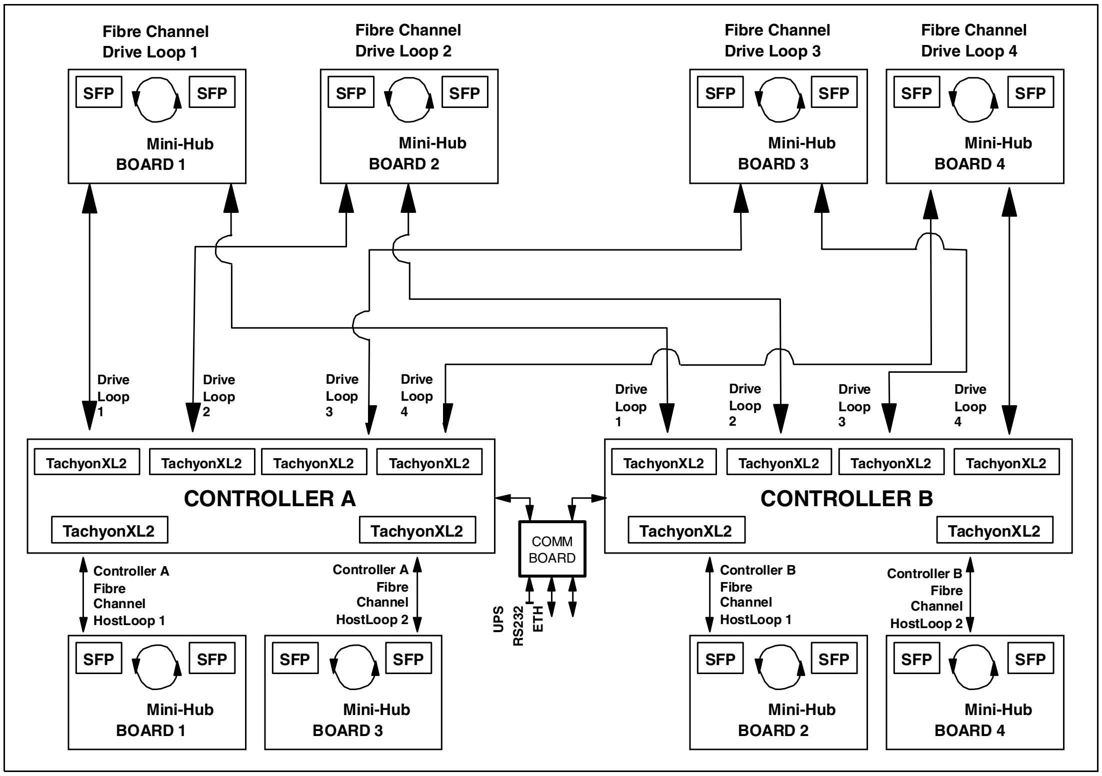

Figure 2-5 Connectivity between FAStT and the mini-hubs

The mini-hubs support a 1Gb/sec. and 2Gb/sec. Fibre Channel. The link speed is selected with a DIP switch as shown in Figure 2-9 on page 21.

2.1.4 Host-side connections

On the host side, each RAID controller is connected to two FC loops using two mini-hubs. You can have a total of four mini-hubs, and each mini-hub can take two SFP modules. This means you can establish up to eight host connections without needing an external hub or switch.

To assure redundancy, the usual way is to install a pair of host adapters inside each server and connect each of the two adapters to a different RAID controller. For example, if one host adapter is connected to mini-hub 1 or 3 (these are connected to RAID controller A), the other host adapter must be connected to mini-hub 2 or 4 (RAID controller B). This way you can connect up to four servers directly to FAStT700. For more servers, you have to resort to using external FC hubs or switches.

2.1.5 Drive-side connections

On the drive side, each RAID controller is connected to four FC loops. These four FC loops are used as two pairs of redundant loops. We use mini-hubs 1 and 2 to form one redundant loop and mini-hubs 3 and 4 to form another. You can connect each redundant loop to up to eight EXP700 or 11 EXP500 expansion enclosures. This gives you a total of up to 16 or 22 enclosures (224 or 220 disk drives) on both loop pairs. For more detailed information on drive side connectivity, see 2.2.5, “FAStT700 and EXP700 cabling” on page 26, for the EXP700 expansion enclosure. You should also refer to the redbook *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246, for information regarding the EXP500.

2.1.6 FAStT700 diagnostic LEDs

FAStT700 provides several diagnostic LEDs that are located in the module they are monitoring.

Front view

In the front, several LEDs show the overall status of FAStT700; they are located in the controller fan module. The battery module itself has a couple of LEDs that monitor the two separate batteries backing the cache of the controllers (see Figure 2-6).

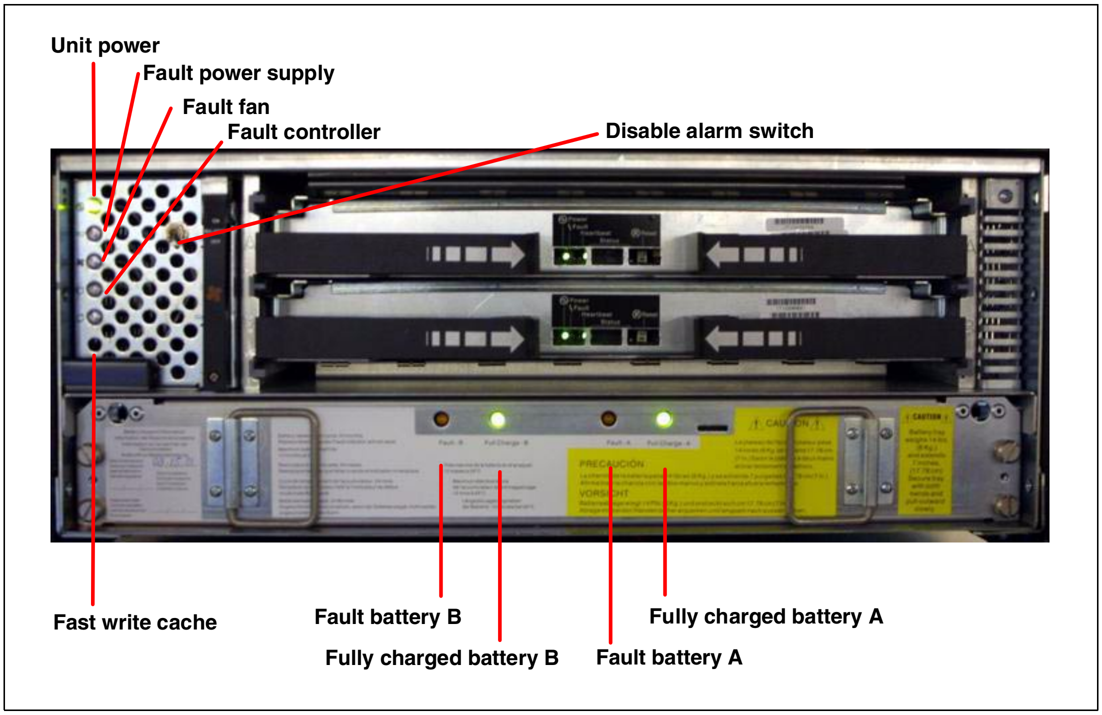

Figure 2-6 FAStT700 front diagnostic LEDs

The alarm switch disables the acoustic alarm if one of the components of FAStT700 fails.

Controller fan module LEDs

These LEDs are also visible when the front bezel is not detached:

- **Unit power LED** (green): This LED indicates that FAStT is powered on and is normally ON.
- **Fault power supply LED** (amber): This LED indicates the failure of one of the two power supplies and is normally OFF.
- **Fault fan LED** (amber): This LED indicates the failure of one of the fans and is normally OFF.
- **Fault controller** (amber): This LED indicates a faulty RAID controller is normally OFF.
- **Fast write Cache LED** (green): This LED indicates that the cache contains data that has not been written to disk.

Battery module LEDs

The module houses four LEDs for the status of the two batteries.

- **Faulty battery LED** (amber): The LED indicates that the battery has failed and needs to be replaced. The LED is normally OFF.
- **Fully charged battery LED** (green): The LED has two modes. If the LED is constantly ON, the battery is completely charged. If the LED is flashing, the battery is being charged at the moment. The FAStT Storage Manager reports an error as long the battery is being charged.

RAID controller

Each controller has its own control panel as shown in Figure 2-7.

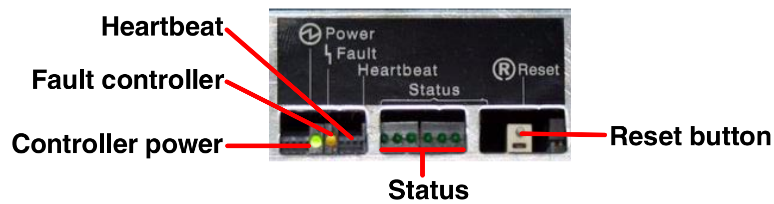

Figure 2-7 FAStT controller LEDs

Each controller can be reset separately by pressing the white reset button.

Controller LEDs

The controller LEDs include:

- **Controller power LED** (green): This LED indicates that the controller is powered up. It is normally ON.
- **Fault controller LED** (amber): This LED indicates a failure of the controller. It is normally OFF.
- **Heartbeat LED** (green): This LED indicates the controller is working properly. It normally flashes every second.
- **Status LED** (green): This array of eight indicators (including the heartbeat LED) may be used to determine the nature of a fault condition indicated by the controller fault LED.

Rear view

The back of FAStT700 also provides some diagnostic information via LEDs as shown in Figure 2-8.

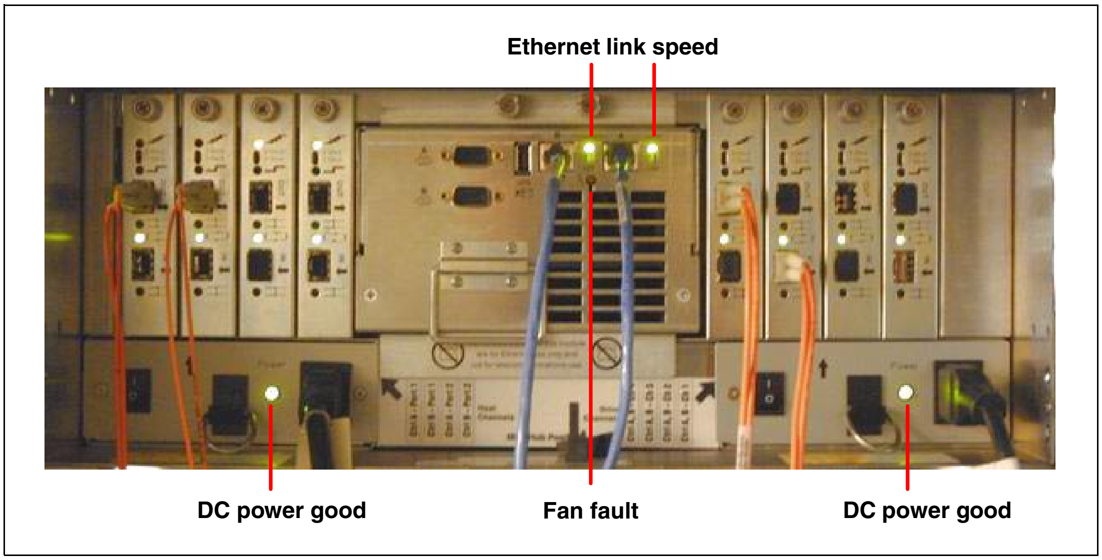

Figure 2-8 FAStT rear diagnostic LEDs

Power supply LEDs

Each power supply contains one LED – DC power good LED (green). This LED indicates the DC power is OK and is normally ON.

Fan and communications module LEDs

The fan and communications module contains five LEDs:

- **Ethernet link speed LED** (green): There are two LEDs per Ethernet port (for a total of four LEDs), which are dedicated to controller A and B. The upper LEDs indicate a proper link with 100 Mb/sec., and the lower ones indicate a link with 10 Mb/sec. If FAStT700 is attached to the Ethernet, one of the LEDs per Ethernet connection is normally ON. In Figure 2-8, two of the LEDs are lit and the other two are not.
- **Fan fault LED** (amber): This LED indicates a failure of the fan in the module. It is normally OFF.

Mini-hub

The mini-hub provides several LEDs to quickly detect hardware failures and link problems. The DIP switch allows the link speed selection for the SFP module between the 2Gb/sec. and 1Gb/sec. Fibre Channel (see Figure 2-9).

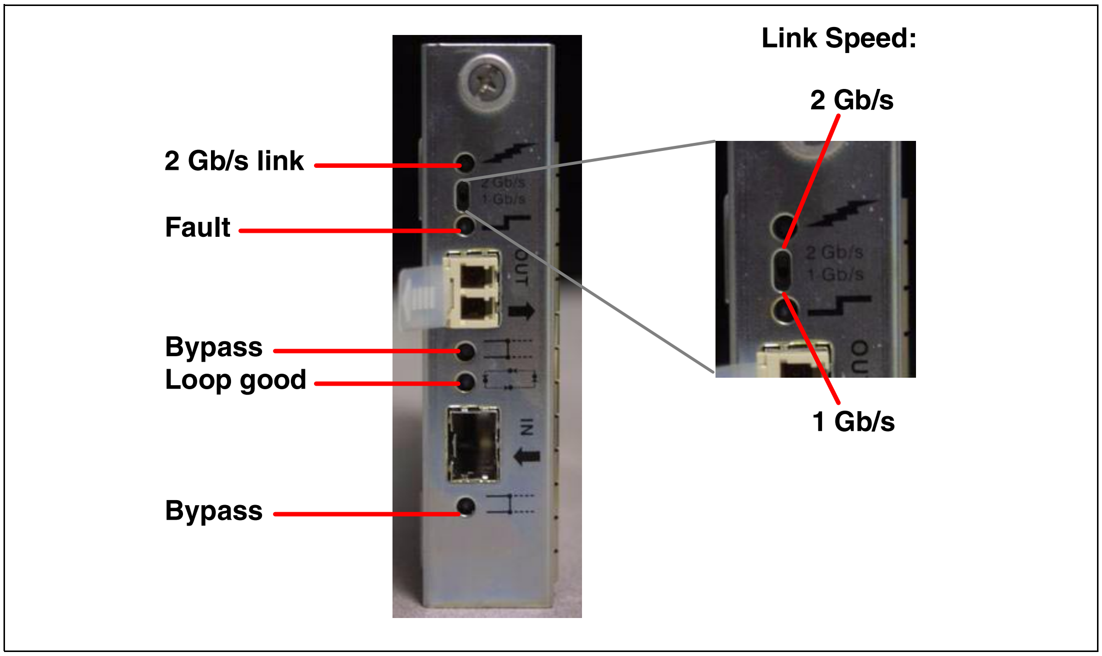

Figure 2-9 FAStT mini-hub explained

Important: The Fibre Channel link speed must be set manually via the DIP switch. The mini-hub does not auto detect this setting.

Mini-hub LEDs

The LEDs present on the mini-hub are:

- **2 Gb/s link LED** (green): This LED is ON when the selected link speed is 2Gb/s and a link is up.
- **Fault LED** (amber): This LED should normally be OFF. If it is ON, it indicates a fault of the mini-hub or one of the SFP modules.
- **Bypass LED** (amber): This LED should normally be OFF if no SFP module is installed. But if a SFP module is present and a link error is detected, such as no cable or faulty cable, it comes on. There is one Bypass LED for each SFP module.
- **Loop Good LED** (green): This LED should be ON normally. It may be OFF if you have link errors.

2.2 FAStT EXP700 expansion enclosure

The EXP700 is a rack-mountable storage expansion enclosure that contains 14 bays for slim-line hot-swappable FC disk drives. It occupies 3U inside a rack and features hot-pluggable and redundant power supplies and fans. In addition, it contains two Enclosure Service Monitor (ESM) boards and can be connected to FAStT700 in a fully redundant FC loop.

2.2.1 Front view

From the front, you can access up to 14 hot-swappable FC hard disk drives and obtain status information via diagnostic LEDs (see Figure 2-10).

Figure 2-10 EXP700 Storage Enclosure

2.2.2 Rear view

You access the ESM boards, power supplies, and fans from the rear side of the EXP700 (Figure 2-11). Two hot-swappable power supplies and fan units provide redundancy and, therefore, offer a higher availability level. If a fan unit fails, you should not remove it from the storage server until the replacement is available. The cooling airflow is not optimal when a fan unit is missing.

The same is true for power supplies. Do not remove a failed power supply unit before you have a replacement available. If you do, the cooling is no longer efficient.

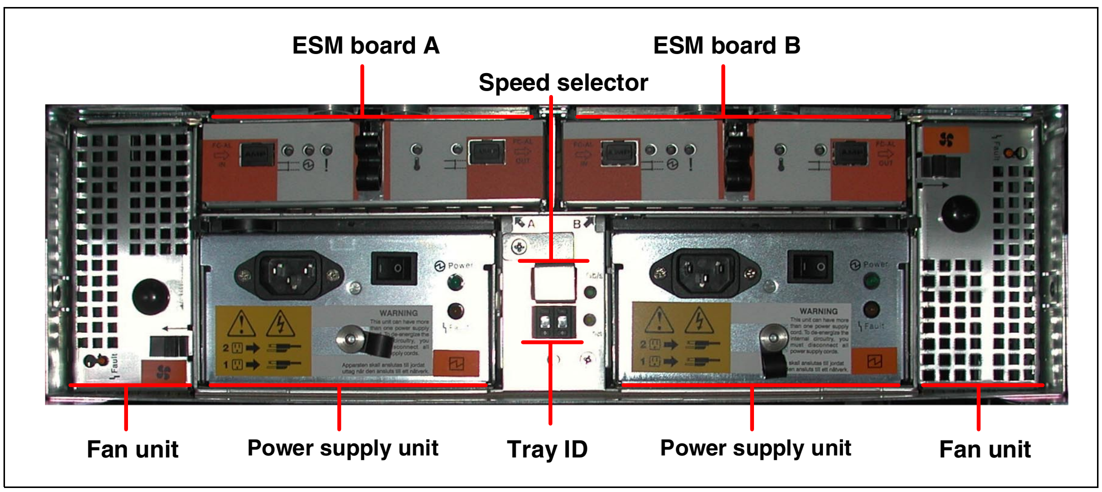

Figure 2-11 EXP700 Rear View

ESM redundancy is provided through the redundant dual loop configuration that FAStT700 provides.

2.2.3 EXP700 ESM boards

The ESM boards provide the connectivity of the drives inside the EXP700 to FAStT700. As you can see in Figure 2-12, both ESM boards are connected to all 14 disk drive bays in the EXP700. This provides redundancy in case of a failed ESM board or faulty loop connection. All disk drives remain connected through the other ESM board or FC loop. Each ESM board has an incoming and outgoing SFP module port.

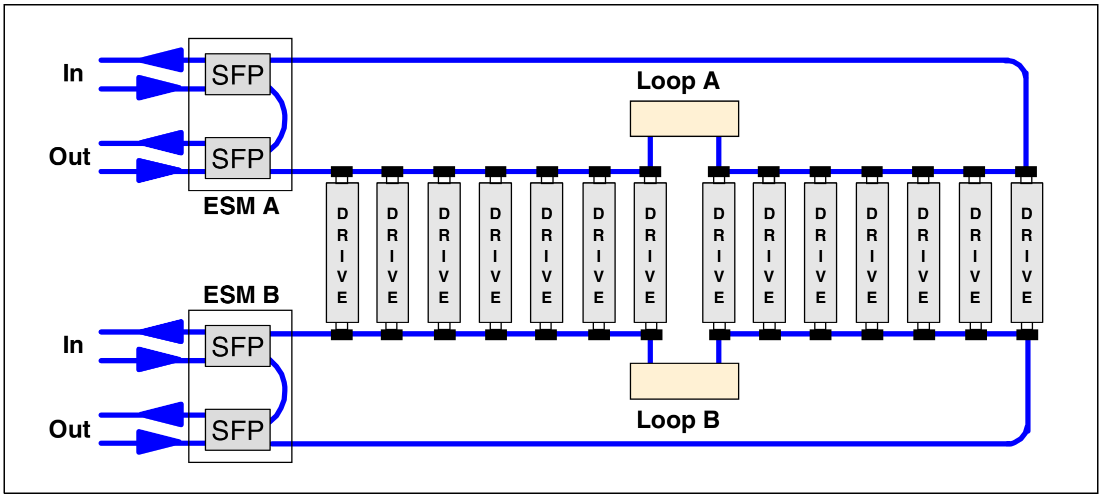

Figure 2-12 EXP700 internal cabling

One SFP module port is marked as IN, and the other one is marked as OUT. A close look at Figure 2-12 reveals that both ports actually function identically, so it is not mandatory to stick to their incoming or outgoing role. However, we recommend that you use them in the suggested manner. If you always connect outgoing ports on FAStT700 to incoming ports on EXP700, you will introduce clarity and consistency to your cabling implementation. This allows for easier, quicker, and more efficient troubleshooting.

It is very important to correctly set the tray ID switches on ESM boards. They are used to differentiate between multiple EXP700 enclosures that are connected to the same FAStT700. Each EXP700 must be use a unique value. The FAStT Storage Manager utility uses the tray IDs to identify each EXP700 enclosure. In addition, the FC loop ID for each disk drive is automatically set according to:

- The EXP700 bay where the disk drive is inserted
- Tray ID setting

Two switches are available to set the tray ID (Figure 2-13):

- A switch for tens (x10)
- A switch for ones (x1)

Therefore, you can set any number between 0 and 99.

Important: Each EXP attached to the same FAStT must have a different tray ID:

Another important setting you have to verify is the link speed selection. The link speed selector is located just above the tray ID. It is normally secured by a small metal cover to prevent from accidental usage. If you are attaching the EXP700 to a FAStT700, you can use the 2Gb/sec. switch setting, which is the default setting. However, if you connect the EXP700 to a FAStT500 or you are going to intermix EXP500 and EXP700 on the same FAStT700, you have to choose the 1Gb/sec. link speed (see Figure 2-13).

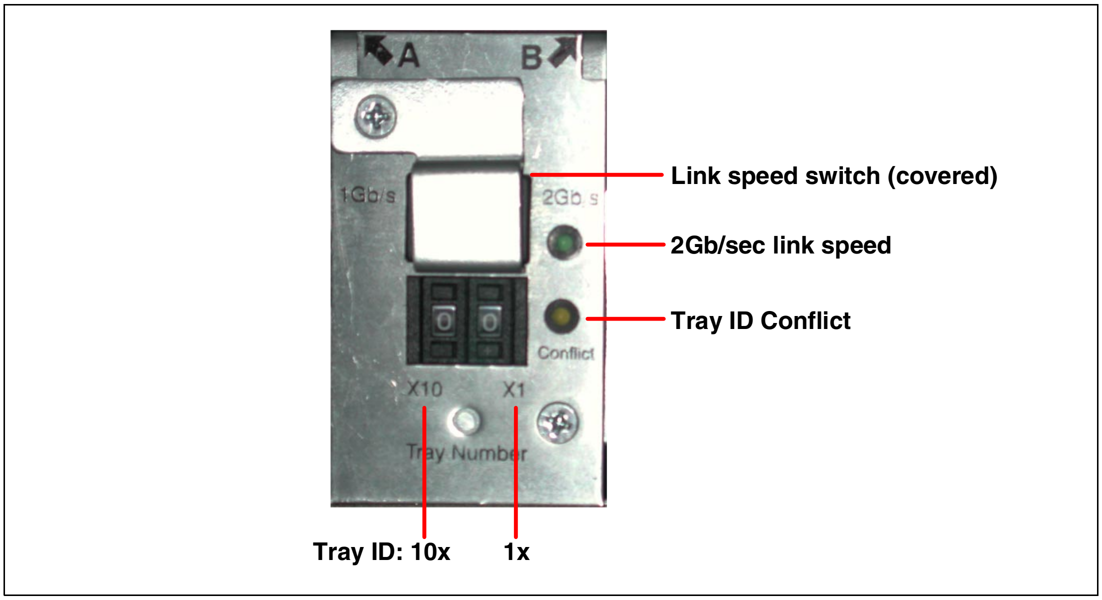

Figure 2-13 EXP700 tray ID and link speed setting

2.2.4 EXP700 diagnostic LEDs

The EXP700 provides several diagnostic LEDs in the different components.

Front view

Two diagnostic LEDs are located in the top left corner of the EXP700. Each disk drive contains another two LEDs, which are actually not situated in the disk tray itself but on the backplane of the EXP700. They are only shown through the drive tray. In this way, it is assured that the drive tray is completely passive. You can see the location of the LEDs in Figure 2-14.

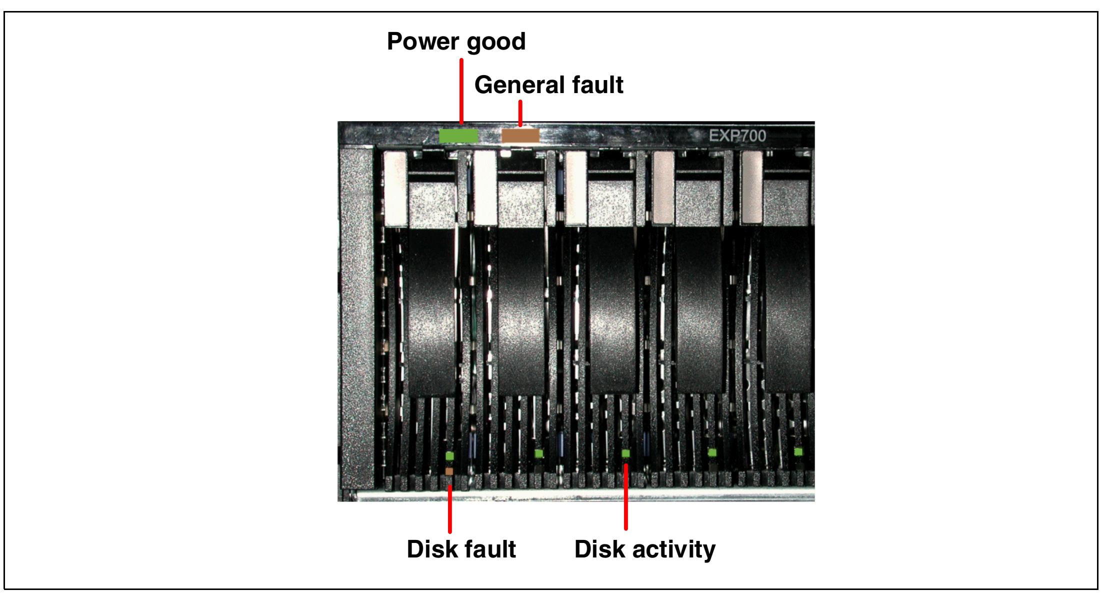

Figure 2-14 EXP700 front diagnostics LEDs

EXP700 diagnostic LEDs

The EXP700 diagnostic LEDs include:

- **Power LED** (green): This indicates the EXP700 is powered up. It is normally ON.
- **General system fault LED** (amber): This indicates whether any component in the EXP700 fails. This LED light is normally OFF.

FC disk drive LEDs

The FC disk drive LEDs include:

- **Activity LED** (green): When this LED flashes, it indicates disk drive activity. When it is steadily ON, it means the disk drive is properly functioning. The LED is normally ON.
- **Disk drive fault LED** (amber): This LED indicates a defunct drive. It is normally OFF. If the LED is flashing, the drive is identified through the FAStT Storage Manager.

Rear view

This section explains the different diagnostic LEDs on the rear of the EXP700 module by module.

Tray ID and Link Speed Module LEDs

There are two diagnostic LEDs (see Figure 2-13):

- **Tray ID Conflict** (amber): This LED indicates that another EXP with the same tray ID is attached to the FAStT Storage Server. The LED is normally OFF.
- **2Gb/sec Link Speed** (green): This LED indicates a link speed of 2Gb/sec. and is lit when such a link is detected.

ESM board LEDs

Each ESM board contains five LEDs as shown in Figure 2-15 and as explained in the following list:

- **Bypass LED** (green): This LED should normally be OFF if no SFP module is installed or an SFP module is installed and a proper link is detected. But if a SFP module is present and a link error is detected, for example, no cable or faulty cable, it comes on. There is one Bypass LED for each SFP module or two as shown in the example.
- **Power Good LED** (green): This indicates the EXP700 is powered up. It is normally ON.
- **Fault LED** (amber): This indicates if a failure in the ESM board occurs. Normally it is OFF.
- **Over Temperature LED** (amber): This LED indicates an over temperature condition. Normally it is OFF.

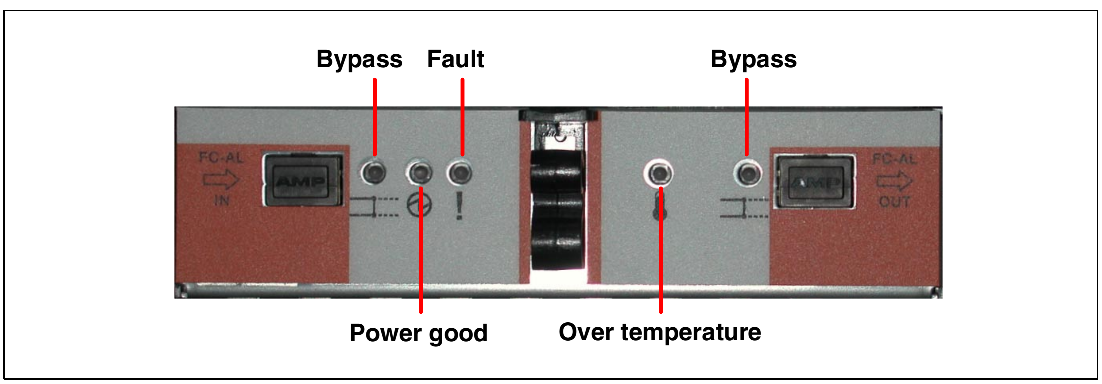

Figure 2-15 EXP700 ESM board diagnostic LEDs

Fan module LED

Each fan module provides one LED – Fan fault LED (amber) – as shown in Figure 2-16. This LED indicates a fan failure of the module. It is normally OFF.

Figure 2-16 EXP700 fan and power supply module diagnostic LED

Power supply LEDs

Each power supply houses one diagnostic LED (see Figure 2-16).

EXP700 power supply diagnostic LEDs

There are two diagnostic LEDs:

- **DC power good LED** (green): This LED indicates whether the DC power is OK. It is normally ON.
- **Fault LED** (amber): This LED indicates a failure of the power supply. It is normally OFF.

2.2.5 FAStT700 and EXP700 cabling

The FAStT700 Storage Server offers flexibility and scalability in attaching disk enclosures. You can connect up to eight EXP700 or 11 EXP500 enclosures in each redundant loop, so that gives you a maximum of 16 or 22 enclosures. Using only EXP700, this gives you a maximum of 224 drives attached to a single FAStT700, using only EXP500 the maximum is 220 drives.

Restrictions apply because the EXP500 is working with 1Gb/s Fibre Channel compared to the possible 2Gb/s Fibre Channel of the EXP700. If you only connect EXP700 enclosures to the FAStT Storage Server, you can use the 2Gb/s link speed. As soon as an EXP500 is connected to this FAStT, in another drive loop or on the same drive loop, all drive side mini-hubs must be set to 1GB/s Fibre Channel, because all four drive channels must work on the same speed.

Restrictions:

- If an EXP500 is attached to FAStT700, all drive side mini-hubs must be set to 1Gb/s link speed.
- In a mixed environment, using EXP500 and EXP700 on the same drive loop, a maximum of 110 drives is supported in a single redundant drive loop.

Figure 2-17 shows four EXP700 enclosures connected to a FAStT700 through two fully redundant FC loops. The two EXP700s in the upper part are connected to mini-hubs 1 and 2 and the lower two EXP700s are connected to mini-hubs 3 and 4.

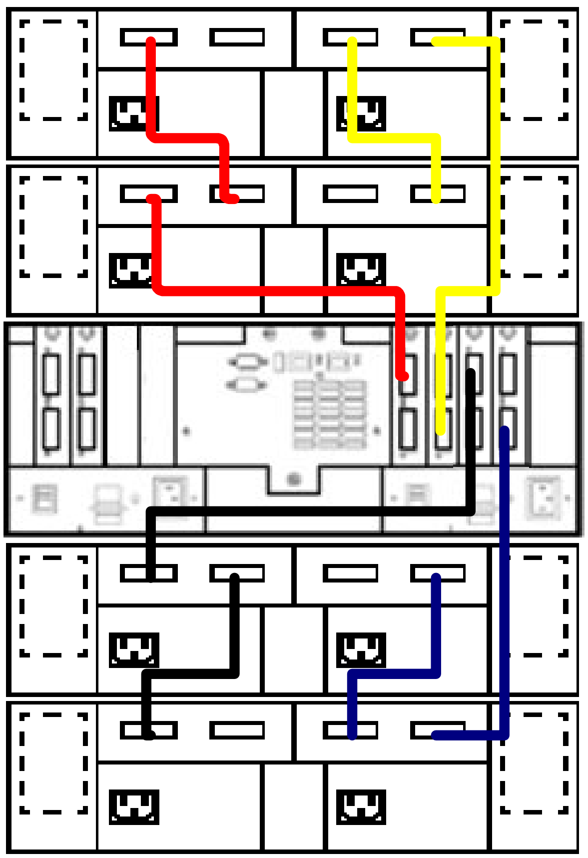

Figure 2-17 Cabling of FAStT700 and four EXP700

Important: Please note that on FAStT700, you only use one SFP module port on each drive side mini-hub – either the incoming or outgoing one (whichever is consistent with your cabling strategy), but not both.

2.3 General cabling and connectivity issues

There are several key points to remember when using cabling:

- Light is transmitted down the cable. Therefore, you should ensure that the cable is not kinked and that you do not use tight bends since these can damage the cable.
- The mini-hubs in FAStT700 and EXP700 use LC connectors. They have a switch that determines at what speed the fibre loop will run. See “Mini-hub” on page 20. The

mini-hubs in FAStT500 and EXP500 use SC connectors. The differences in the two types of connectors are shown in Figure 2-18.

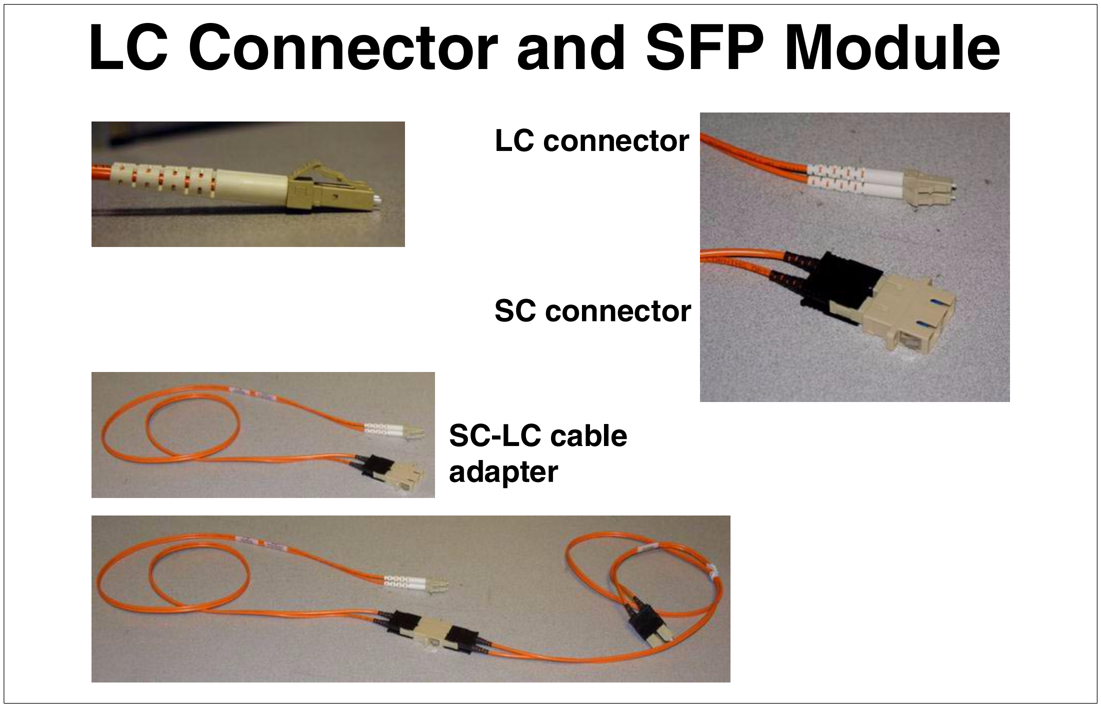

Figure 2-18 LC and SC connectors

- When planning the cabling from FAStT700, use care to ensure that the correct cables are ordered with the correct connectors, SC or LC. Figure 2-19 shows the types of connectors that are used by different components. Therefore, you need converters to connect FAStT700 and EXP500, as well as FAStT700 and older HBAs and switches.

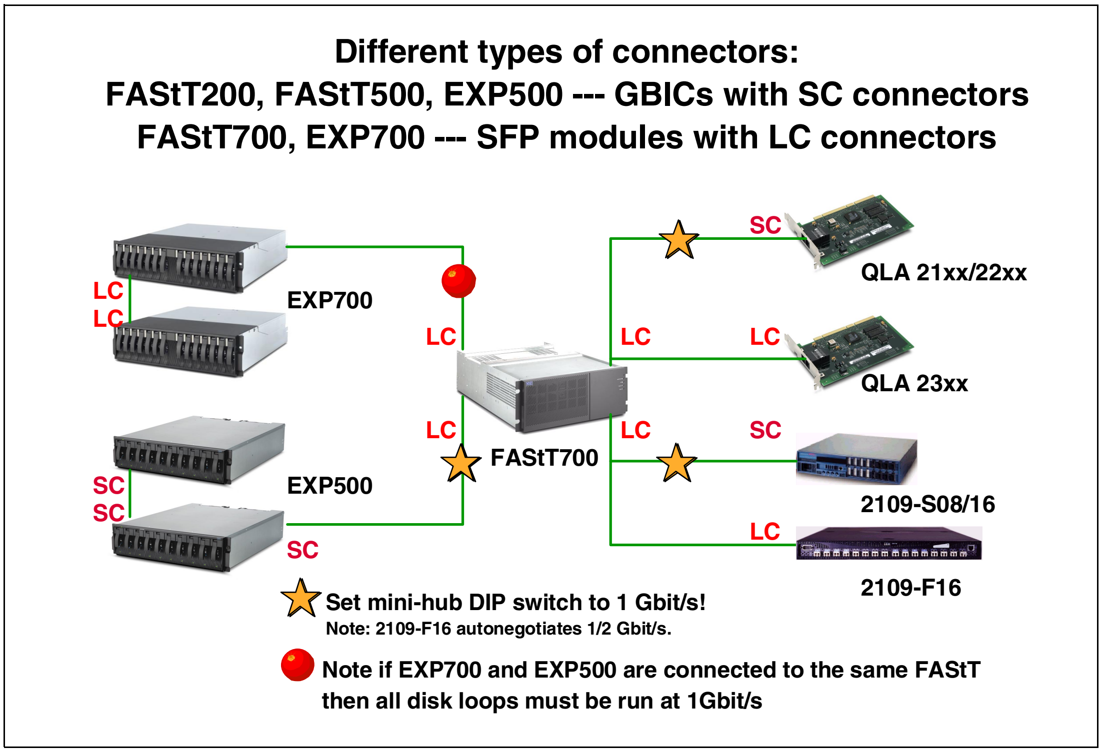

Figure 2-19 Connectors and cable types

2.4 FAStT FC-2 Host Bus Adapter

This high-performance host bus adapter (Figure 2-20) delivers data transfer of up to 2 Gb/s. The adapter is a half-length PCI-X adapter, which means it runs at 66 MHz on a 64-bit PCI-X bus. It is compatible with 64-bit PCI slots running at 33 MHz. It may be used in 32-bit PCI slots running at 66 MHz or 33 MHz.

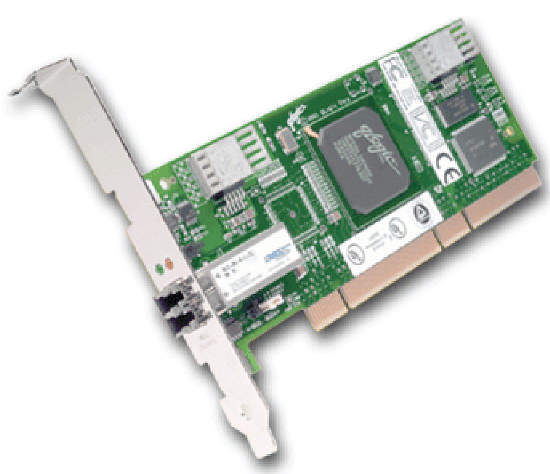

Figure 2-20 FAStT FC-2 Host Adapter

The adapter auto-senses the data transfer rate to 2 GB/s or 1 GB/s, depending on whether you connect to FAStT200, FAStT500, or FAStT700.

For a higher level of availability, the adapter supports Active PCI-X functionality. In a Windows 2000 environment, you can configure two adapters as a fault-tolerant pair using the PCI-X hot plug functionality. When one fails and another takes over, the failed adapter can be replaced without bringing the server down.

Hot-add functionality is supported as well. It can further increase system availability. Using hot-add, you can avoid system downtime when adding new adapters.

The adapter supports short-wave fibre cabling with a SFP module built-in. If you want to use long-wave cables in your FC network, you have to use a FC hub or a switch with long-wave GBIC. The connection between the adapter and the hub or switch must be short-wave.

2.5 TotalStorage SAN Switch F16 and F08

You can use the IBM TotalStorage SAN Switches to interconnect multiple host systems with many storage servers and other SAN devices. They can provide the throughput required for large SAN configurations. Two versions exist – one for eight ports and the other for 16 ports. A TotalStorage SAN Switch F16 is shown in Figure 2-21.

Figure 2-21 IBM TotalStorage SAN Switch F16

Each switch port provides bandwidth of up to 200 MB/s with maximum latency of 2 microseconds. The non-blocking architecture allows multiple simultaneous connections. Switch ports can operate in either of these modes: F, FL, or E. The F16 switch comes with eight short-wave SFP modules, and the F08 switch comes with four short-wave SFP modules. You can install either short or long-wave SFP modules into the additional switch ports.

If you cascade several switches, you can achieve complex SAN connections with a massive number of switch ports. You can also increase the distance by connecting the switches in a series. Up to seven hops are allowed. This means up to 70 km. if you use long-wave SFP modules and cables.

The switches are self-learning. This feature allows the fabric to automatically discover and register host and storage devices. Another important capability is self-healing, which enables the fabric to isolate a problem port and reroute traffic onto alternate paths. For performance and flexibility, an internal processor provides services such as name serving, zoning, and routing. For availability, you can add an optional second hot-plug power supply to provide redundant power, so that you can avoid downtime when a power supply fails. You can also perform dynamic microcode upgrades.

To configure the FC switch, use the IBM TotalStorage SAN Fibre Channel Switch Specialist management tool. The switch features embedded Web browser interface for configuration, management, and troubleshooting.

Because designing and configuring a SAN fabric using zoning techniques is a complex topic, they are not covered in this redbook. However, this book illustrates several solutions that need zoning.

You can find in-depth information in several redbooks including:

- *Designing an IBM Storage Area Network*, SG24-5758
- *Implementing an Open IBM SAN*, SG24-6116
- *Implementing an Open IBM SAN Featuring the IBM 2109, 3534-1RU, 2103-H07,* SG24-6412
- *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127

3

Managing the FAStT Storage Server

This chapter introduces the IBM FAStT Storage Manager software, which is used to configure, manage, and troubleshoot the FAStT Storage Servers. It is a Java-based GUI utility that is available for various operating systems.

You can use the FAStT Storage Manager software to:

- Create new arrays and logical drives
- Expand existing arrays and logical drives
- Migrate to a different RAID level
- Configure storage partitioning
- Perform diagnostic and troubleshooting tasks

Attention: Some of the operating systems that are discussed in this section, at the time this redbook was written, are not formally supported. You must check the FAStT supported servers Web site prior to connecting hosts to the FAStT Storage Server or using the advanced functions. You can locate the FAStT700 supported servers Web site at:

`http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm`

3.1 The FAStT Storage Manager software

The IBM FAStT Storage Manager software is used to configure arrays and logical drives, assign your logical drives into storage partitions, replace and rebuild failed disk drives, expand the size of arrays and logical drives and convert from one RAID level to another. It also allows you to perform troubleshooting and management tasks, like checking the status of the FAStT Storage Server components, updating the firmware of RAID controllers, and similar activities. It also allows you to configure and manage FlashCopy Volumes and Remote Volume Mirroring (RVM). FlashCopy and RVM are premium features that you must purchase.

3.1.1 Supported host systems

At the time this redbook was written, the current publicly available version of Storage Manager was 8.2 with support on FAStT700 for:

- RedHat Linux Versions 7.1 and 7.2 (kernel version 2.4.9-21)
- SuSE Linux Version 7.3 (kernel version 2.4.10)
- TurboLinux Version 7.0 (kernel version 2.4.9-3)
- Microsoft Windows NT4 (with service pack 6a)
- Microsoft Windows 2000 (with service pack 2)
- Novell NetWare 5.1 and 6

Sometime during the second half of 2002, support is due to be added for:

- IBM AIX
- HP-UX
- Sun Solaris

IBM Storage Manager V8.2 is available as a chargeable upgrade for FAStT200 and FAStT500. Note that on FAStT200, Remote Volume Mirroring is not available. For an update on supported operating systems, servers, firmware, and drivers, see the FAStT Storage Family Web site:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

You should also check this support Web site for:

- Supported operating system levels and any required patches when support for these UNIX systems becomes available.
- Cluster solutions support. Currently only Microsoft Cluster Services (MSCS) and Novell Cluster Services are supported.

3.1.2 What's new

This section describes the new premium features and new standard features that are part of the IBM FAStT Storage Manager Version 8.2 software.

There are two premium features that can be enabled by purchasing a premium feature key:

- *FlashCopy* supports the creation and management of *FlashCopy logical drives*. A FlashCopy logical drive is a logical point-in-time image of another logical drive, called a *base logical drive*, in the storage subsystem. A FlashCopy is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space.

  Because a FlashCopy is a host addressable logical drive, you can perform backups using FlashCopy while the base logical drive remains online and user-accessible. In addition, you can write to the FlashCopy logical drive to perform application testing or scenario development and analysis. The maximum FlashCopy logical drives allowed is one half of the total logical drives supported by your controller model.

- *Remote Volume Mirroring (RVM)* is used for online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or unrecoverable error on one storage subsystem, the Remote Volume Mirror option enables you to promote a second storage subsystem to take over responsibility for normal input/output (I/O) operations. When the Remote Volume Mirror option is enabled, the maximum number of logical drives per storage subsystem is reduced.

In addition to the new premium features, the storage-management software includes the following new standard features:

- **Dynamic logical-drive expansion:** Enables you to increase the capacity of an existing logical drive. You can use the free capacity of an existing array or add unconfigured capacity (new or unassigned drives) to that array. You can expand a logical drive dynamically without losing access to it or to any other logical drives.
- **2048 logical drive support**: Enables you to increase the number of defined logical drives up to 2048 for each FAStT700 storage subsystem. The number of logical drives is reduced to 256 if Remote Volume Mirroring is enabled. (On FAStT500, this number is 512 and 128 respectively).
- **Storage partitioning**: Gives you support for up to 64 storage partitions. Specifically, there is support for up to two host ports in each host and up to eight ports in each host group (supporting a four-way host group of dual-adapter hosts).
- **Script Engine and Command Line Interface (CLI) Language**: Expands previous version support to include:
  - Creating and deleting a FlashCopy logical drive
  - Setting FlashCopy logical drive attributes
  - Supporting dynamic logical drive expansion for standard or repository logical drives
  - Disabling or recreating a FlashCopy
  - Creating or deleting logical drive-to-LUN mappings
  - Uploading Read Link Status data from a storage subsystem to a comma delimited file on a management station
  - Downloading a feature key file to a storage subsystem to enable a premium feature
- **Read Link Status Diagnostic (RLS)**: Enhances controller diagnostics enabling you to isolate the source of link problems on a Fibre Channel loop. During communication between devices, RLS error counts are detected within the traffic flow of the loop. Error count information is accumulated over a period of time for every component and device on the loop. The baseline is automatically set by the controller. However, a new baseline is set manually through the Read Link Status Diagnostics window.
- **Redesigned user interface**: Supports a more integrated process of creating logical drives and then mapping them to hosts or host groups and logical unit numbers to create storage partitions. The previous Storage Partitions window is now a Mappings View that you can switch to from the Logical/Physical View. In the Subsystem Management window, the Mappings View displays storage partition topology such as the default and defined host groups, hosts, host ports in a tree, and FlashCopy logical drives.
- **Controller default IP address**: Unless a different IP address is assigned to the controller by the administrator, or if a DHCP/BOOTP server is found, controllers A and B are automatically assigned the following IP addresses:
  - Controller A (192.168.128.101)
  - Controller B (192.168.128.102)

– Subnet mask 255.255.255.0.

Tip: After being assigned a fixed IP address, you can revert the controllers back to DHCP/BOOTP by entering 0.0.0.0 as the IP address.

Figure 3-1 shows how Storage Manager has evolved.

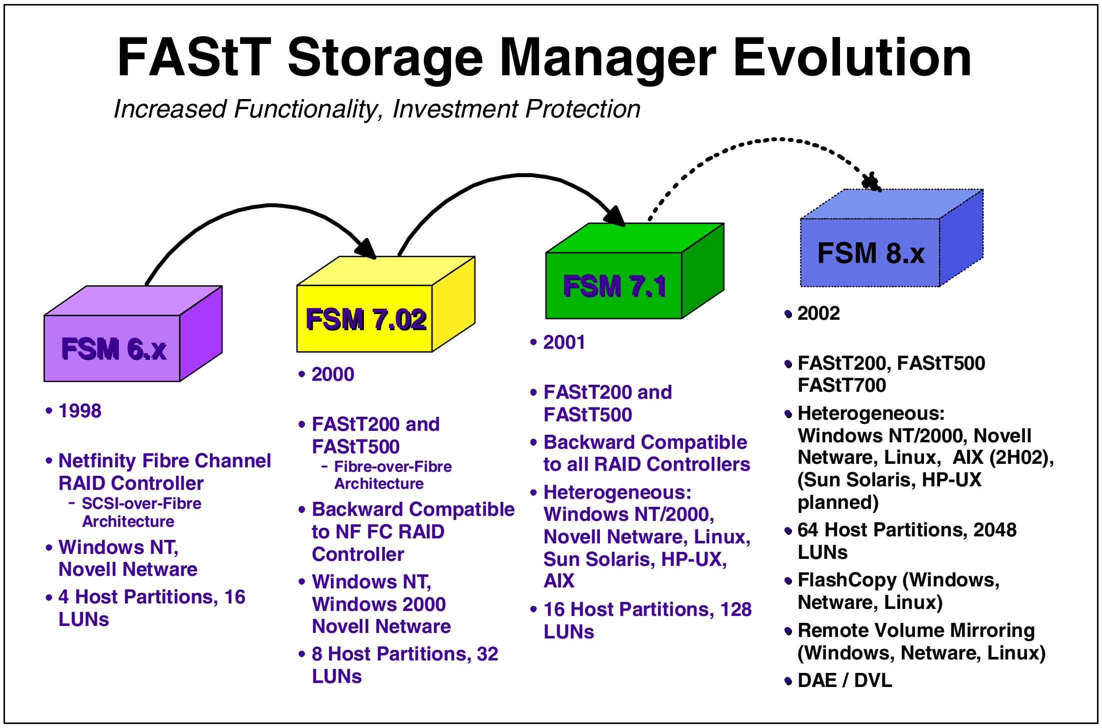

Figure 3-1 FAStT Storage Manager Evolution

3.1.3 Storage subsystem management methods

The storage-management software provides two methods for managing storage subsystems:

- Host-agent (in-band) management method
- Direct (out-of-band) management method

Depending on your specific storage-subsystem configurations and host systems, you can use either or both methods. The management methods you select determine where you need to install the software components.

Host-agent (in-band) management method

When you use the host-agent (in-band) management method, you manage the storage subsystems through the Fibre Channel I/O path to the host. The management information can either be processed in the host or passed to the management station through the network connection, as shown in Figure 3-2.

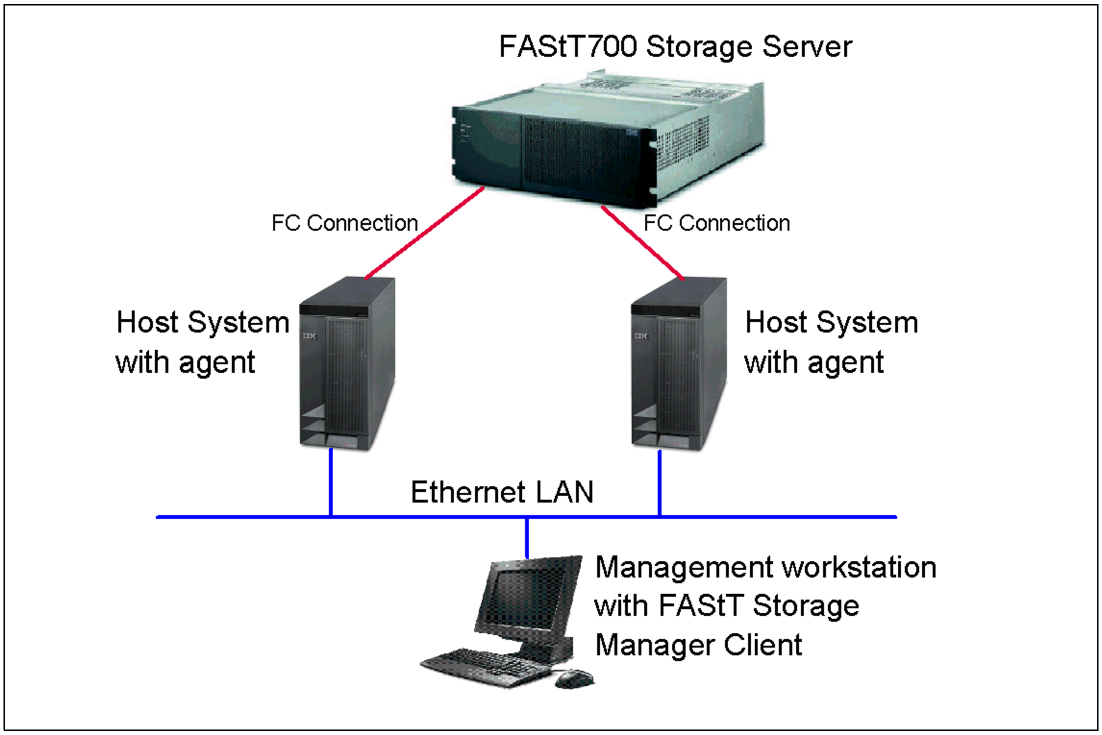

Figure 3-2 In-band management

Managing storage subsystems through the host agent has the following advantages:

- Ethernet cables do not need to be run to the controllers.
- A host name or IP address must only be specified for the host instead of for the individual controllers in a storage subsystem. Storage subsystems that are attached to the host are automatically discovered.

Managing storage subsystems through the host-agent has the following disadvantages:

- The host-agent requires a special logical drive, called the *access logical drive*, to communicate with the controllers in the storage subsystem. Therefore, you are limited to configuring one less logical drive than the maximum number that is allowed by the operating system and the host adapter that you are using. Not all operating systems support the access logical drive. In-band management is not supported on these systems.
- If the connection through the Fibre Channel is lost between the host and the subsystem, the subsystem cannot be managed or monitored.
- Security is limited to the FAStT Storage Server password. That is, anyone who installs the client can access the FAStT Storage Server.

Notes:

- The access logical drive is also referred to as the *Universal Xport Device*.
- Linux and AIX do not support the access LUN and do not support in-band management.

Important: If your host already has the maximum number of logical drives configured, either use the direct management method or give up a logical drive for use as the access logical drive.

Direct (out-of-band) management method

When you use the direct (out-of-band) management method, you manage storage subsystems directly over the network through a TCP/IP Ethernet connection to each controller. To manage the storage subsystem through the Ethernet connections, you must define the IP address and host name for each controller. The controller must be attached, and a cable must be attached to the Ethernet connectors on each of the storage subsystem controllers, as shown in Figure 3-3.

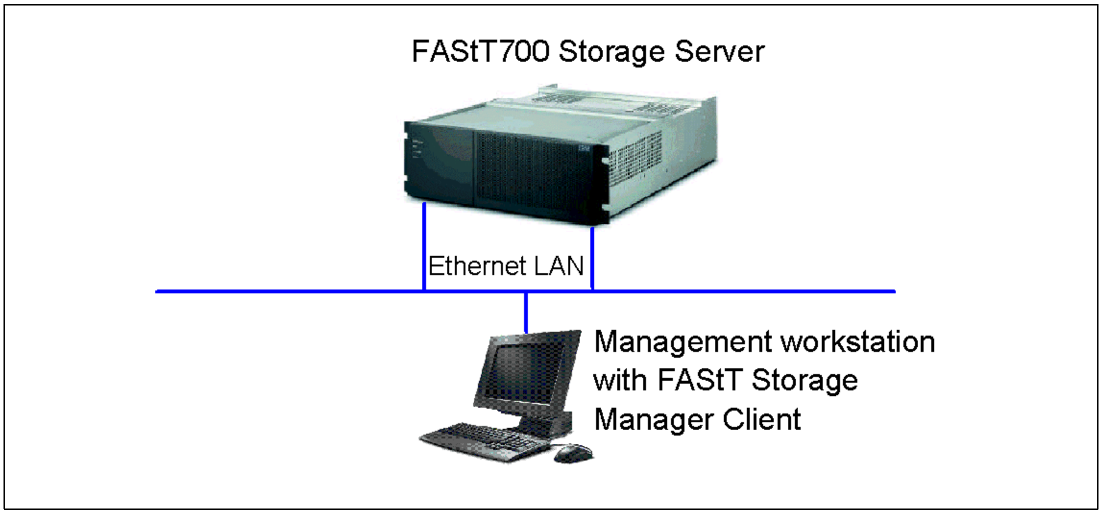

Figure 3-3 Out-of-band management

Managing storage subsystems using the direct (out-of-band) management method has the following advantages:

- The Ethernet connections to the controllers enable a management station running SMclient to manage storage subsystems that are connected to a host running any of the operating systems that are supported by Storage Manager 8.2.
- An access logical drive is not needed to communicate with the controllers. You can configure the maximum number of logical drives that are supported by the operating system and the host adapter that you are using.
- There is a constant management connection to the subsystem.
- Security is enhanced because you can restrict by IP address the workstations that are allowed to connect. It's also possible to create management LANs/VLANs and to use more advanced solutions, such as VPN, to manage the system remotely.

Managing storage subsystems using the direct (out-of-band) management method has the following disadvantages:

- You need two Ethernet cables to connect the storage subsystem controllers to a network.
- When adding devices, you must specify an IP address or host name for each controller.
- You might need a DHCP/BOOTP server, and network preparation tasks are required. You can avoid DHCP/BOOTP server and network tasks by assigning static IP addresses to the controller, or by using the default IP address.
- To assign a static IP addresses, see the IBM support Web site for RETAIN Tip H171389, "Unable To Setup Networking Without DHCP/BOOTP" at:

  `http://www.pc.ibm.com/qtechinfo/MIGR-4MMP53`

  Or refer to 4.6.11, "Network setup of the controllers" on page 129.

- If the storage subsystem controllers have firmware version 05.00.xx or later, they and FAStT700 will have the default settings shown in Table 3-1, only if no DHCP/BOOTP server is found.

Table 3-1 Fixed IP addresses

Controller	IP address	Subnet mask
A	192.168.128.102	255.255.255.0
B	192.168.128.102	255.255.255.0

3.1.4 The FAStT Storage Manager Client

The FAStT Storage Manager Client uses two main window types to give you control over your Storage Servers:

- The Enterprise Management window
- The Subsystem Management window

The Enterprise Management window

This window (Figure 3-4) appears when you start the FAStT Storage Manager Client. It displays a list of all FAStT Storage Servers that the client can access either directly or through the host agents. If you can access a certain FAStT in both ways, and possibly through several host agents, you see it listed not just once, but many times in the Enterprise Management window.

The utility can automatically detect new storage servers or you can add them to the Enterprise Management window manually. The name, status, and management type (through Ethernet or through host agent) are displayed for each listed storage server. You can also perform various tasks, like executing scripts, configuring alert notification destinations or selecting a particular storage server (or subsystem) you want to manage.

> **Note:** Although a single FAStT could appear listed several times in the left pane, accessed by various host agents or directly attached, it only appears once in the right pane.

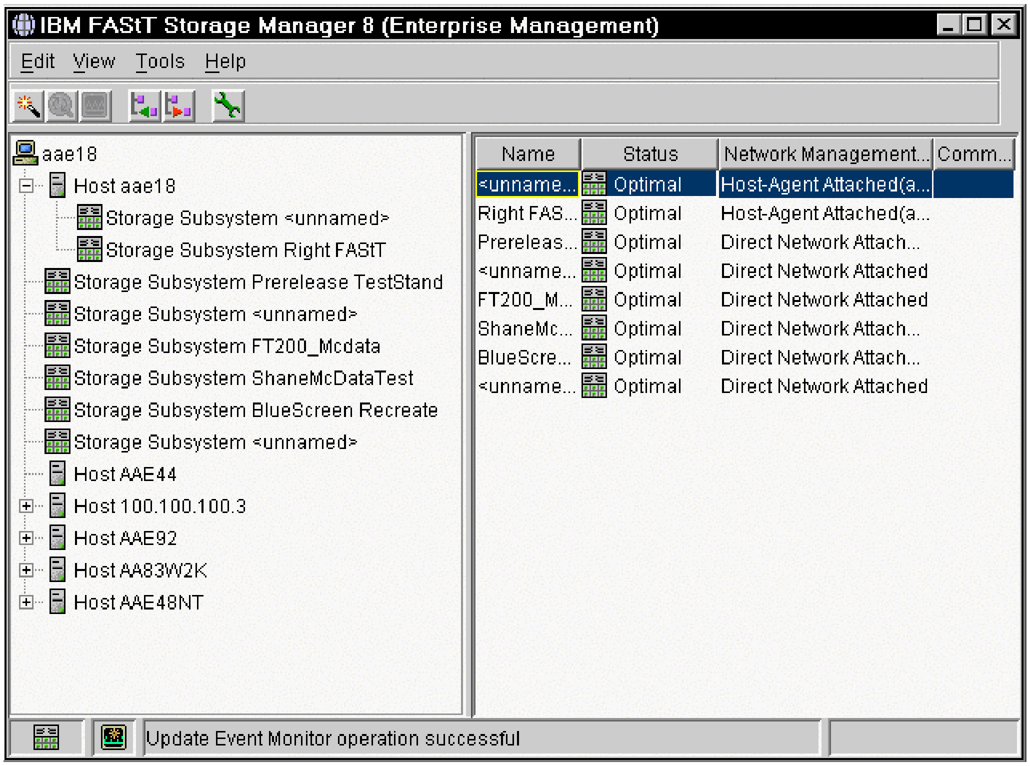

Figure 3-4 Enterprise Management window

The Subsystem Management window

Once you select a system you want to manage in the Enterprise Management window, the Subsystem Management window for that particular system opens. As you can see in Figure 3-5 and Figure 3-6, this window has two tabs that can be alternated. A specific Subsystem Management window allows you to manage one particular storage server, but you can have multiple windows open at any one time.

Logical/Physical View

The left pane shows the logical view. This is a tree-like structure that shows all arrays and logical drives configured on the storage server. You can select any array or logical drive object in the tree and perform various tasks on it.

The right pane shows the physical view. This view shows the RAID controllers installed in the storage server. It also displays the physical disk drives in all attached storage enclosures. The controllers and disk drives are selectable and enable you to perform the tasks on them.

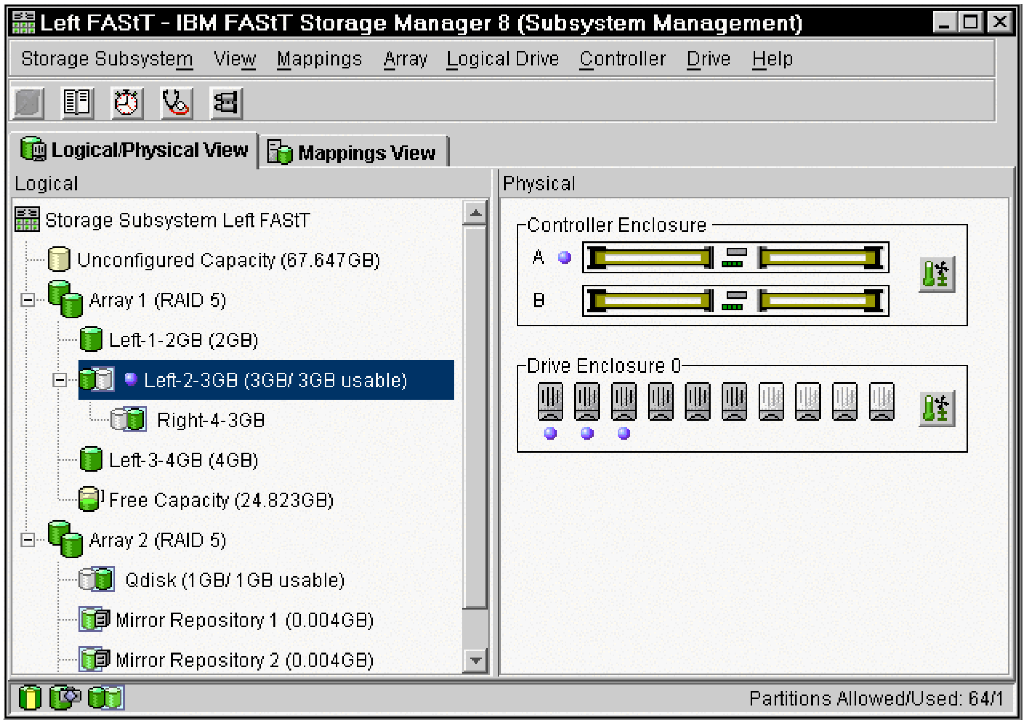

Figure 3-5 The Logical/Physical View

Mappings View

This view used to be launched as a separate window in FAStT Storage Manager 7.10. The left pane shows the defined storage partitions, hosts and host ports, while the right pane shows the LUN mappings for a particular host or host group. From this window, you can see and re-configure who is allowed to see which logical drive. Figure 3-6 shows the Mappings View of the logical drives that have been created in the FAStT Storage Server.

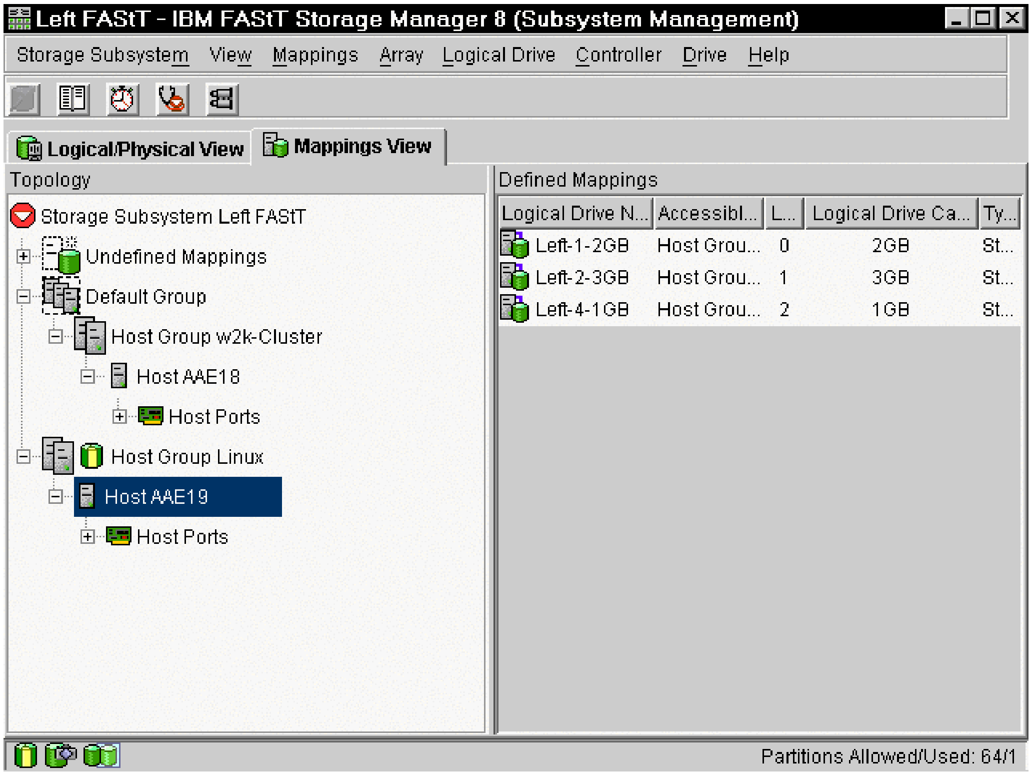

Figure 3-6 The Mappings View

3.1.5 Event Monitor

Storage Manager v7.10 introduced the Event Monitor. This monitor provides storage server monitoring functions when the Enterprise Management window is not active. It is still maintained in Storage Manager v8.2. Monitoring is performed either by Enterprise Management window or, when Enterprise Management window does not run, by the Event Monitor.

You cannot install the Event Monitor separately. It can only be installed on a managing workstation as a part of the whole FAStT Storage Manager Client package. Once installed, it runs in the background and checks for possible critical problems. If it detects a problem, it notifies a remote system through e-mail or Simple Network Management Protocol (SNMP) or both.

Use the Enterprise Management window (by clicking **Edit-> Alert Destinations**) to configure alert actions of the Event monitor. Be aware that when you monitor a FAStT Storage Server from multiple management systems, each running its own instance of the Storage Manager Client or Event Monitor, each Event Monitor sends its own alert to the configured destination when a critical event occurs. In such a case, the destination system receives several alert notifications for a single FAStT Storage Server event.

To set up alert notification to a Network Management Station (NMS) using SNMP traps, you must follow these steps:

1. Insert the IBM FAStT Storage Manager Version 8.2 CD into the CD-ROM drive on the NMS. You need to set up the designated management station only once.
2. Copy the **SM820.MIB** file from the SM820.mib directory to the NMS.

3. Follow the steps required by your NMS to compile the management information base (MIB) file. (For details, contact your network administrator, or see the documentation specific to the Storage Management product you are using.)
4. Add the NMS IP address and community name in the alert destinations section in the Enterprise Management window.

3.2 Basic management

This section describes some common administration concepts such as:

- Configuring new storage space
- Expanding the capacity of existing arrays and logical drives
- Assigning logical drives to hosts
- Creating FlashCopy drives
- Configuring Remote Volume Mirroring

3.2.1 Arrays and logical drives

An array is a set of drives that the controller logically groups together to provide one or more logical drives to an application host or cluster. An array can contain a maximum of 30 drives.

Creating arrays and logical drives

This is one of the most basic steps. Before you can start using the physical disk space, you must configure it. That is, you divide your disk drives into arrays and create one or more logical drives inside each array. When you create a logical drive from unconfigured capacity, you create an array and the logical drive at the same time. When you create a logical drive from free capacity, you create an additional logical drive on an already existing array. With the advent of larger capacity drives and the ability to distribute logical drives across controllers, creating more than one logical drive per array is a good way to use your storage capacity and protect your data.

The Create Logical Drive Wizard allows you to create one or more logical drives on the storage subsystem. Using the Wizard, you select the capacity you want to allocate for the logical drive (free capacity or unconfigured capacity) and then define basic and optional advanced logical drive parameters for the logical drive. For more information, see 4.3, “Creating arrays and logical drives” on page 100.

Important: The host operating system can have specific limits on the number of logical drives the Host can access. Consider this when creating logical drives for use by a particular host.

Note: When the Remote Volume Mirroring premium feature is in use, there are limits to the number of logical drives that are supported on a given storage subsystem. For the FAStT700 storage subsystem, up to 256 logical drives are supported. For the FAStT500 storage subsystem, up to 128 logical drives are supported.

Expanding arrays and logical drives

The ability to increase array size without needing to restart the host system is a very important feature. In today’s IT environment, the need for storage space grows constantly. Many customers exhaust their existing space sooner or later and have to expand their storage capacity. It is essential that this process is non-disruptive and does not cause any downtime.

With FAStT Storage Manager, you can simply add new disk drives to the storage server and start the expansion procedure while the system remains fully operational. Once the procedure starts, you cannot stop it. Be aware that you may see some performance impact because the expansion process competes with normal disk access. We recommend that, where possible, you carry out this type of activity when I/O activity is at a minimum. You can use this free capacity to create additional logical drives. Existing logical drives in the array do not increase in size as a result of this operation.

Attention: It is not possible to use more than 30 disk drives in one array. Once the maximum number of drives is reached, you obviously cannot add new drives anymore.

With Storage Manager 8.x, it is now possible to increase the size of the logical drive. This is called *Dynamic Volume Expansion*. The capacity of standard logical drives and FlashCopy repository logical drives may be increased using one or both of the following capacities:

- Free capacity available on the array of the standard or FlashCopy repository logical drive
- Unconfigured capacity (in the form of unused drives) on the array of the standard or FlashCopy repository logical drive

Increasing the capacity of a FlashCopy repository logical drive does not increase the capacity of the associated FlashCopy logical drive. The FlashCopy logical drive's capacity is always based on the capacity of the base logical drive at the time the FlashCopy is created.

Note: Increasing the capacity of a standard logical drive is only supported on certain operating systems. If you increase the logical drive capacity on a host operating system that is unsupported, the expanded capacity will be unusable and you cannot restore the original logical drive capacity.

The following operating systems support an increase in logical drive capacity for a standard logical drive, where logical drive-to-LUN mappings have been defined:

- Linux - ext2 and ReiserFS file systems
- NetWare
- Windows 2000 - Dynamic Disks
- Windows NT4 - Basic Disks

Tip: If a logical drive-to-LUN mapping has not yet been defined, you can increase the capacity for a standard logical drive on any host operating system.

Important: A maximum of two drives may be added at one time to increase logical drive capacity in one step, combining DAE/DVE.

A standard logical drives storage capacity cannot be increased if:

- One or more hot spare drives are in use in the logical drive.
- The logical drive has *Non-Optimal* status.
- Any logical drive in the array is in any state of modification.
- The controller that owns this logical drive is in the process of adding capacity to another logical drive (each controller can add capacity to only one logical drive at a time).
- No free capacity exists in the array, *and* no unconfigured capacity (in the form of drives) is available to be added to the array.

You can find more information in 4.6.8, “Maintaining arrays” on page 123, and in 4.6.9, “Maintaining logical drives” on page 125.

Defragmenting an array

To consolidate all free capacity on a selected array, click **Array-> Defragment**. Doing so lets you create additional logical drives from the maximum amount of free capacity. A fragmented array can result from logical drive deletion or from not using all available free capacity in a Free Capacity node during logical drive creation.

You cannot cancel the operation once it begins, but your data remains accessible during the defragment operation. To use this option, all logical drives in the array must be online and have *Optimal* status. Also, there must not be any logical drive modification operations, such as Changing the Segment Size of a Logical Drive, in progress. See 4.6.8, “Maintaining arrays” on page 123, for more information.

3.2.2 Storage partitioning

Storage partitioning adds a high level of flexibility to the FAStT Storage Server. It allows you to connect a much higher number of host systems, either in standalone or clustered mode.

Intel processor-based host systems connected to the FAStT Storage Server usually run an operating system with limited storage handling capabilities. Most of these operating systems can only treat the storage as if it was locally attached to the host system. Two or more individual host systems or clusters cannot access the same storage space, at least not without disastrous results, unless third-party file sharing software is used. This is in conflict with the idea of SAN, where the storage is supposed to be globally accessible to many host systems.

Without storage partitioning, the logical drives configured on a FAStT Storage Server can only be accessed by a single host system or by a single cluster. This can surely lead to inefficient use of storage server hardware.

Storage partitioning, on the other hand, allows you to create sets, containing the hosts with their host bus adapters and the logical drives. We call these sets *storage partitions*. Now the host systems can only access their assigned logical drives, just as if these logical drives were locally attached to them. You can connect multiple hosts or clusters running the popular Intel-platform operating systems to the same FAStT Storage Server. Storage partitioning adapts the SAN idea of globally-accessible storage to the local-storage-minded operating systems.

A storage partition contains several components:

- Logical drive mappings
- Hosts or host groups
- Host ports

A *host group* is a collection of hosts that are allowed to access the same logical drives, for example a cluster of two systems.

A *host* is a single system that can be contained in a host group.

A *host port* is the FC port of the host bus adapter in the host system. The host port is identified by its world-wide name (WWN). A single host can contain more than one host port. If you attach the servers in a redundant way (highly recommended), each server needs two host bus adapters. That is, it needs two host ports within the same host system.

The FAStT Storage Server only communicates through the use of the WWN. The storage subsystem is not aware of which host bus adapters are in the same server or in servers that have a certain relationship, such as a cluster. The host groups, the hosts, and their host ports actually reflect a logical view of the physical connections of your SAN as well as the logical connection between servers, such as clusters.

With the logical setup defined previously, mappings are specific assignments of logical drives to particular host groups or hosts.

The storage partition is the combination of all these components. It ensures proper access to the different logical drives even if there are several hosts or clusters connected.

The Default Host Group is visible to all hosts. By default, it contains the access LUN. If you map logical drives to the default host group, be aware that, unless you are using a type of file locking software (such as Tivoli SANergy), you may suffer data corruption.

Every unassigned logical drive is mapped to the undefined mappings group. This means no host (or host port, to be exact) can access these logical drives until they are mapped.

New with FAStT Storage Manager v8.2 is the ability of the FAStT Storage Servers to support up to 64 storage partitions. Every mapping of a logical drive to a new host or host group creates a new storage partition. If you map additional logical drives to the same host or host group, this does not count as a new storage partition.

For example, a cluster with two nodes with redundant I/O paths would be configured as one host group with two hosts. Each host would have two host ports for redundancy. Several logical drives would be mapped to this host group. All these components represent one storage partition. If you attach another single host system to the same storage subsystem and map a logical drive to this host, you create another storage partition. If you then define a new logical drive and map it to either the cluster or the single host, you are still using two storage partitions.

For a step-by-step guide, see 4.4, “Configuring storage partitioning” on page 105.

Mapping logical drives to the host systems

If you do not intend to use storage partitioning (for example, you only intend to attach a single host or cluster), you still have to define your host or hosts and host ports to the Default Host Group. You also have to define mappings for any logical drives created because Storage Manager 8.2 no longer defaults to putting new logical drives in the default host group.

A valid reason for not using storage partitioning is if you use file sharing software, such as Tivoli SANergy, or if you were to use one of the third-party storage virtualization packages that are now available.

If you defined multiple hosts to the default host group, each host system has equal access to all defined logical drives. Since operating systems do not usually allow multiple hosts to access the same logical drives, you must create storage partitions. You do this by mapping specific logical drives to the host ports of host systems. You have to identify the host ports by the WWN of the host bus adapters.

We recommend that you use host groups even if they only contain a single host. This simplifies the configuration so you follow the same steps for clustered hosts or single systems. The only difference is the number of hosts in the host group. Otherwise, you have some mappings to the host groups and others to individual hosts, which may look less consistent.

You create a host group first. The next step is to define a new host and its host ports. This is where you must know the WWN of your host bus adapters. Then map one or more logical drives to this host group. You must also select a LUN for each logical drive.

Attention: Be careful when changing the mapping and LUN of the access logical drive. If you use host-attached management through the host agent component, you may immediately lose the management connection when altering the access logical drive setting.

Storage partitioning considerations

There are several reasons to use storage partitioning and keep the default host group empty except for the access logical drive.

Alternatively, you can map the access logical drive to a storage partition. However, keep in mind that, in this case, you can only manage through the host systems that belong to the storage partition that contains the access logical drive. This applies only to in-band management, not out-of-band management.

Even if only one host system is attached to the storage subsystem, you should define one partition for this host. If you want to attach other host systems at a later time, you only need to define new host groups and map the logical drives to those new host groups. This process does not interfere with the existing host system. On the other hand, if you keep the original host system in the default host group, you definitely have to change this when adding additional hosts. This is obviously much more disruptive.

Attention: If you ever have to replace a host bus adapter, the WWN of the new adapter will be different. Storage partitioning assignments are based on the WWN. Since the new WWN does not appear in any of the storage partitioning assignments, after replacement, this host system will have no access to any logical drives through this adapter.

In a security-sensitive environment, you can also assign the access logical drive to a particular storage partition and ensure host-based management access only through the servers in this storage partition. In this environment, you may assign a password to the storage subsystem as well.

Heterogeneous hosts

Heterogeneous host support is maintained in Storage Manager v8.2, but is currently limited to the Intel Platform. FAStT Storage Manager 8.2 currently supports:

- Windows NT
- Windows 2000
- Novell NetWare
- Linux

Support for AIX, Sun Solaris, and HP-UX is planned.

You can use a mixture of different operating systems and clustered and non-clustered variants of the same operating systems. However, all logical drives in a single storage partition must be configured for the same operating system. Also, all hosts in that same storage partition must run the same defined operating system.

3.2.3 FlashCopy

A FlashCopy logical drive is a point-in-time image of a logical drive. It is the logical equivalent of a complete physical copy, but you create it much more quickly than a physical copy. Plus it requires less disk space. In FAStT Storage Manager 8.2, the logical drive from which you are basing the FlashCopy, called the *base logical drive*, must be a standard logical drive in your storage subsystem. Typically, you create a FlashCopy so that an application, for example a backup application, can access the FlashCopy and read the data while the base logical drive remains online and user-accessible. When the backup completes, the FlashCopy logical drive is no longer needed.

You can also create several FlashCopies of a base logical drive and write data to the FlashCopy logical drives to perform testing and analysis. Before you upgrade your database management system, for example, you can use FlashCopy logical drives to test different configurations. Then you can use the performance data provided by the Storage Management software to help you decide how to configure your live database system.

When you take a FlashCopy, the controller suspends I/O to the base logical drive for a few seconds. Meanwhile it creates a physical logical drive called the *FlashCopy repository logical drive* to store FlashCopy meta data and copy-on-write data. When the controller finishes creating the FlashCopy repository logical drive, I/O write requests to the base logical drive can continue. However, before a data block on the base logical drive is modified, a copy-on-write occurs, copying the contents of blocks that are to be modified into the FlashCopy repository logical drive for safekeeping. Since the FlashCopy repository logical drive stores copies of the original data in those data blocks, further changes to those data blocks write directly to the base logical drive without another copy-on-write. And, since the only data blocks that are physically stored in the FlashCopy repository logical drive are those that have changed since the time of the FlashCopy, the FlashCopy technology uses less disk space than a full physical copy.

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive, its capacity, and other parameters. You can disable the FlashCopy when you are finished with it, for example after a backup completes. Then, you can re-create the FlashCopy the next time you do a backup and reuse the same FlashCopy repository logical drive. Using the Disable FlashCopy and Re-create FlashCopy menu options provides a shortcut to create a new FlashCopy logical drive of a particular base logical drive. You do not need to create a new FlashCopy repository logical drive. You can also delete a FlashCopy logical drive, which also deletes the associated FlashCopy repository logical drive.

The Storage Management software provides a warning message when your FlashCopy repository logical drive nears a user-specified threshold (a percentage of its full capacity; the default is 50%). When this condition occurs, you can use the Storage Management software to expand the capacity of your FlashCopy repository logical drive from free capacity on the array. If you are out of free capacity on the array, you can even add unconfigured capacity to the array to expand the FlashCopy repository logical drive.

Important: FlashCopy logical drives can only be mapped back to the source host on Windows NT4 regular disks, Windows 2000 basic disks, and Linux standard disks. Due to operating system restrictions, FlashCopy logical drives of Windows 2000 dynamic disks *must* be mapped to a different host.

See Chapter 5, “FlashCopy and the FAStT Storage Server” on page 139, for detailed information.

3.2.4 Remote Volume Mirroring (RVM)

The Remote Volume Mirroring premium feature is used for online, real-time data replication between storage subsystems over a remote distance. The mirroring is managed by the storage subsystem controllers and is transparent to host machines and applications. You create one or more mirrored logical drive pairs that consist of a primary logical drive at the primary site and a secondary logical drive at a secondary, remote site. After you create the mirror relationship between the two logical drives, the controller owner of the primary logical drive copies all of the data from the primary logical drive to the secondary logical drive. This is called a *full synchronization*.

Note that since replication is managed on a per-logical drive basis, you can mirror individual logical drives in a primary storage subsystem to appropriate secondary logical drives in several different remote storage subsystems. Or you can create mirrors in both directions between two FAStT Storage Servers.

The secondary, remote logical drive is unavailable to host applications while mirroring is in progress. In the event of a disaster at the primary site, you can failover to the secondary site by performing a role reversal to promote the secondary logical drive to a primary logical drive. Then the recovery host can access the newly promoted logical drive, and business operations can continue.

When a primary controller (the controller owner of the primary logical drive) receives a write request from a host, the following actions occur:

1. The controller logs information about the write operation to a special logical drive called a *mirror repository logical drive*.
2. The controller writes the data to the primary logical drive.
3. The controller initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site.
4. After the host write to the primary logical drive is complete and the data is copied to the secondary logical drive at the remote site, the controller removes the log record on the mirror repository logical drive.
5. The controller sends an I/O completion indication back to the host system.

Because the controller does not send the I/O completion to the host until the data is copied to both the primary and secondary logical drives, this mirroring operation is called *synchronous*.

When write caching is enabled on either the primary or secondary logical drive, the I/O completion is sent when data is in the cache on the side (primary or secondary) where write caching is enabled. When write caching is disabled on either the primary or secondary logical drive, then the I/O completion is not sent until the data is stored to physical media on that side.

When a storage server in a mirror receives a read request from a host system, the read request is handled by the primary controller, as though no mirror were present, and no communication takes place between the primary and secondary storage subsystems.

Sometimes a primary controller receives a write request from a host that it can write to the primary logical drive, but a link interruption prevents communication with the secondary controller. In this case, the remote write cannot complete to the secondary logical drive, and the primary and secondary logical drives are no longer appropriately mirrored. The primary controller transitions the mirrored pair into an *unsynchronized* state and sends an I/O completion to the primary host. The primary host can continue to write to the primary logical drive but remote writes do not take place.

When connectivity is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, a full synchronization takes place. The mirrored pair transitions from an *unsynchronized* state to a *synchronization in progress* state.

The primary controller also marks the mirrored pair as unsynchronized when a logical drive error on the secondary side prevents the remote write from completing. For example, an offline or a failed secondary logical drive can cause the Remote Volume Mirror to become unsynchronized. When the logical drive error is corrected (the secondary logical drive is placed online or recovered to an optimal state), then a full synchronization automatically begins and the mirrored pair transitions to a *synchronization in progress* state.

A primary controller only attempts to communicate with its matching controller in the secondary storage subsystem. For example, Controller A in the primary storage subsystem only attempts communication with Controller A in the secondary storage subsystem. The controller (A or B) that owns the primary logical drive determines the controller owner of the secondary logical drive. If the primary logical drive is owned by Controller A on the primary side, the secondary logical drive is, therefore, owned by Controller A on the secondary side. If primary Controller A cannot communicate with secondary Controller A, no controller ownership changes take place.

When an I/O path error causes a logical drive ownership change on the primary side, or if the storage administrator changes the controller owner of the primary logical drive, the next remote write processed automatically triggers a matching ownership change on the secondary side. For example, if a primary logical drive is owned by Controller A and then you change the controller owner to Controller B, the next remote write changes the controller owner of the secondary logical drive from Controller A to Controller B. Because controller ownership changes on the secondary side are controlled by the primary side, they do not require any special intervention by the storage administrator.

Sometimes a remote write is interrupted by a controller reset or a storage subsystem power cycle before it can be written to the secondary logical drive. The storage subsystem controller does not need to perform a full synchronization of the mirrored logical drive pair in this case. A controller reset causes a controller ownership change on the primary side from the preferred controller owner to the alternate controller in the storage subsystem. When a remote write has been interrupted during a controller reset, the new controller owner on the primary side reads information stored in a log file in the preferred controller owner's mirror repository logical drive. The information is used to copy the affected data blocks from the primary logical drive to the secondary logical drive, eliminating the need for a full synchronization of the mirrored logical drives.

Like other premium features, the Remote Volume Mirroring feature is enabled by purchasing a feature key file from IBM or your IBM Business Partner. You must enable the feature on both primary and secondary storage subsystems.

Unlike other premium features, you must also activate the feature after you enable it, using the Activate Remote Mirroring Wizard in the Subsystem Management window. Each controller in the storage subsystem must have its own mirror repository logical drive for logging write information to recover from controller resets and other temporary interruptions. The Activate Remote Mirroring Wizard guides you to specify the placement of the two mirror repository logical drives (on newly created or existing free capacity in the storage subsystem).

After you activate the feature, one Fibre Channel host side I/O port on each controller is solely dedicated to Remote Volume Mirroring operations. No host-initiated I/O operations are accepted by the dedicated port. I/O requests received on this port are accepted only from remote controllers that are participating in Remote Volume Mirroring operations with the controller.

Important: Dedicated Remote Volume Mirroring ports must be attached to a Fibre Channel Fabric environment with support for the Directory Service and Name Service interfaces.

You can use a Fabric configuration that is dedicated solely to the remote logical drive mirroring ports on each controller. In this case, host systems can connect to the storage subsystems using Fabric, FC-AL, or Point-to-Point configurations that are totally independent of the dedicated Remote Volume Mirroring fabric.

Alternatively, you can use a single Fibre Channel Fabric configuration that is split into zones for both the Remote Volume Mirroring connectivity and for the host I/O paths to the controllers.

The maximum distance between primary and secondary sites is 10 km., using single mode Fiber and Optical Long-Wave Giga-bit Interface Converters (GBICs) or Small Form Factor Pluggable Modules (SFPs), depending on the switches in use.

The following restrictions apply to mirrored logical drive candidates and storage subsystem mirroring:

- RAID level, caching parameters, and segment size can be different on the two mirrored logical drives.
- The secondary logical drive needs to be at least as large as the primary logical drive.
- The only kind of logical drive that may participate in a mirroring relationship is a standard logical drive. FlashCopy logical drives cannot participate.
- You can take a FlashCopy of a primary logical drive, but not of a secondary logical drive. Role reversals that cause a primary logical drive to reverse to a secondary logical drive automatically fail all associated FlashCopies.
- A given logical drive may participate in only one mirror relationship.

On a given FAStT700 storage subsystem, up to 16 logical drives may participate in mirror relationships. On the FAStT700 storage subsystem, up to 256 logical drives are supported when the Remote Volume Mirroring premium feature is used.

On a given FAStT500 storage subsystem, up to eight logical drives may participate in mirror relationships. On the FAStT500 storage subsystem, up to 128 logical drives are supported when the Remote Volume Mirroring premium feature is being used.

See Chapter 6, “Remote Volume Mirroring” on page 191, for detailed information.

3.3 Advanced management

This section describes some advanced management functions of the FAStT Storage Manager. These include the concepts of data protection, tuning and performance, the FAStT Utilities, and diagnostics and troubleshooting. It also covers command line support.

3.3.1 Data protection

Different types of data protection can be implemented on the FAStT Storage Server. Which type you decide to use depends on application and performance requirements.

RAID levels

Redundant Array of Independent Disks (RAID) describes a storage solution in which multiple disks are bundled together and behave like one to create a large capacity drive or to improve performance or to provide redundancy. RAID relies on a series of configurations, called *levels*, to determine how user or redundancy data is written and retrieved from the drives.

FAStT Storage Manager 8.2 offers four formal RAID level configurations – RAID level 0, 1, 3, and 5. Each RAID level provides different redundancy and performance features. Select the level that best meets requirements for data availability and performance in your particular case.

You configure a single RAID level across a single array. All redundancy data for that array is stored within the array. The capacity of the array is the aggregate capacity of the member drives, minus the capacity reserved for redundancy data. The amount of capacity needed for redundancy depends on the RAID level used.

RAID 0

RAID 0 is a technique that stripes data evenly across all disk drives in the array. Strictly, it is not a RAID level, because no redundancy is provided. On average, accesses are random, keeping each drive equally busy. SCSI has the ability to process multiple, simultaneous I/O requests, and I/O performance is improved because all drives can contribute to system I/O throughout.

Since RAID 0 has no fault tolerance, when a single drive fails, the entire array becomes unavailable. RAID 0 offers the fastest performance of any RAID strategy for random commercial workloads. RAID 0 also has the lowest cost of implementation because no parity informations must be stored. Single disks can be configured as RAID 0, referred to as JBOD (simply a bunch of disks).

RAID 0 is normally not recommended for storing server data files. Because the data most of the time is so important to your business, losing that data could be devastating. Because a RAID 0 array does not protect you against a disk failure, you should not use it for any critical system component, such as the operating system, files, database files, or transaction log files. It may be used for data that need a very fast access but can easily be restored. If you cannot tolerate downtime, do not use RAID 0.

RAID 1

RAID 1 traditionally means the mirroring of data between two disk drives. But on the FAStT Storage Server, you are not limited to just two disk drives for RAID 1. You can use any even number up to 30. This implementation is actually called *RAID 10* and it combines RAID 1 (data mirroring) with RAID 0 (data striping).

RAID 1 offers both fault tolerance and very good performance, but the disk space efficiency is lower than with RAID 5. It also performs well with a failed disk drive, because it uses mirroring for data protection. This is generally faster than RAID 5, where the controller must read from all surviving disk drives and then calculate the missing data. Write performance is somewhat reduced because both drives in the mirrored pair must complete the write operation. For example, two physical write operations must occur for each write command generated by the operating system.

RAID 1 is a great solution when one disk drive can hold all of the data. Some recommendations for using RAID 1 are:

- Use RAID 1 for the disks that contain your operating system. It is a good choice because the operating system can usually fit on one disk.

- Use RAID 1 for transaction log. Typically, the database server transaction log can fit on one disk drive. In addition, the transaction log performs mostly sequential writes. Only rollback operations cause reads from the transaction logs. Therefore, you can achieve a high rate of performance by isolating the transaction log on its own RAID 1 array.
- Use write caching on RAID 1 arrays. Because a RAID 1 write does not complete until both writes are done (two disks), a write cache can improve performance of writes.

RAID 5

RAID-5 offers an optimal balance between price and performance for most commercial server workloads. RAID 5 provides single-drive fault tolerance by implementing a technique called *single equation single unknown*. This technique says that if any single term in an equation is unknown, the equation can be solved to exactly one solution.

The RAID 5 controller calculates a checksum using a logic function known as an *exclusive-or (XOR) operation*. The checksum is the XOR of all data elements in a row. The XOR result can be performed quickly by the RAID controller hardware and is used to solve for the unknown data element.

A significant benefit of RAID 5 is the low cost of implementation, especially for configurations requiring a large number of disk drives. To achieve fault tolerance, only one additional disk is required. The checksum information is evenly distributed over all drives, and checksum update operations are evenly balanced within the array. When a disk drive fails, the data within RAID 5 array is obviously still accessible. Reading performance decreases, because the data that resided on the failed disk drive must now be calculated from corresponding data blocks on all other disks. Instead of reading just one data block, the controller needs to read a block from each surviving disk drive.

RAID 5 usually performs better in a read-intensive environment, than in write-intensive. Write performance of RAID 5 arrays is usually not the best, due to the fact that the controller must issue a couple of reads before it can write the data. This can be improved by efficient caching. In general, uncached RAID 5 performance is 30 to 50% lower than with RAID 1.

RAID 3

RAID 3 is similar to RAID 5 in that it stripes data and parity across a set of disks. However, it is more suitable for large data transfers than RAID 5.

You may want to configure several arrays that use different RAID levels on the same FAStT Storage Server to accommodate various types of data with different performance and availability requirements.

Note: All arrays, regardless of the RAID level, are restricted to 30 drives.

Migrating RAID levels

Changing the RAID level of an array is performed in a non-disruptive manner. The system remains fully operational while the process takes place. A few possible reasons why customers may want to do this operation are:

- The storage requirements have changed over time and existing RAID levels are no longer optimal for a particular environment.
- The performance tuning process has indicated that a different RAID level would be more appropriate than the existing one.

You can change any RAID level to any other one. Be aware there are some restrictions that apply to the new arrays:

- RAID 1 or 10 requires an even number of disk drives.
- RAID 3 and 5 require at least three drives.
- There is a limit of 30 drives per array.

There are limitations if you do not have enough free space in the array. For example, a RAID 5 array of four disk drives with no free space cannot be migrated directly to RAID level 1. If you start this migration, you will receive an error message stating that you do not have enough free space. You need to add new drives to the array first to increase the free capacity and then you can change the RAID level. Also if the array has an odd number of drives and you want to migrate to RAID 1, you must add a drive first to have an even number.

By doing the opposite, changing from RAID 1 to RAID 5, you gain free space in the array that can be used to define new logical drives or expand existing ones.

Starting this procedure is simple. Use FAStT Storage Manager to select your array, and perform the RAID level migration. When the procedure starts, it reorganizes the data segments in the array according to the new RAID level.

Because this requires a large amount of I/O to be performed, there is an impact on performance while the migration lasts. You can influence the performance impact by changing the value of the modification priority. This parameter is set on a logical drive basis and you should change it for all logical drives in the array. The higher modification priority means the shorter migration time, but the performance impact will be higher. You may change the modification priority to a low value during the migration process to minimize performance degradation. When the migration finishes, change the value to a higher one again to reduce the time for a rebuild in the case of a drive failure. This minimizes the critical time of non-redundant operation caused by the disk drive fault.

Once the migration starts, you cannot stop it.

Cache memory

RAID 5 arrays are slow for write operations. A write to a basic RAID 5 array actually generates two reads from the disks and two writes to the disks. A write to a RAID 5 array generates four physical I/O operations to the disks. Therefore, the write capacity of RAID 5 array is equivalent to the capacity of a quarter of the disk drives in the array. That is unless the RAID 5 implementation uses write caching.

Write caching can increase I/O performance, but it requires a battery backup. Use write cache mirroring to protect data during a controller or cache memory failure. When you enable write cache mirroring, cached data is mirrored across two redundant controllers with the same cache size. (Data left in cache is protected by an on board cache battery.) The data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other controller can complete all outstanding write operations.

Sometimes write caching is disabled when batteries are low or discharged. If you enable the Write Caching Without Batteries parameter on a logical drive, write caching continues even when batteries in the controller enclosure are discharged. If you do not have an Uninterruptible Power Supply (UPS) for power protection, do not enable this parameter. Otherwise data in the cache will be lost during a power outage when the controller enclosure does not have working batteries.

To prevent data loss or corruption, the controller periodically writes cache data to disk (flushes the cache). When the amount of unwritten data in cache reaches a certain level, called a *start percentage*, the software signals the controller to write the data to disk. The controller writes to disk until the amount of data in cache drops to a stop percentage level. For example, you can specify that the controller start flushing the cache when the cache reaches 80% full and stop flushing the cache when the cache reaches 16% full.

For maximum data protection, you can choose low start and stop percentages. However, in both cases, this increases the chance that data needed for a host read will not be in the cache, which decreases performance. Choosing low start and stop percentages also increases the number of disk writes necessary to maintain the cache level, increases system overhead, and further decreases performance.

For protection against power outages, batteries in the controller enclosure protect data in the cache that has not been written to disk. Change the controller enclosure batteries at the recommended time intervals. The Storage Management software features a battery age clock that you can set when you replace a battery. This clock keeps track of the age (in days) of the battery so you know when it is time to replace the battery. You receive critical alert notification when the battery is nearing expiration and when it has reached expiration.

Hot spare disks

Hot-spare disk drives provide additional protection that might prove to be essential in case of a disk drive fault in a fault tolerant array (RAID 1, 3, or 5). A hot-spare drive is like a replacement drive installed in advance. The data from the failed disk drive is automatically rebuilt to the hot spare when one exists. Hot-spare drives are global. They can replace a disk drive in any array on the storage server. When assigning disks as hot spares, make sure they have enough storage capacity. If the failed disk drive is larger than the hot spare, reconstruction is obviously not possible.

If a drive fails, the controller attempts to find a global hot spare on the same channel as the failed drive. The global hot spare must be at least as large as the configured capacity of the failed drive.

If a global hot spare does not exist on the same channel, or if it is already in use, the controller checks the remaining global hot spare drives, beginning with the last global hot spare configured. For example, the drive at location 1:4 (EXP tray number 1, disk drive number 4) may have failed and the global hot spare drives may be configured in the order 0:12, 2:12, 1:12, 4:12, 3:12. In this case, the controller checks the global hot spare drives in the order 1:12, 3:12, 4:12, 2:12, 0:12. Then it selects the first available global hot spare drive closest to the configured capacity of the failed drive. Typically there is a limit of 15 global hot spare drives per FAStT700.

Media scan

Magnetic media errors can cause various problems in RAID arrays, including complete data loss. Media scan is a background process that checks all physical disk drives in the array for possible media errors. FAStT Storage Manager allows you to run media scan. You can also configure how to run this process. For example, you can enable redundancy check.

Fault-tolerant RAID levels 1, 3, and 5 provide protection against a single disk drive failure. If the disk drive fails, its data is rebuilt to a replacement drive, using information from other drives in the array. This only works if all data and either mirrored or parity blocks are correct on the surviving disk drives.

An error on either of those drives causes data loss because it is not possible to rebuild the data that resided on the failed disk drive. Therefore, it is vital to ensure that there are no magnetic media errors on the disk drives in the array and that all redundancy information corresponds to the data blocks.

Auto Logical Drive Transfer (ADT)

I/O data path protection to redundant controllers in a storage subsystem is accomplished with the ADT feature and a host multi-path driver. A multi-path driver, such as Redundant Disk Array Controller (RDAC), is an I/O path failover driver installed on host computers that access the storage subsystem. ADT is a built-in feature of the controller firmware that allows logical drive-level failover rather than controller-level failover.

With Storage Manager 8.2, ADT is set by the host type and on a per LUN basis. This means heterogeneous support is now extended across all operating system types. (With FAStT Storage Manager 7.10, ADT had to be disabled on a controller basis if an operating that didn't support ADT, such as AIX, was used. This restricted heterogeneous support.)

A pair of active controllers are located in a storage subsystem. When you create a logical drive, you assign a controller to own the logical drive (called *preferred controller ownership*) and to control the I/O between the logical drive and the application host along the I/O path. The preferred controller normally receives the I/O requests to the logical drive. If a problem along the data path (such as a component failure) causes an I/O operation to fail, the multi-path driver issues the I/O request to the alternate controller. When ADT is enabled and used in conjunction with a host multi-path driver, it helps ensure an I/O data path is available for the storage subsystem logical drives. The ADT feature changes the ownership of the logical drive receiving the I/O to the alternate controller.

Once the I/O data path problem is corrected, the preferred controller automatically re-establishes ownership of the logical drive as soon as the multi-path driver detects the path is normal again. All current supported operating systems fit into this category. However Linux uses the FAStT Management Suite Java (MSJ) software, instead of the RDAC. It does not support the use of the access LUN, so the access LUN must be removed from any Linux storage partition. For this reason, Linux does not support in-band (host agent) management. NetWare also does not use an RDAC, but its multi-path driver is compatible with the access LUN.

Password

Executing destructive commands on a storage subsystem can cause serious damage, including data loss. A specified password protects any options that the controller firmware deems destructive. (These options include any functions that change the state of the storage subsystem, such as the creation of logical drives, modification of cache settings, and so on.) Without password protection, *all* options are available within this Storage Management software.

When using passwords, consider these points:

- If no password was set previously, no current password is required to establish a new password.
- The password is stored on the storage subsystem. Each storage subsystem that you want to be password protected needs a password.
- You cannot change a storage subsystem password unless you have supplied the current password first. If you have forgotten the password, contact IBM Technical Support.
- The maximum password length is 30 characters.
- Passwords are case sensitive. Remember your use of uppercase and lowercase letters when you change a password.
- Trailing white spaces are not stripped from passwords. Be careful not to include trailing spaces in the new password, because they can be difficult to enter accurately in the future.
- Only asterisks are displayed when you type a password.

3.3.2 Tuning and performance

This section discusses how you can tune FAStT700 to gain maximum performance. It explains a number of different parameters that can effect overall system performance.

Performance monitor data

The performance monitor data can be used to make storage subsystem tuning decisions. The data that is collected includes:

- Total I/Os
- Read percentage
- Cache hit percentage
- Current KB/sec and maximum KB/sec
- Current I/O per sec and maximum I/O per sec

See 8.1.1, "Performance Monitor" on page 250, for detailed information.

Cache parameters

The FAStT Storage Manager utility enables you to configure various cache settings, including:

- Start and stop cache flushing levels
- Cache block size
- Enabling or disabling write cache mirroring
- Cache read-ahead multiplier
- Write-back and write-through mode

These settings have a large impact on the performance of the FAStT Storage Server and on the availability of data. Be aware that performance and availability are often in conflict with each other. If you want to achieve maximum performance, in most cases, you must sacrifice system availability and vice versa. Most customers are usually more motivated by the highest availability they can get, rather than the maximum performance.

See 8.1.2, "Cache parameters" on page 252, for detailed information.

Controller ownership

You can statically load balance between controllers by monitoring I/O and re-assigning preferred ownership of LUNS. This is discussed further in 8.1.3, "Controller ownership" on page 254.

Segment size

The segment size is specified per each logical drive. You should choose the segment size very carefully, because it can have a large impact on performance. You can find more information in 8.1.4, "Segment size" on page 254.

RAID levels

Performance varies based on the RAID level set. The Performance Monitor should be used to obtain logical drive application read and write percentages. For a brief summary of RAID levels, refer to 8.1.5, "RAID levels" on page 254.

Logical drive modification priority

The modification priority defines how much processing time is allocated for logical drive modification operations relative to system performance. You can increase the logical drive modification priority, although this may affect system performance. For more information, see 8.1.6, "Logical drive modification priority" on page 255.

Remote Volume Mirroring

When a storage subsystem logical drive is a primary logical drive and a full synchronization is necessary, the controller owner performs the full synchronization in the background. Because the full synchronization diverts controller processing resources from I/O activity, it can impact performance on the host application. For details on settings, see 8.1.7, “Remote Volume Mirroring” on page 255.

3.3.3 FAStT Utilities

FAStT Storage Manager v8.2 comes with three command line utilities that are installed separately from the other components. These vary by operating system type but generically are:

- **hot_add**: This utility is used to scan for new disks available to the operating system after they are defined and mapped in FAStT Storage Manager. This is especially useful for operating systems that normally have to be re-booted, such as Windows NT4.
- **SMdevices**: This utility lists all logical drives available to the host, including target ID and logical drive name (as defined in the FAStT Storage Manager). This is useful if you have several logical drives of the same size defined for a given host because it allows you to identify which is which before mounting and formatting under the operating system.
- **SMflashcopyassist**: This utility has two uses:
  – Running it against a specific drive or mount point causes the buffers to be flushed to disk.
  – For Windows NT4 only, it writes a new disk signature so the drive can be used on the same host.

3.3.4 Script Editor and command line interface

Many storage management options available through the Subsystem Management window can be sent to the storage subsystem using statements in scripts. You can use the Script Editor to create or edit a script file, save a script file to the Storage Management station’s local disk, or load a script file from disk. You can also use the command line interface (CLI) to issue individual commands to the scripting engine from the host operating system command shell or to call complete pre-written scripts.

Script Editor

To open the Script Editor, follow these steps:

1. Select a storage subsystem in the Device Tree View or Device Table.
2. Click **Tools-> Execute Script** as shown in Figure 3-7, or right-click and select **Execute Script**.

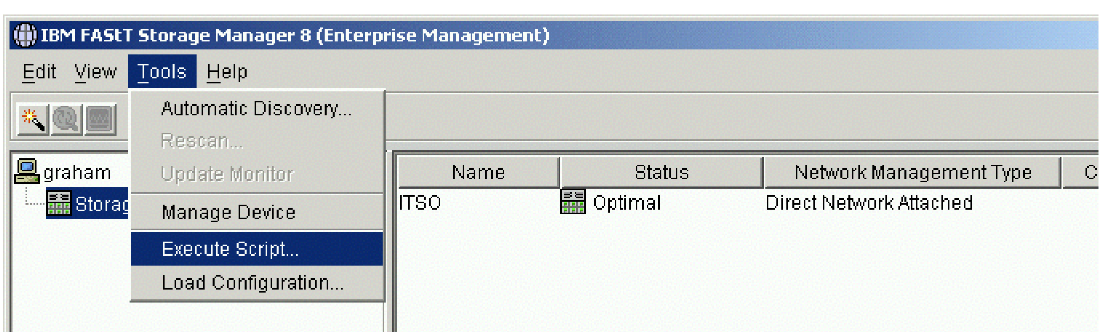

Figure 3-7 Starting the Script Editor

The Script Editor opens as shown in Figure 3-8. There are two views in the window:

- **Script view**: Provides an area for inputting/editing script commands.
- **Output view**: Displays verification or execution results.

A splitter bar divides the window between script view and output view. You can use the splitter bar to resize the views.

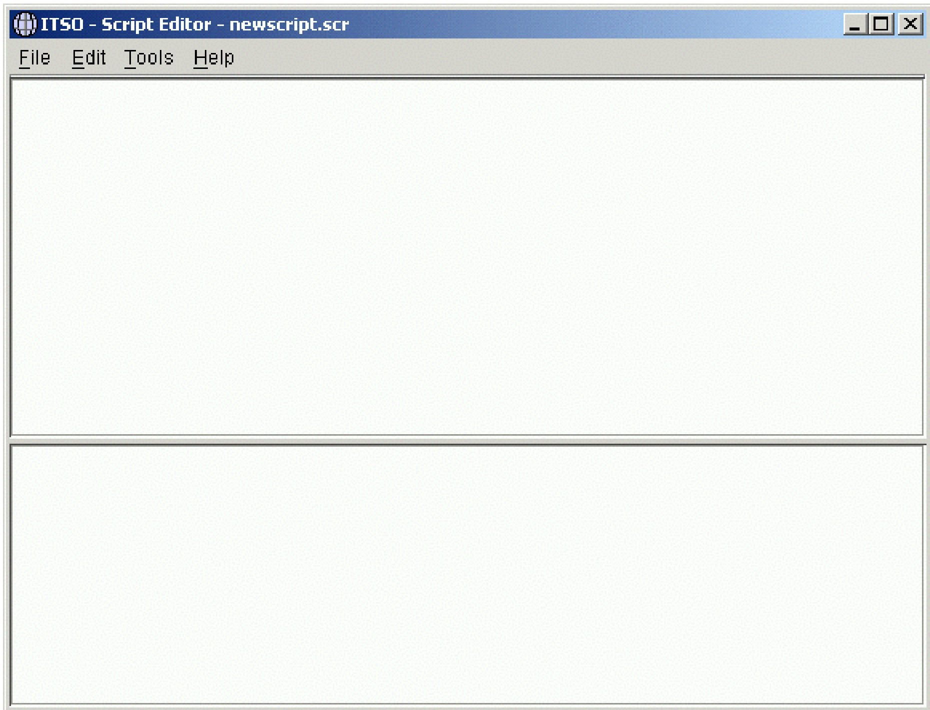

Figure 3-8 The Script Editor

Usage guidelines

Consider the following guidelines when using the Script Editor:

- All statements must end with a semicolon (;).
- Each base command and its associated primary and secondary parameters must be separated with a space.
- The Script Editor is not case sensitive.
- Each statement must be on a separate line.
- Comments can be added to your scripts to make it easier for you and future users to understand the purpose of the command statements.

Adding comments to a script

The Script Editor supports the following comment formats:

- Text contained after two forward-slashes // until an enter character is reached.

  For example, the comment "The following command assigns hot spare drives." is included for clarification and is not processed by the Script Editor:

  ```
  //The following command assigns hot spare drives.
  set drives [1,2 1,3] hotspare=true;
  ```

> **Important:** You must end a comment beginning with // with an end-of-line character, which you insert by pressing the Enter key. If the script engine does not find an end-of-line character in the script after processing a comment, an error message is displayed and the script execution is terminated. This error commonly occurs when a comment is placed at the end of a script and you forget to press Enter.

- Text contained between the characters /* and */.

  For example, the comment "The following command assigns hot spare drives." is included for clarification and is not processed by the Script Editor:

```
/* The following command assigns hot spare drives.*/
set drives [1,2 1,3] hotspare=true;
```

> **Important:** The comment must start with /* and end with */. If the script engine does not find both a beginning and ending comment notation, an error message is displayed and the script execution is terminated.

Using the show statement

Use the show statement to embed comments in your script that display in the output view during script execution.

For example, including a `Show` **"**`setting controller mode`**"** statement in your script results in the display of `setting controller mode` in the output view when this line is processed during script execution.

Interpreting the script execution results

During script execution, messages are displayed in the output view beginning with:

```
Executing script...
```

After a successful script execution, you see the message:

```
Script execution complete.
```

If there is an error during the parse phase, an error is displayed in the output view giving the line and column number and a description of the syntax error.

If there is an error during execution, a message is displayed in the output view stating that the command failed and reporting a description of the error.

> **Important**: Certain execution errors, including the inability to communicate with the storage subsystem, always causes script execution to halt. In these cases, execution stops even if you used the On Error Continue statement.

Command line interface

The command line interface gives you direct access to the script engine from the command line shell of your operating system. The CLI provides an efficient way to edit, send, and execute Storage Management commands on multiple network storage subsystems. You can also access the script engine using the Enterprise Management window. In doing so, you can only edit or execute script commands on the storage subsystem that you selected in the Enterprise Management window instead of multiple storage subsystems.

The command line interface can be used to:

- Access the script engine directly instead of through the Enterprise Management window.
- Select multiple storage subsystems on which to execute script commands.
- Supply script language commands either directly to the command line interface screen or to a script file.
- Create batch files of script commands for execution on multiple storage subsystems.
- Execute script commands on host-agent managed or directly managed storage subsystems, or a combination of both.
- Execute mass operations on multiple storage subsystems, such as firmware downloads and upgrades.
- Display configuration information about the network storage subsystems.
- Add storage subsystems to the management domain.
- Perform an automatic discovery of all storage subsystems attached to the local subnet.
- Add or delete SNMP trap destinations and e-mail alert notifications.
- Specify the mail server and sender e-mail address or SNMP server for alert notifications.
- Display the alert notification settings for storage subsystems currently configured in the Enterprise Management window.
- Direct the output to standard command line display or to a named file.

To use the CLI, go to the command line shell of your operating system. At the command prompt, enter `SMcli`, followed by either the controller name, host-agent name, world-wide name (WWN) or user-supplied name of the specific storage subsystems. The name you enter depends on your storage subsystem management method:

- **Directly managed**: Enter the *hostname or IP address* of the controller(s).
- **Host-agent managed**: Enter the *hostname or IP address* of the host.

Important: You must use the -n option if more than one host-agent managed storage subsystem is connected to the host, for example:

```
SMcli hostmachine -n sajason
```

If you specify a host name, or the IP address, the command line utility verifies the existence of a storage subsystem. If you specify the user-supplied storage subsystem name or WWN, the utility ensures that a storage subsystem with that name exists at the specified location and can be contacted.

Note: If the storage subsystem is configured in the Enterprise Management window, you can specify the storage subsystem by its user-supplied name only using the -n option, for example:

```
SMcli -n Storage Subsystem London
```

The name must be unique to the Enterprise Management window.

Now you are in interactive mode and can enter one or more commands or the name of a script file. SMcli first verifies the existence and locations of the specified storage subsystems and, if applicable, the script file. Next, it verifies the script command syntax and then executes the commands. This is shown in the following examples:

- One or more commands:

```
-c "<command>;[<command2>;...]"
```

- Script file:

```
-f <scriptfile>
```

You don't have to run the CLI in interactive mode. You can pass parameters to CLI as you run it, for example:

```
SMcli -n ITSO -f scriptfile.scr -e
```

This executes the commands found in the script file named scriptfile.scr in the storage subsystem named "ITSO" without performing a syntax check.

You can use this method to create, for example, scripts that automatically create (by mixing native operating system commands with the CLI) a FlashCopy logical drive, mount it under the operating system, and perform a backup on it.

For detailed information on the CLI parameters, see Appendix C, "Command line reference" on page 317, or the Enterprise Management window online help.

Command reference

The Script Editor and the command line interface each support the use of the following twelve commands. These commands, used in conjunction with specific parameters, let you perform various storage subsystem management tasks. For detailed information on the command parameters, see Appendix C, "Command line reference" on page 317, or the Enterprise Management window online help.

create

- Create a standard logical drive on a set of drives (unconfigured capacity), turning those drives into a new array, or create logical drives on an existing array (free capacity).
- Create FlashCopy logical drives, using either unconfigured or free capacity for the associated repository logical drives.
- Create storage partition topology definitions (host group, host, or host port).
- Create a logical drive-to-LUN mapping.

delete

- Delete an array or logical drive.
- Delete a FlashCopy logical drive (and its associated repository logical drive).
- Delete a storage partition topology definition (host group, host, or host port).
- Delete a logical drive-to-LUN mapping.

disableFlashCopy

- Disable a FlashCopy logical drive.

download

- Download new controller firmware or NVSRAM to the storage subsystem.
- Download a feature key file to the storage subsystem.
- Download drive firmware or ESM firmware to the storage subsystem.

on error

- Dictate the script behavior when execution fails. The default behavior is for the script to continue executing subsequent commands after a command has failed. The other parameter you can use is `stop`, so the script engine stops after a failed command.

recreateFlashCopy

- Recreate a previously disabled FlashCopy logical drive.
- Disable and recreate a FlashCopy logical drive.

reset

- Restore logical drives to their preferred controller owners.
- Reset a controller.

set

- Change the properties of a storage subsystem component, controller, array, standard logical drive, FlashCopy logical drive, or drive (you can set multiple properties for a specific storage subsystem element using one `set` command).
- Revive a failed drive.
- Set the Performance Monitor polling interval and number of samples, to be used in conjunction with the upload command for uploading Performance Monitor statistics.
- Change the role of a mirrored logical drive from primary to secondary or from secondary to primary.

show

- Display the properties of the different logical and physical components comprising the storage subsystem.
- Embed text strings (comments) in your command line that display in the output view during command execution.
- Show storage partition topology and mappings information.

start

- Start a long-running operation. Currently, the only long-running operation you can start using this command is a logical drive expansion operation.

upload

- Upload a file containing Read Link Status statistics from the storage subsystem to your storage management station.
- Upload a file containing storage subsystem configuration data from the storage subsystem to your storage management station.
- Upload a file containing performance data from the storage subsystem to your storage management station.
- Upload a file containing events from the Major Event Log (all events or just critical events) from the storage subsystem to your storage management station.
- Upload a file containing controller state dump information from the storage subsystem to your storage management station.

use

- Enter the password to use for destructive commands. Currently there is only one option for the `use` command. This command does *not* set the password. In the `set` command, there is a password parameter for the storage subsystem. This command is only required once in a script, not in front of each destructive command.

3.3.5 Diagnostics

The FAStT Storage Server provides a high level of system availability thanks to the built-in features and components used in the system. Even now a failure may still happen. It is very important that you can identify and fix the failure as fast as possible.

The FAStT Storage Server logs all error conditions in its own event log stored in the controller. With the FAStT Storage Manager Software, this event log can be read and saved to a local system. In this event log, the error is logged including time, exact location, and cause of the error. Refer to 8.2.2, “Event Viewer” on page 257, for more information. You can find a complete list of all critical events reported in Appendix B, “Critical event descriptions” on page 307.

The *Recovery Guru*, another part of the FAStT Storage Manager, can interpret the event log and present a detailed step-by-step procedure of how to recover from a particular failure. An example of an error condition is a failed power supply in one of the drive enclosures. See 8.2.3, “Recovery Guru” on page 258, for more information.

Also the controllers have a built-in test that can check the basic connectivity within the Fibre Channel paths to identify for example cable problems or defect SFPs. See 8.3.1, “Controller runtime diagnostics” on page 260, which discusses this further.

4

Step-by-step procedures for the FAStT Storage Server

This chapter provides a step-by-step procedure on how to setup and configure the FAStT Storage Server using the FAStT Storage Manager software. This process, which is explained in the first part of the chapter, is similar for all supported operating systems. Only the installation of the drivers and the software is operating system dependent.

Then it explains how to configure the FAStT Storage Server, define arrays and volumes, and enable storage partitioning. It presents the different ways to monitor the FAStT Storage Server. Next it discusses how to maintain the FAStT Storage Server and adapt the already existing configuration to new requirements of resizing logical drives to provide more capacity or changing the segment size to meet new usage patterns.

Finally, we introduce the FAStT Management Suite Java (MSJ). It allows monitoring and configuring of the FAStT Host Bus Adapters through a client/server-based application.

Prior to reading this chapter, you should:

- Understand the basic architecture of the FAStT Storage Server and the different ways of managing it with the FAStT Storage Manager.
- Have installed your FAStT Storage Server according to Chapter 2, "Hardware details" on page 13. You must also have cabled the FAStT Storage Server in a way that a connection to each of the two controllers can be established from any management system, via a Fibre Channel or Ethernet connection.
- Have installed the host systems attached with the operating system and configured them according to your needs, except for the driver for the host bus adapter and the FAStT Storage Manager software.

Attention: Some of the operating systems that are discussed in this section, at the time this redbook was written, are not formally supported. You must check the FAStT supported servers Web site prior to connecting hosts to the FAStT Storage Server or using the advanced functions. You can locate the FAStT700 supported servers Web site at:

`http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm`

4.1 Driver and host software installation

The host software package consists of different packages as explained in 3.1, “The FAStT Storage Manager software” on page 34. The FAStT Client package and the FAStT Runtime Environment are available for all operating systems; the runtime environment may be incorporated into the client package. For Windows NT 4.0 and 2000, AIX, and Solaris, the FAStT RDAC packages provides redundancy within Fibre Channel paths. For NetWare, IBM provides a FAStT specific driver called *IBMSAN*, which allows multipath I/O and path failover. Linux uses the FAStT MSJ, which allows the definition of a preferred path and an alternate path for each logical drive, to achieve load balancing and path protection. HP-UX is multipath-I/O aware by using the PV-Links package, a part of the operating system.

Because in-band management is not supported for all operating systems, the FAStT Agent is only available for those systems that support in-band management.

Since the installation of the host software varies for each operating system, we explain the procedure for all operating systems that can be used on a FAStT. The package names used in this documentation should be considered as generic ones. You may have to replace them with the correct package names as found on the Web or CD.

Important: Required steps may change with a newer release of the drivers or software compared to the one used while writing this book. Therefore, always refer to the readme file of the driver and software packages you want to use.

It may also be necessary to change some parameters in the NVSRAM of the host bus adapter or its configuration file.

Compared to the older version 7.10 of the FAStT Storage Manager, the package contents and the file locations are slightly changed. With the new code, only one version of the Java Runtime Environment (JRE) is installed once, when you install the first package of the host software. You are also only asked once for the installation directory while installing the first package. All of the following packages are installed beneath the specified directory according to:

- **IBM_FAStT/**

  The main directory that holds all packages.

  - **Agent/**

    The directory holding the agent package.

  - **Client/**

    The directory holding the client package.

  - **JRE/**

    The directory holding the java runtime package.

  - **Util/**

    The directory holding the utility package

4.1.1 Microsoft Windows NT 4.0 and Windows 2000 installation

This section covers the installation for Windows NT 4.0 and Windows 2000 since the process is nearly the same for each one.

The host software for Windows consists of four packages:

- FAStT Client
- FAStT RDAC
- FAStT Host Agent
- FAStT Utilities

Windows NT 4.0 has another package, the Microsoft Java Virtual Machine, which must be installed.

The driver for the host bus adapter is delivered in a separate package. If two host bus adapters are installed in the system, redundancy is achieved with the FAStT RDAC package.

If you want to use the FAStT Event Monitor with SNMP, you have to install the Microsoft SNMP service first, since the FAStT Event Monitor uses its functionality.

Because you are installing new software, including new drivers, you need to log on with administrator rights.

Installing the host adapter driver

The driver for the host bus adapter is available for download at:

`http://www.storage.ibm.com/hardsoft/disk/fastt/index.html`

The packages differ for Windows NT 4.0 and Windows 2000. Be sure to check the readme for additional information regarding settings of the host adapter and driver.

Installing the driver for Windows NT 4.0

To install the Windows NT 4.0 driver, follow these steps:

1. From the Control Panel, select **SCSI Adapters**.
2. Click the **Drivers** tab.
3. To install a new driver, click **Add** to open the dialog, choose **Have Disk** and point to the driver location.
4. After you install the driver, reboot the system.

Installing the driver for Windows 2000

To install the Windows 2000 driver, follow these steps:

1. Open the **Device Manager**.
2. Look for unknown devices. The FAStT FC-2 Host Bus Adapter is not recognized by Windows 2000.
3. Highlight the unknown adapter.
4. Right-click and choose **Properties**.
5. Click the **Driver** tab and select **Update Driver**.
6. On the following window, click **Other source** and point to the location of the driver. You must follow this procedure for each adapter installed in the system.
7. After you install the driver, reboot the system.

Installing the host software

You must maintain the following order while installing the different parts of the host software:

1. Microsoft Java Virtual Machine (only for Windows NT 4.0)
2. FAStT Client for Windows (same for Windows NT 4.0 and Windows 2000)

3. FAStT RDAC for Windows (different for Windows NT 4.0 and Windows 2000)
4. FAStT Host Agent for Windows (same for Windows NT 4.0 and Windows 2000)
5. FAStT Utilities for Windows (same for Windows NT 4.0 and Windows 2000)

If you are installing the software on management station only, you only need the first two parts:

- Java Virtual Machine for Windows NT 4.0
- FAStT Client

Installing the Java Virtual Machine

For Windows NT 4.0 systems, the Java Virtual Machine package must be installed first, because FAStT Storage Manager is written in Java.

Run the *MSJavx86.exe* file and follow the instructions on the panel. Do not reboot the system at this time, but continue with the installation of the FAStT Client.

Installing the FAStT Client

Run the setup.exe file to install the software and follow the instructions on the panel.

> **Important:** When you install FAStT Client on a stand-alone host and manage storage subsystems through the Fibre Channel I/O path, rather than through the network, you must install the TCP/IP software on the host and assign an IP address to the host.

During the installation, you are prompted with the window shown in Figure 4-1 if you want to install the Event Monitor Service.

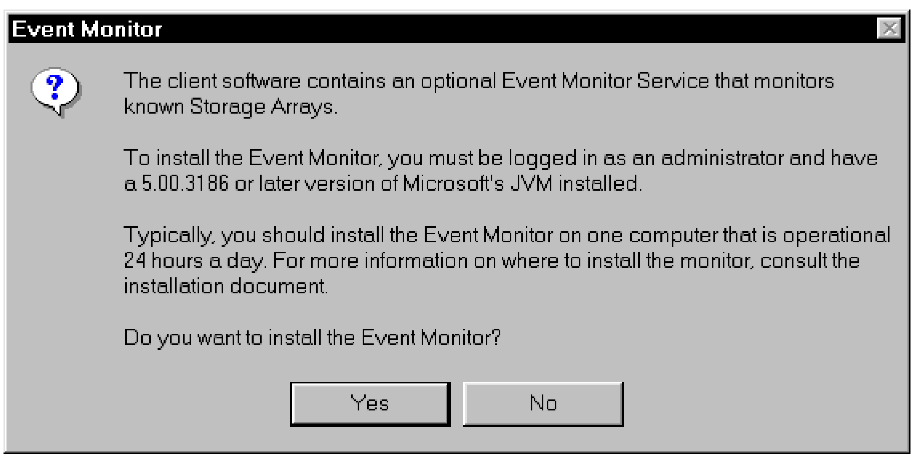

Figure 4-1 Installing the FAStT Storage Manager Event Monitor

This service provides stand-alone alerting and monitoring of the FAStT Storage Server through SNMP traps or e-mail alerts. You should install and configure this service at least on one server that is operational 24 hours and directly attached to a FAStT Storage Server either via Fibre Channel or Ethernet. If you want to monitor FAStT through Fibre Channel, the service requires the server to have access to the access logical drive.

If you install the client software on a management workstation, which is not directly attached to the FAStT Storage Server or is not operated 24 hours, you don't need to install the service.

The client package also installs the command line utility **SMcli**, which we use in the FlashCopy scenarios in Chapter 5, "FlashCopy and the FAStT Storage Server" on page 139.

Installing FAStT RDAC

The FAStT RDAC package provides failover and load balancing to the FAStT Storage Server.

Important: If you are installing on a Windows NT 4.0 system, you must assign static drive letters to existing local drives. Otherwise problems may occur during installation or upgrade. The reason is that the RDAC driver is installed before the Windows NT class driver. If static drive letters are previously assigned to local drives, the system detects FAStT drives first and assigns them drive letters before drive letters are assigned to local disks. This can lead to problems especially with installed applications.

Use the following procedure to manually assign static drive letters on host computers running Windows NT 4.0:

1. Determine how many partitions are on the system using the disk administrator.
2. Verify that the following conditions are met:
   - A boot partition exists.
   - The last partition was created with Disk Administrator.
   - The partition has an assigned drive letter with either an Unknown status or a status other than Unformatted.

If all conditions are met, you can install the FAStT RDAC package. Otherwise follow these steps:

1. Select a drive with fee space on your local system or on a storage subsystem.
2. Use Disk Administrator to create a new partition. It is created with a status of *Unformatted*, and a drive letter is assigned. The partition may be deleted at a later time without causing problems.
3. Use Commit Changes Now. The status of the partition changes to *Unknown*, and static drive letters are assigned to all existing partitions.
4. Continue the installation of the FAStT RDAC.

To install the RDAC call setup and install the product, reboot the system.

Installing the FAStT Host Agent

The FAStT Host Agent allows in-band management of the FAStT Storage Server through Fibre Channel. It also allows other management workstations access the FAStT Storage Server via Fibre Channel and the TCP/IP connection of the system. To install the package, invoke the setup and follow the instructions.

The FAStT Host Agent can be disabled at a later time if it is not needed anymore by setting the startup type of the Windows services “FAStT Host Agent” to Manual.

Installing the FAStT Utilities

The FAStT Utilities package delivers two command tools. hot_add allows the addition of new LUNs while the operating system is up. SMdevices dumps a list of devices (controllers, LUNs) that are attached to the server. The package is installed by invoking setup. Reboot the system after the installation.

Enabling multipath I/O

Since you already installed the RDAC package, there is no need to configure anything for path redundancy.

4.1.2 Linux (RedHat 7.1 and 7.2, SuSE 7.3, and TurboLinux 7.0)

The Linux host software consists of three packages:

- FAStT Runtime Environment
- FAStT Client
- FAStT Utilities

There is no FAStT Agent for Linux, since in-band management is not available for Linux hosts.

> **Restriction:** The access logical drive must not be mapped to a Linux host because this interferes with the multipath I/O solution.

The FAStT Management Suite Java, FAStT MSJ, is part of the host software. It provides redundancy in the Fibre Channel paths. After the installation, you need to define preferred and alternate paths for each logical drive. The AVT feature of the FAStT Storage Server moves the logical drive from the preferred to the alternate controller if a failure occurs along the I/O path.

The drivers are delivered in a separate package, which is available for download at:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

These driver are for the FAStT Host Adapter and the FAStT FC-2 Host Adapter only. They do not support the Fibre Channel Host Adapter anymore. The drivers only work with certain combinations of distribution and kernel version. Therefore, it may be necessary to upgrade your system to a specific kernel version, according to the readme file contained in the driver package. Updated kernel packages are available for download for the different distributions:

- RedHat:

  http://www.redhat.com
  ftp://updates.redhat.com/

  Because this directory only contains the very newest version, it may be necessary to also look under:

  ftp://ftp.redhat.com/pub/redhat/support/enterprise/isv/kernel-archive/

- SuSE:

  http://www.suse.com
  ftp://ftp.suse.com/pub/suse/i386/update/

- TurboLinux:

  http://www.turbolinux.com
  ftp://ftp.turbolinux.com/pub/TurboLinux/turbolinux-updates/

You need to log on as `root` since you are installing new software, including new kernel modules, and changing the bootup procedure.

> **Important:** If an older version of the FAStT host software is already installed, you must remove it before you install a new one.

To remove an older version of the FAStT host software, use the command:

```
rpm -e <package-name>
```

Here *package name* is the name of the installed packages of the host software. By issuing the following command, you see a list of all installed packages of the FAStT host software:

```
rpm -ql | grep SM
```

Installing the FAStT MSJ

The FAStT MSJ is the software used to configure the host adapter driver to provide redundant paths to the storage subsystem. To install the software, you need to be running X-Windows. Then change to the directory where the installer is located and run:

```
./FAStTMSJ_install.bin
```

Be sure that the file has executable rights. Otherwise use the following command to allow execution of the file:

```
chmod o+x ./FAStTMSJ_install.bin
```

The FAStT MSJ is installed with the components you choose. You need at least the client part to manage this system from another one. You may install all components to do stand-alone management.

Installing host adapter driver

The drivers can now all be installed after meeting all requirements stated in the readme file of the driver package regarding kernel versions, host adapter settings, etc. Follow these steps:

1. Install the package by issuing the following command:

   ```
   rpm -i qla2x00.i386.rpm
   ```

   This installs the kernel modules.

2. Configure the startup procedures of your system to load the module by default. Add the following lines to the file /etc/modules.conf:

   ```
   alias scsi_hostadapter2 qla2200
   alias scsi_hostadapter3 qla2300
   options scsi_mod max_luns=32
   ```

3. Invoke the following command to generate the new module dependencies:

   ```
   depmod -a
   ```

   Again, since these settings may change with a newer driver version, check the readme file.

4. Rebuild the initial ram disk to include the module at boot time. Because the exact names vary from distribution to distribution, we use generic names here:

   ```
   /sbin/mkinitrd /boot/newinitrd-image 2.4.9
   ```

   The name 2.4.9 represents the exact kernal version that the system is running. This is the subdirectory name within the /lib/modules directory from which the kernal modules were taken. While writing this book, we worked on a RedHat 7.2 distribution with kernel 2.4.9-21smp.

5. When generating a new initial ram disk in the step above, modify the settings of the boot loader, lilo or grub, depending on the distribution. In our example, we explain the necessary steps for lilo. Again we use generic names here. Modify the following line in the file /etc/lilo.conf:

   ```
   initrd=/boot/old-initrd
   ```

   to

   ```
   initrd=/boot/newinitrd-image
   ```

6. Apply the new configuration of lilo by running:

   ```
   /sbin/lilo
   ```

7. Reboot the system and the driver loads automatically with the new settings.

Installing the host software

Now you are ready to install the FAStT Storage Manger software. You must maintain the following installation order:

1. FAStT Runtime Environment
2. FAStT Client
3. FAStT Utilities

All packages are installed under the /opt directory, added to the path, and can be invoked just after installing the software.

Installing the FAStT Runtime Environment

Since the entire FAStT software package is written in Java, a runtime environment is provided. To install the package change to the directory where the package resides, enter the following command.

```
rpm -ivh SMruntime.rpm
```

Now the installation starts.

Installing the FAStT Client

Now the client package can be installed. Change to the directory where the rpm-package of the software is located. Then issue the following command:

```
rpm -ivh SM7client.rpm
```

The client requires X-Windows to be running. Start it by typing the following command:

```
SMclient
```

The Event Monitor is always installed. The daemon starts automatically after the next reboot. If you don't want to use the Event Monitor on this system, you can disable the automatic startup by changing its settings in /etc/rc.d/init.d.

The client package also installs the command line utility **SMcli**, which is used in the FlashCopy scenarios in Chapter 5, "FlashCopy and the FAStT Storage Server" on page 139.

Installing the FAStT Utilities

The last package to install is the utility package. To install it, enter:

```
rpm -ivh SMutils.rpm
```

Enabling multipath I/O

After you define the logical drives and the storage partitioning, you need to configure the FAStT Management Suite Java if you want to ensure redundancy in the I/O paths. For each logical drive, you need to declare a preferred and an alternate path.

If the access logical drive is mapped to the Linux host, you cannot use the FAStT MSJ. It stops with an error message as shown in Figure 4-2.

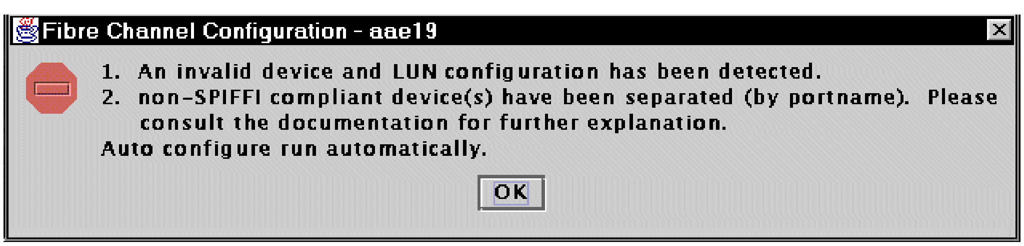

Figure 4-2 FAStT MSJ error message for the access logical drive

See 4.4, “Configuring storage partitioning” on page 105, to learn how to change the mapping of the access logical drive.

The following procedure is necessary whenever the logical drive configuration changes. It is best to finish configuring the FAStT Storage Server completely before continuing with the FAStT MSJ.

1. Before you can use the software, you need to start the client part of it by entering the following command in a command window as the root:

   ```
   qlremote
   ```

2. Ensure that no I/Os to FAStT are occurring when **qlremote** is running.
3. Start the FAStT MSJ and connect to your Linux system you want to administer, refer to 4.7, “The FAStT Management Suite Java” on page 133, for the general usage of the FAStT MSJ.
4. Highlight one of the adapters under the **HBA** tab, and click **Configure**, as shown in Figure 4-3.

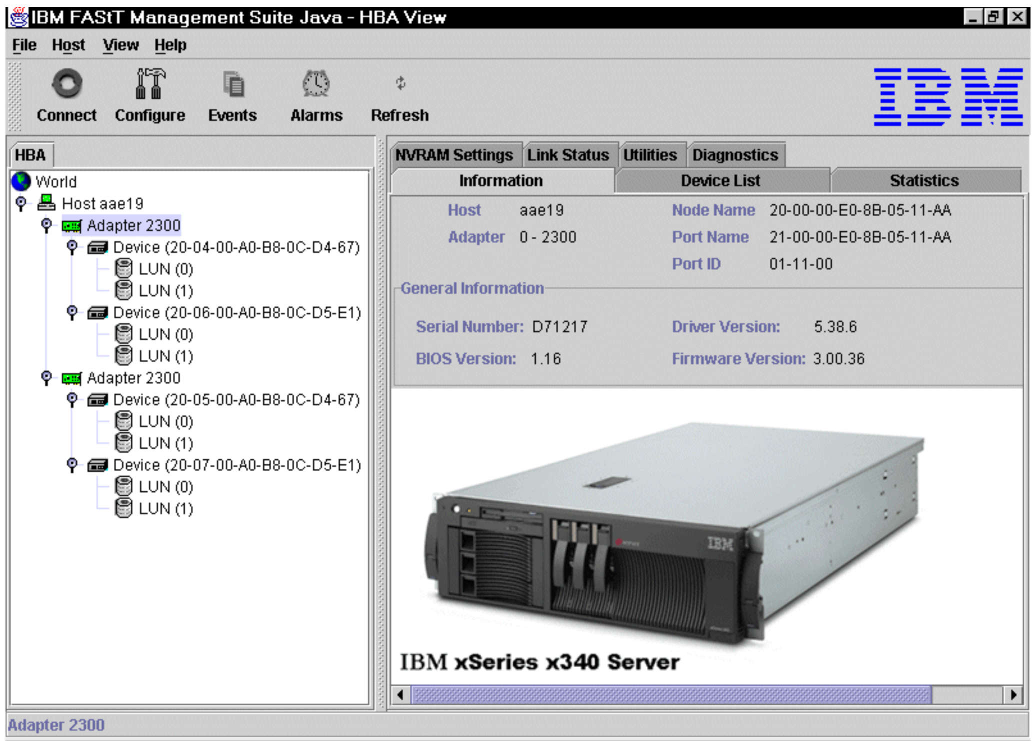

Figure 4-3 FAStT MSJ connected to a Linux Host

5. If this is the first time the FAStT MSJ is used on the host, or the configuration was deleted or changed since the last use, you see an error message (Figure 4-4) stating an invalid configuration.

Figure 4-4 FAStT MSJ invalid configuration detected

6. The configuration of the host adapters is corrected automatically, so that the first adapter is set to *visible* and the second to *hidden* for the operating system for all attached subsystems. It is also possible to configure the LUNs automatically (Figure 4-5). This leads to a configuration where all LUNs are enabled and the first adapter is used as the preferred path.

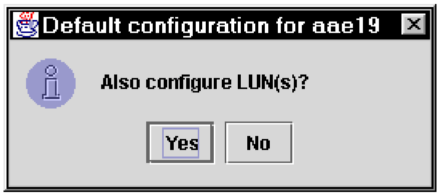

Figure 4-5 FAStT MSJ Automatic LUN configuration

Because the automatic procedure leads to a configuration without load balancing, you may want to click **No** here and configure the LUNs yourself.

7. The port configuration windows gives you a view of the Fibre Channel subsystems that are seen through each host adapter. The left column shows the two FAStT700 attached to the system with node names. Each FAStT houses two individual controllers, each with its own Fibre Channel port name, which is listed in the second column. The last two columns represent the host adapters in the system. Since Linux is not multipath aware, one adapter is always visible to the operating system, noted by a small eye in the particular entry, and the other one is hidden (Figure 4-6). Configure both adapters: one as visible and the other as hidden.

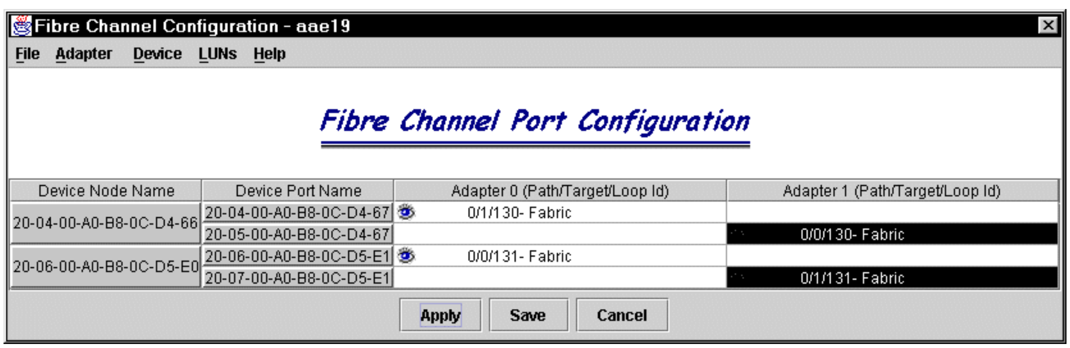

Figure 4-6 FAStT MSJ Fibre Channel port configuration

If one of the adapters is unconfigured, highlight the appropriate entry in the table and set the adapter to either hidden or visible. You receive an error message if any adapter in your system is unconfigured. At the time this redbook was written, a maximum of two host adapters is supported for Linux.

Even if one adapter is hidden to the operating system, it doesn't mean that this adapter is only used as a failover adapter. Both adapters can be used at the same time to load balance the logical drives throughout the adapters.

If you chose to automatically configure your LUNs (Figure 4-5), all LUNs use the first adapter as the preferred path; the other adapter is failover only.

> **Attention:** If you are using load balancing, be sure that the preferred path defined for a particular LUN in the host system matches the preferred controller ownership of the logical drive on the FAStT Storage Server.

If the preferred path in the host system does not match the preferred controller defined in FAStT, you see the error message *Logical Drives not on preferred Path* on the FAStT Storage Server. Refer to 8.4.3, “Logical drives not on the preferred path” on page 265, for more information. Redistributing the logical drives to their preferred controllers on the FAStT Storage Server does not solve the problem since the host will access again through its preferred path. This forces FAStT to move the LUN to the alternate controller again. To recover from that situation, you either can change the preferred controller on FAStT or you change the preferred path in the operating system.

8. To enable load balancing, highlight the storage subsystem represented by the node name and click **LUNs-> Configure LUNs** from the menu. Then the LUN configuration window appears as shown in Figure 4-7.

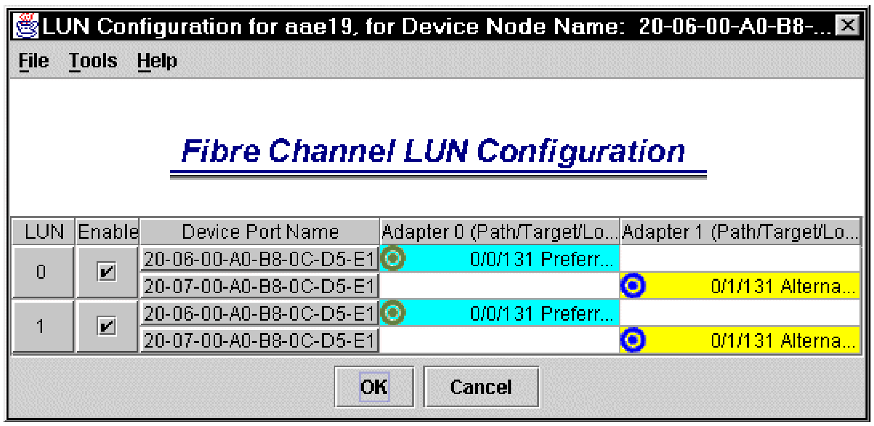

Figure 4-7 FAStT MSJ LUN configuration

You see a color coded representation of the logical drives available from the selected storage subsystem. You can enable or disable each LUN separately with the check box. Each LUN you want to use on the host system must be enabled.

The load balancing can be configured automatically be choosing **Tools-> Load Balance**. You can balance only the enabled LUNs or all LUNs. This mechanism alters the preferred path from LUN to LUN.

If you want to adapt the paths to your needs, right-click the entry you want to change and reverse the role of the paths for that LUN (Figure 4-8).

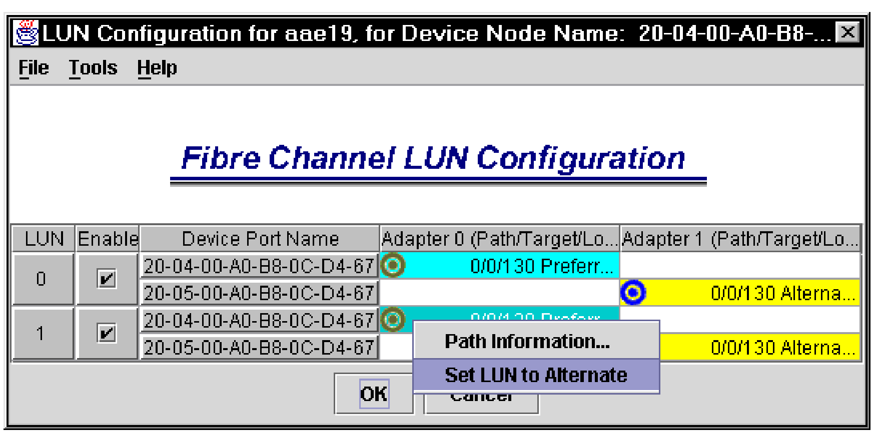

Figure 4-8 FAStT MSJ load balancing

9. If you finished configuring your logical drives, close the LUN configuration window and save the configuration to the host system in the Port Configuration window (Figure 4-6). A confirmation appears if the configuration saved successfully.

   We had problems saving the configuration when we used a different kernel version than the stated one. If the configuration is not saved correctly, your logical drives are not protected by path redundancy.

10. After you save the new configuration, you can exit the FAStT MSJ and stop the qlremote daemon by pressing Ctrl+C in the command line.

> **Important:** The new configuration cannot be applied dynamically. The adapter driver must be reloaded to work with the new configuration.
>
> Also the initial ram disk must be recreated to reflect the configuration changes after the next reboot.

11. To make the new configuration actually work, reload the driver modules qla2200 or qla2300. You cannot do this if a disk is mounted on the Fibre Channel subsystem attached to the adapter that is to be reloaded.

    If you load the host adapter driver through an initial ram disk, you need to recreate this ram disk because the configuration data for redundant disks is applied when loading the adapter module. This procedure is explained in “Installing the host adapter driver” on page 67.

Whenever the configuration of the logical drives changes, you must follow this procedure to save a valid configuration to the system.

4.1.3 Novell NetWare 5.1 and 6.0

The host software consists of three packages

- FAStT Client
- FAStT Host Agent
- FAStT Utilities

For NetWare 5.1, you also have to install a Java Virtual Machine. The drivers are available as a separate package for download from:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

The installation order for NetWare is to install the host bus adapter first, followed by the FAStT software and finishing with the multipath I/O driver.

Installing the host adapter driver

To install the driver load the appropriate NetWare program by invoking LOAD NWCONFIG, follow these steps:

1. Click **Driver-> Configure disk and storage device drivers** from the menu.
2. Choose **Select an additional driver** and point the installation for an unlisted driver to the location where the *HAM* driver for the host bus adapter is located. The driver should be copied from that location to the server directory during installation.

Installing the host software

Maintain the following order while installing the different packages of the host software on a NetWare server:

1. Java Virtual Machine (NetWare 5.1 only)
2. FAStT Client for NetWare
3. FAStT Agent for NetWare
4. FAStT Utilities for NetWare

Installing the Java Virtual Machine

For NetWare 5.1, install this package first, because the FAStT Storage Manager is written in Java. Refer to the readme file of the FAStT Storage Manager for which the Novell Support Pack must be installed. To install the JVM, refer to the Novell Technical Information Document (TID) 2961564, which you can find on the Novell Web site at:

`http://www.novell.com/support/`

Installing the FAStT Client

To install the FAStT Client, follow these steps:

1. Open the Novell Install Products dialog.
2. Click **Add** to install the host software.
3. Point the installer to the directory containing the FAStT Client package and select the file **product.ni**.
4. Accept the license agreement during installation.

The client package also installs the command line utility **`SMcli`**, which we use in the FlashCopy scenarios in Chapter 5, “FlashCopy and the FAStT Storage Server” on page 139.

Installing the FAStT Host Agent

The FAStT Host Agent allows in-band management of the FAStT Storage Server through the Fibre Channel. It also allows other management workstations access the FAStT Storage Server via Fibre Channel and the TCP/IP connection of the system.

To install the FAStT Host Agent, complete these steps:

1. Open the Novell Install Products dialog.
2. Click **Add** to install the host software.
3. Point the installer to the directory containing the FAStT Agent package and select the file **product.ni**.
4. Accept the license agreement during installation.

Installing the FAStT Utilities

The FAStT Utilities package delivers two command tools. hot_add allows the addition of new LUNs while the operating system is up, SMdevices dumps a list of devices (controllers, LUNs) that are attached to the server.

To install the FAStT Utilities, complete these steps:

1. Open the Novell Install Products dialog.
2. Click **Add** to install the host software.
3. Point the installer to the directory containing the FAStT Utility package and select the file **product.ni**.
4. Accept the license agreement during installation.

Enabling multipath I/O

IBM provides a special driver called IBMSAN.CDM. This driver provides multipath I/O for a NetWare server attached to a FAStT Storage Server. You can download the driver from the Web at:

http://www.pc.ibm.com/support

Check the included readme file for any requirements such as certain driver and BIOS levels.

> **Restriction:** If you are using the IBMSAN driver, PCI hot-plug functionality for the host bus adapters is not supported.
>
> In NetWare 6 with Support Pack 1 installed, you must disable the multipath I/O driver Novell provides. To disable the driver, the following line must be the first line in the startup.ncf file of the system:
>
> `MM Set Multi-path Support = Off`

To install the driver, copy the file ibmsan.cdm file to the C:\NWSERVER directory on the server. Edit the file C:\NWSERVER\STARTUP.NCF and locate the following line:

```
LOAD SCSIHD.CDM
```

After the previous line, add the following line to load the driver during startup:

```
LOAD IBMSAN.CDM
```

The IBMSAN driver automatically configures itself to provide redundant paths to the attached FAStT Storage Server.

The driver needs some additional parameters for the host bus adapter driver, so add the following four options to the parameters for the QL2x00.HAM driver:

- **/LUNS**
- **/ALLPATHS**: Without these two options, the alternate paths are hidden.
- **/PORTNAMES**: This setting allows a proper utilization of the access logical drive for in-band management of the FAStT Storage Server.
- **/XRETRY=400**: This setting extends the retry count for any I/O command experiencing failures.

An example C:\NWSERVER\STARTUP.NCF file is shown in Example 4-1.

Example 4-1 Sample STARTUP.NCF

```
LOAD MPS14.PSM
LOAD IDECD.CDM
LOAD SCSIHD.CDM
LOAD IBMSAN.CDM
LOAD IDEATA.HAM PORT=1FO INT=E
LOAD IPSRAID.HAM SLOT=5
LOAD AIC78U2.HAM SLOT=10005
LOAD AIC78U2.HAM SLOT=10006
LOAD QL2300.HAM SLOT=2 /LUNS /ALLPATHS /PORTNAMES /XRETRY=400
LOAD QL2300.HAM SLOT=1 /LUNS /ALLPATHS /PORTNAMES /XRETRY=400
```

Restart the system for the changes to become active.

4.1.4 AIX 4.3.3 and 5.1

Since AIX is only supported with out-band management, the host software consists of only one package – the FAStT Client. The disk array driver, which provides redundancy in the I/O paths, is available as a PTF as part of the AIX operating system.

Because the access logical drive is not needed for an AIX host, ensure that no mapping exists for the access logical drive to an AIX host. Refer to 4.4, "Configuring storage partitioning" on page 105, to learn how to define mappings to a logical drive.

Refer to the readme file to see which PTF levels are required for the host bus adapter driver and the disk array driver. There may be other PTFs that must be installed first. You can download all the information and PTFs from:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

Since you are installing new software and changing the boot up procedure, you need to log on as `root`.

Installing a host adapter driver

The driver for the host adapter is part of AIX and is available as a PTF. Be sure that you use the correct version of the PTF when you install it.

Installing the host software

Only the client package is available for AIX because currently no in-band management support is available.

Installing the FAStT Client

To install the package, issue the following command (the exact name of the package differs between AIX 4.3.3 and AIX 5.1, so we use generic names here):

```
#installp -a -d <complete-path>/SMclient-aix.bff SM8client.aix.rte
```

The installation process begins. The process displays information as it runs, including an installation summary when the process is finished.

To verify the installation, use the command:

```
lslpp -ah SMclient.aix.rte
```

The verification process returns a table that describes the software installation, including the installation package file name, version number, action, and action status.

The clients only run under X-Windows and can be started with the SMclient command.

Enabling multipath I/O

Since you already installed the RDAC package, there is no need to configure anything for path redundancy.

4.1.5 HP-UX v11

The host software for HP-UX consists of four packages:

- FAStT Runtime Environment
- FAStT Client
- FAStT Host Agent
- FAStT Utilities

Check the readme file for any host specific requirements for specific patches or adapter settings.

The driver for the host bus adapter is included in the latest revision of the HP-UX. If two host bus adapters are installed in the system, redundancy is achieved with the PV-Links package, which is part of the operating system.

Because you are installing new software and changing the boot up procedure, you need to log on as `root`.

Installing host adapter driver

Since the driver for the host bus adapter is included in the HP-UX kernel, there is no need to install a separate package.

After you install the adapters, the system recognizes the adapters automatically.

Installing the host software

Now you are ready to install the FAStT Storage Manger software. You must maintain the following order of installation:

1. FAStT Runtime Environment for HP-UX
2. FAStT Client for HP-UX
3. FAStT Host Agent for HP-UX
4. FAStT Utilities for HP-UX

All packages are added to the path and can be invoked after you install the software.

Installing FAStT Runtime Environment

To install the runtime package, follow these steps:

1. Change to the directory where the software package located. Then issue the following command:

   ```
   swinstall -s <complete path>/SMruntime-HP.pkg
   ```

2. Choose to install all packages. If you did not copy the software package in the HP-UX software depot, you must enter the full path where the package is located.
3. Verify the software installation:

   ```
   swverify -v SMruntime
   ```

4. If there are any error messages, check the file /var/adm/sw/swagent.log and follow the instructions outlined there.

Installing FAStT Client

Next you need to install the client package. Follow these steps:

1. Change to the directory where the software package located. Then issue the following command:

   ```
   swinstall -s <complete path>/SMclient-HP.pkg
   ```

2. Choose to install all packages. If you did not copy the software package in the HP-UX software depot, you must enter the full path where the package is located.
3. Verify the software installation:

   ```
   swverify -v SMclient
   ```

4. If there are any error messages, check the file /var/adm/sw/swagent.log and follow the instructions outlined there.

Installing the FAStT Host Agent

Then you need to install the FAStT Host Agent. Follow these steps:

1. Enter the following command:

```
swinstall -s <complete path>/SMagent-HP.pkg
```

2. The agent installs a daemon that starts directly after the installation. The daemon is also included in the startup scripts and starts with every system reboot. The Event Monitor is also installed with the functionality matching the system configuration. If you have SNMP and SMTP installed, it is installed with both options. If you only have one, it only installs support for this.
3. Verify the software installation:

```
swverify -v SMagent
```

4. If there are any error messages, check the file /var/adm/sw/swagent.log and follow the instructions outlined there.

Installing the FAStT Utility

The last package to install is the utility package. Follow these steps:

1. Enter the following command:

```
swinstall -s <complete-path>/SMutil-HP.pkg
```

2. Verify the software installation:

```
swverify -v SMutil
```

3. If there are any error messages, check the file /var/adm/sw/swagent.log and follow the instructions outlined there.

Enabling multipath I/O

If the HP-UX system is attached with two host bus adapters to the FAStT Storage Server, you can take advantage of this layout to provide redundant access to the storage. You can do this using PV-Links, which is part of the HP-UX operating system. The redundancy is achieved by using volumes with a primary and secondary path to the same device.

You must log on as `root`, because new devices are added in this procedure. Then follow these steps:

1. Determine the primary and alternate path for each logical drive. Make sure, that all logical drives are on the preferred path in the subsystem management window.
2. Execute the `SMdevices` command. A dump similar to the one shown in Example 4-2 is listed. Notice that every logical drive, HP_0 to HP_3, and the access volume, are listed twice because you have two paths to each logical drive.

Example 4-2 Sample output of SMdevices

```
# SMdevices
IBM FAStT Storage Manager Devices, Version 08.20.45.00
Built Wed Mar 20 00:58:59 GMT+00:00 2002
(C) Copyright International Business Machines Corporation, 2002 Licensed Material
Program Property of IBM.  All rights reserved.

/dev/rdsk/c26t0d0 [Storage Subsystem ITSO, Logical Drive HP_0, LUN 0, Logical Drive WWN
<600a0b80000c2e86000000393cebc4ea>, Preferred Path (Controller-A): In Use]
/dev/rdsk/c26t0d1 [Storage Subsystem ITSO, Logical Drive HP_1, LUN 1, Logical Drive WWN
<600a0b80000c2e11000000393cebc25d>, Alternate Path (Controller-A): In Use]
```

```
/dev/rdsk/c26t0d2 [Storage Subsystem ITSO, Logical Drive HP_2, LUN 2, Logical Drive WWN
<600a0b80000c2e86000000353cebc471>, Preferred Path (Controller-A): In Use]
/dev/rdsk/c26t0d3 [Storage Subsystem ITSO, Logical Drive HP_3, LUN 3, Logical Drive WWN
<600a0b80000c2e110000003d3cebc4b0>, Alternate Path (Controller-A): In Use]
/dev/rdsk/c26t3d7 [Storage Subsystem ITSO, Logical Drive Access, LUN 31, Logical Drive WWN
<600a0b80000c2e8600000000100000000>]
/dev/rdsk/c34t0d0 [Storage Subsystem ITSO, Logical Drive HP_0, LUN 0, Logical Drive WWN
<600a0b80000c2e86000000393cebc4ea>, Alternate Path (Controller-B): In Use]
/dev/rdsk/c34t0d1 [Storage Subsystem ITSO, Logical Drive HP_1, LUN 1, Logical Drive WWN
<600a0b80000c2e11000000393cebc25d>, Preferred Path (Controller-B): In Use]
/dev/rdsk/c34t0d2 [Storage Subsystem ITSO, Logical Drive HP_2, LUN 2, Logical Drive WWN
<600a0b80000c2e86000000353cebc471>, Alternate Path (Controller-B): In Use]
/dev/rdsk/c34t0d3 [Storage Subsystem ITSO, Logical Drive HP_3, LUN 3, Logical Drive WWN
<600a0b80000c2e110000003d3cebc4b0>, Preferred Path (Controller-B): In Use]
/dev/rdsk/c34t3d7 [Storage Subsystem ITSO, Logical Drive Access, LUN 31, Logical Drive WWN
<600a0b80000c2e1100000000100000000>]
#
```

As you can see, the list contains the path, the name of the logical drive, and the information on the preferred or alternate path for each device seen on the FAStT Storage Server. The access logical drives don’t have a preferred or alternate path because they are each bound to their controller.

3. Extract the information needed to define the paths using PV-Links in Table 4-1.

Table 4-1 Preferred and alternate paths for HP-UX

Logical drive	Preferred path	Alternate path
HP_0	/dev/rdsk/c26t0d0	/dev/rdsk/c34t0d0
HP_1	/dev/rdsk/c34t0d1	/dev/rdsk/c26t0d1
HP_2	/dev/rdsk/c26t0d2	/dev/rdsk/c34t0d2
HP_3	/dev/rdsk/c34t0d3	/dev/rdsk/c26t0d3

4. With this information, you can define a volume group that contains the logical drive HP_0 and that uses the preferred path and the alternate path.

   We can define the primary path for the logical drive HP_0 by issuing:

   ```
   # pvcreate /dev/rdsk/c26t0d0
   Physical volume "/dev/rdsk/c26t0d0" has been successfully created.
   ```

5. The system confirms the creation of the new physical volume /dev/rdsk/c26t0d0.
6. Change to the directory /dev, create a new directory for the volume group and create a new device for this volume group:

   ```
   # cd /dev
   # mkdir vgibm1
   # cd vgibm1
   # mknod group c 64 0x010000
   ```

7. Verify the successful creation by using the **ls -l** command.
8. Create the volume group and extend it with the secondary path as shown in Example 4-3 using the **vgcreate** command.

Example 4-3 Creating the volume group

```
# vgcreate /dev/vgibm1 /dev/dsk/c26t0d0
Increade the number of physical extents per physical volume to 12988.
```

```
Volume group "/dev/vgibm1" has been successfully created.
Volume Group configuration for /dev/vgibm1 has been saved in /etc/lvmconf/vgibm1.conf
# vgextend /dev/vgibm1 /dev/dsk/c34t0d0
Volume group "/dev/vgibm1" has been successfully extended.
Volume Group configuration for /dev/vgibm1 has been saved in /etc/lvmconf/vgibm1.conf
```

9. The volume group is now successfully created with redundant paths by using both host bus adapters. You can verify the configuration of the volume group by using the **vgdisplay** command as shown in Example 4-4.

Example 4-4 Verifying the new volume group

```
# vgdisplay -v /dev/vgibm1

--- Volume groups ---
VG Name                     /dev/vgibm1
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      0
Open LV                     0
Max PV                      16
Cur PV                      1
Act PV                      1
Max PE per PV               12988
VGDA                        2
PE Size (Mbytes)            4
Total PE                    12986
Alloc PE                    0
Free PE                     12986
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

   --- Physical volumes ---
   PV Name                     /dev/dsk/c26t0d0
   PV Name                     /dev/dsk/c34t0d0Alternate Link
   PV Status                   available
   Total PE                    12986
   Free PE                     12986
   Autoswitch                  On
```

10. The volume group may now be expanded by adding the primary and secondary path to another logical drive. The volume group is now ready to be used.

4.1.6 Solaris (SPARC) 2.6, 2.7, and 2.8

The host software for Solaris consists of five packages:

- FAStT Runtime Environment
- FAStT Client
- FAStT RDAC
- FAStT Host Agent
- FAStT Utilities

The driver for the host bus adapter is delivered in a separate package. If two host bus adapters are installed in the system, redundancy is achieved with the FAStT RDAC package.

Because you are installing new software, including new kernel modules, and changing the boot up procedure, you need to log on as `root`.

Installing host adapter driver

Depending on the type of host bus adapter you are using, follow the instructions included with the adapter because several adapters from different manufacturers are available.

Installing the host software

Now you are ready to install the FAStT Storage Manger software. Be sure to maintain the following installation order:

1. FAStT Runtime Environment for Solaris
2. FAStT Client for Solaris
3. FAStT RDAC for Solaris
4. FAStT Agent for Solaris
5. FAStT Utilities for Solaris

All packages are added to the path and can be invoked just after installing the software.

Installing the FAStT Runtime Environment

The first package that needs to be installed is the runtime package. It provides the Java Virtual Machine for the FAStT host software. Change to the directory where the software package is located. Then issue the following command:

```
pkgadd -d ./SMruntime-SPARC.pkg
```

Installing the FAStT Client

The next step is to install the client. Follow these steps:

1. Change to the directory where the software package is located. Then issue the following command:

   ```
   pkgadd -d ./SMclient-SPARC.pkg
   ```

2. Choose the default to install all packages and accept the execution of scripts; they are part of the installation.

Installing the FAStT RDAC

Now you can install the RDAC. It provides redundant paths to the storage subsystem if you have two adapters in the system. This package requires you to reboot the system after installation. Follow these steps:

1. Run the following command:

   ```
   pkgadd -d ./SM7rdac-SPARC.pkg
   ```

2. To ensure a proper reconfiguration, reboot the system using the command:

   ```
   reboot - - r
   ```

Installing the FAStT Agent

After the RDAC is integrated in your system, you can install the agent. Enter the following command:

```
pkgadd -d ./SM7agent-SPARC.pkg
```

The agent installs a daemon that starts directly after the installation. The daemon is also included in the startup scripts and starts with every system reboot.

Installing the FAStT Utilities

The last package to install is the utility package. To install it, enter:

```
pkgadd -d ./SM7util-SPARC.pkg
```

Enabling multipath I/O

Since you already installed the RDAC package, there is no need to configure anything for path redundancy.

4.1.7 Network attached management

Only the FAStT Client and the runtime environment has to be installed on a management workstation. The management workstation needs to be connected to either the FAStT Storage Server or to other hosts running the FAStT Host Agent through Ethernet. The installation is the same as for the host running the whole FAStT software. See the appropriate installation instructions for the operating system you need.

4.2 Configuring the FAStT Storage Server

This section explains how to configure the FAStT Storage Server. The configuration requires:

- One host running Microsoft Windows NT 4.0 with two logical drives
- One host running Linux RedHat 7.2 with two logical drives
- A two-node Windows 2000 Advanced Server Cluster with three logical drives
- One host running Novell NetWare with two logical drives
- An access logical drive that is only available to the Windows hosts

You should also see Table 4-2.

The FAStT Storage Manager is divided into two main parts:

- The Enterprise Management window: Show all the available FAStT Storage Servers
- The Storage Subsystem Management window

In the Enterprise Management window, you choose the storage subsystem you want to manage and define global settings such as the alert settings for the Event Monitor service. We presume that you are already working in the Storage Subsystem Management window on the FAStT Storage Server, if not stated otherwise.

4.2.1 Planning the configuration

A configuration of a FAStT Storage Server can be complex, especially when different operating systems and storage partitioning are involved. Therefore, you should plan the configuration you want to apply in advance.

On one side, you need to define the arrays and the logical drives you need, including considerations such as number of drives in the arrays, size of the logical drives, RAID level and segment size. To plan the disk layout, you need to know the attached hosts, their operating system, and the application using the storage of the FAStT Storage Server.

On the other side, you also need to define the layout of the attached hosts with their host bus adapter and the mappings of the logical drives to specific host groups or hosts. You should prepare a mapping table, as shown in Table 4-2, where you keep all necessary information regarding the storage partitioning. We refer to our mapping table example throughout the entire configuration of the FAStT Storage Server.

Table 4-2 Storage partitioning

Host group	Host name	Port name	WWN	OS type	LUN	LUN#
Windows	SRV_NT40	NT40_A		Windows SP5 non-clustered	Groups	0
		NT40_B			Homes	1
					Access	31
Linux	SRV_RH72	RH72_A		Linux	Internet	0
		RH72_B			Intranet	1
Cluster	Node1	Node1_A	210000e08b050eaa	W2000 clustered	Quorum	0
		Node1_B	210000e08b0510aa		SQL-Log	1
	Node2	Node2_A	210000e08b0511aa		SQL-Data	2
		Node2_B	210000e08b050faa		Access	31
Solaris	Fire	Fire_A		Solaris	Data	
		Fire_B			Test	

The name of a logical drive allows you later to differentiate easily between all the logical drives defined on FAStT. We suggest you use a combination of the server name and its usage. In our case, we couldn't use this name combination because it wouldn't fit into our mapping table.

4.2.2 Starting the FAStT Storage Manager Client

When you start the FAStT Storage Manager Client, it launches the Enterprise Management window. The first time you start the client you are prompted to select whether you want an initial discovery of available storage subsystems (see Figure 4-9).

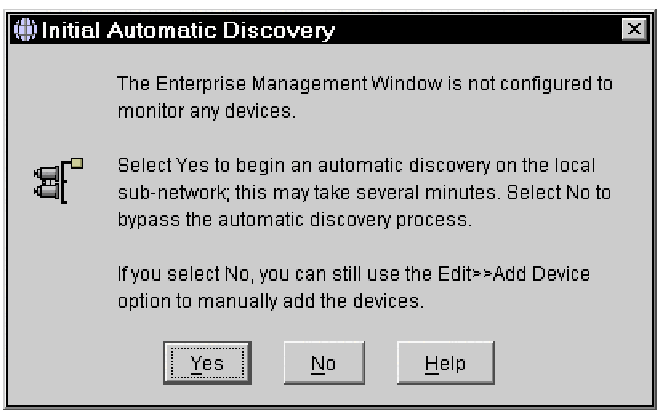

Figure 4-9 Initial Automatic Discovery

The client software sends out broadcasts via Fibre Channel and the IP network, if it finds directly attached storage subsystems or other hosts running the FAStT Storage Manager Host Agent with an attached storage subsystem.

You have to invoke the Automatic Discovery every time you add a new FAStT Storage Server in your network or install new host agents on already attached systems. To have them detected in your Enterprise Management window, click **Edit-> Rescan**. Afterward all FAStT Storage Server are listed in the Enterprise Management window as shown in Figure 4-10.

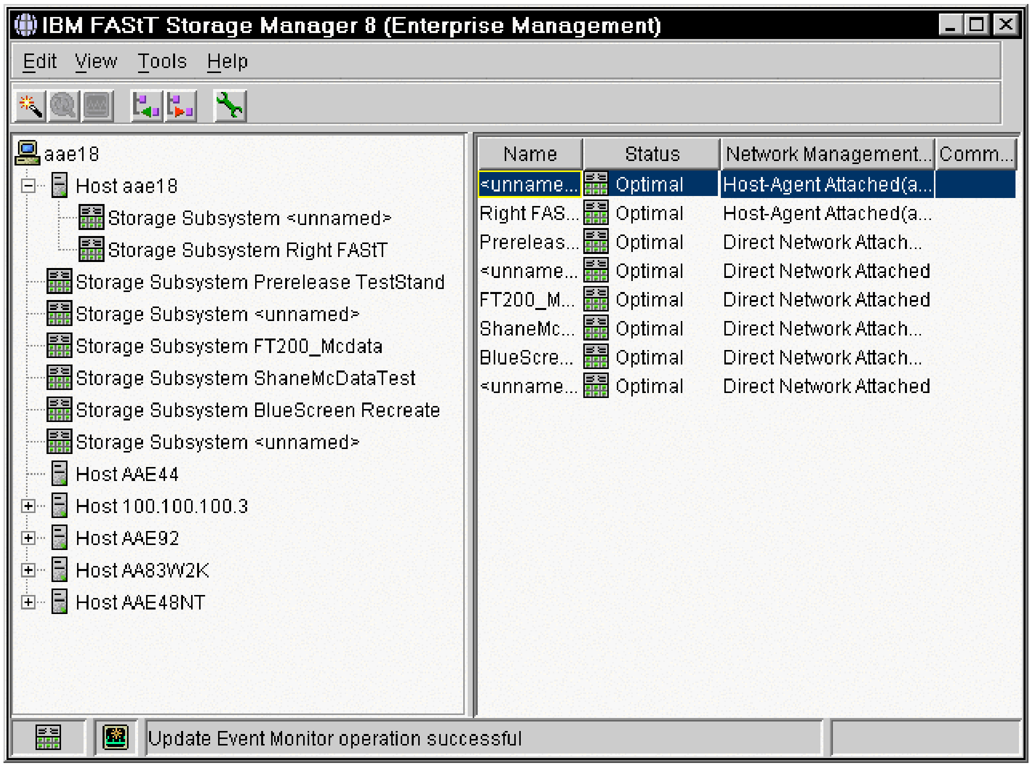

Figure 4-10 The Enterprise Management window

The FAStT Storage Server may be connected via Ethernet. Or you may want to manage through the host agent of another host, which is not in the same broadcast segment as your management station. In either case, you have to add the devices manually. Click **Edit-> Add device** to enter the host name or the IP address you want to attach (see Figure 4-11). If you add a FAStT Storage Server that is directly managed, be sure to enter both IP addresses – one per controller.

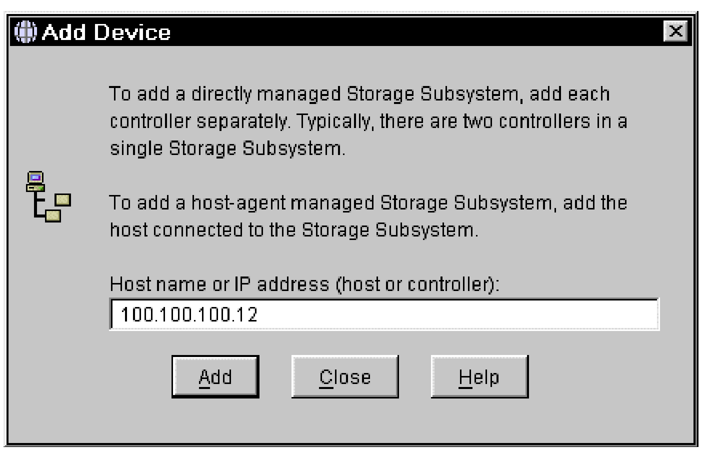

Figure 4-11 Adding a device manually

An overall status is already shown in the Enterprise Management window. You can see all FAStT Storage Servers and how they are managed – either direct or host-agent attached or through both connections. There is a status column. Usually the status is *Optimal* with a green icon next to it. But if there are any problems, the status changes to *Needs Attention* and a red icon appears.

If the FAStT Storage Server is not found by the discovery process and you cannot attach to the device manually, see 8.4, “Common problems” on page 263.

By choosing the storage subsystem you want to manage, right-click and select **Manage Device** for the attached storage subsystem. This launches the Subsystem Management window (Figure 4-12).

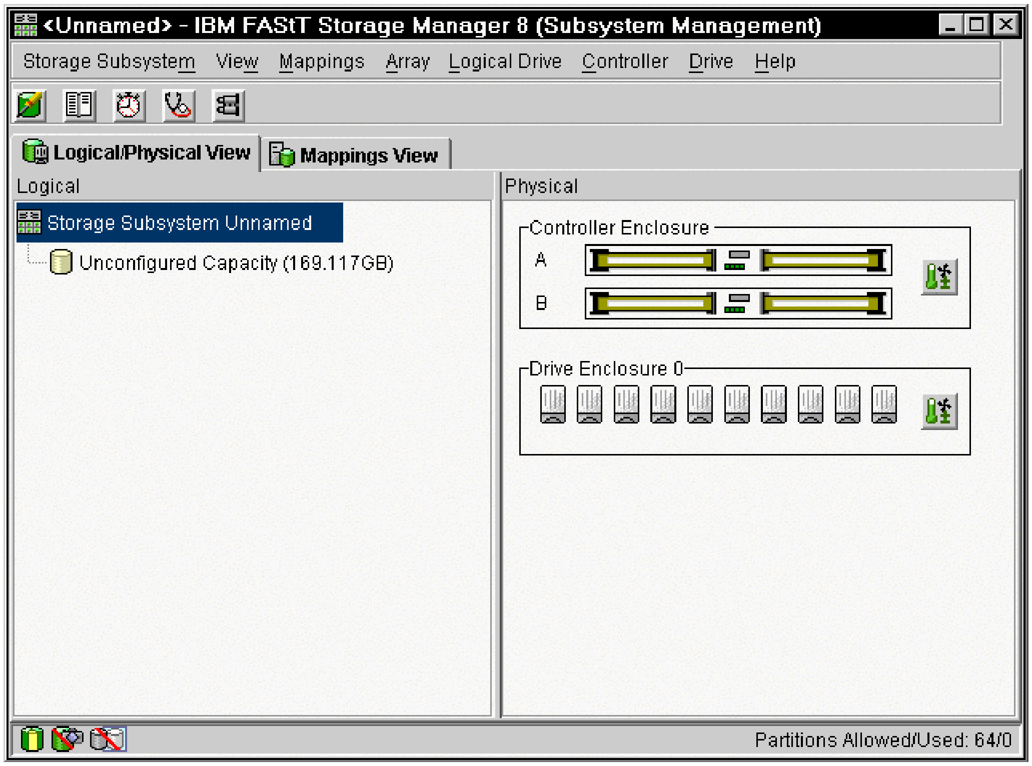

Figure 4-12 First launch of the Subsystem Management window

Verify that the enclosures in the right half of the window reflect your actual physical layout. If the enclosures are listed in an incorrect order, select **Storage Subsystem-> Change-> Enclosure Order** and sort the enclosures according to your site setup (Figure 4-13).

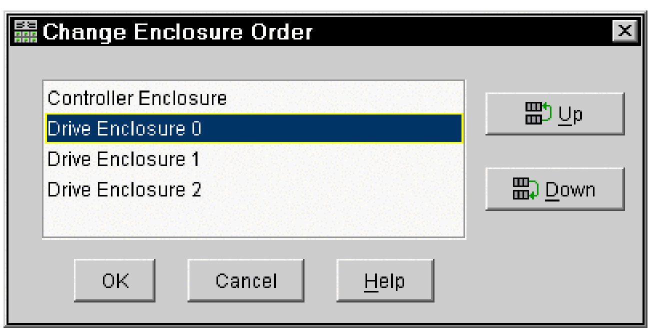

Figure 4-13 Adapting the enclosure order

4.2.3 Updating the controller microcode

The microcode of the FAStT Storage Server consists of two packages:

- The actual firmware
- The NVSRAM package, including the settings for booting the FAStT Storage Server

The NVSRAM is similar to the settings in the BIOS of a host system. The firmware and the NVSRAM are *not* independent. Be sure to install the correct combination of the two packages.

1. The upgrade procedure needs two independent connections to the FAStT Storage Server, one for each controller. It is not possible to perform a microcode update with only one controller connected. Therefore, both controllers must be accessible either via Fibre Channel or Ethernet. Both controllers must also be in the active state.

   If you plan to upgrade via Fibre Channel, make sure that you have a multipath I/O driver installed on your management host, for example the FAStT RDAC package. This is necessary since access logical drive moves from one controller to the other during this procedure and the FAStT Storage Server must be manageable during the entire time.

   Important:

   - Ensure that *all* hosts attached to FAStT have a multipath I/O driver installed.
   - Any power or network/SAN interruption during the update process may lead to configuration corruption. Therefore, do not turn off the power to the FAStT Storage Server or the management station during the update. If you are using in-band management and have Fibre Channel hubs or managed hubs, then make sure any SAN connected devices are not powered up during the update. Otherwise, this can cause a loop initialization process and interrupt the process.

2. Open the Subsystem Management window for the FAStT Storage Server you want to upgrade. To download the firmware, follow these steps:

   a. Highlight the storage subsystem. From the Storage Subsystem menu, click **Download-> Firmware** (Figure 4-14).

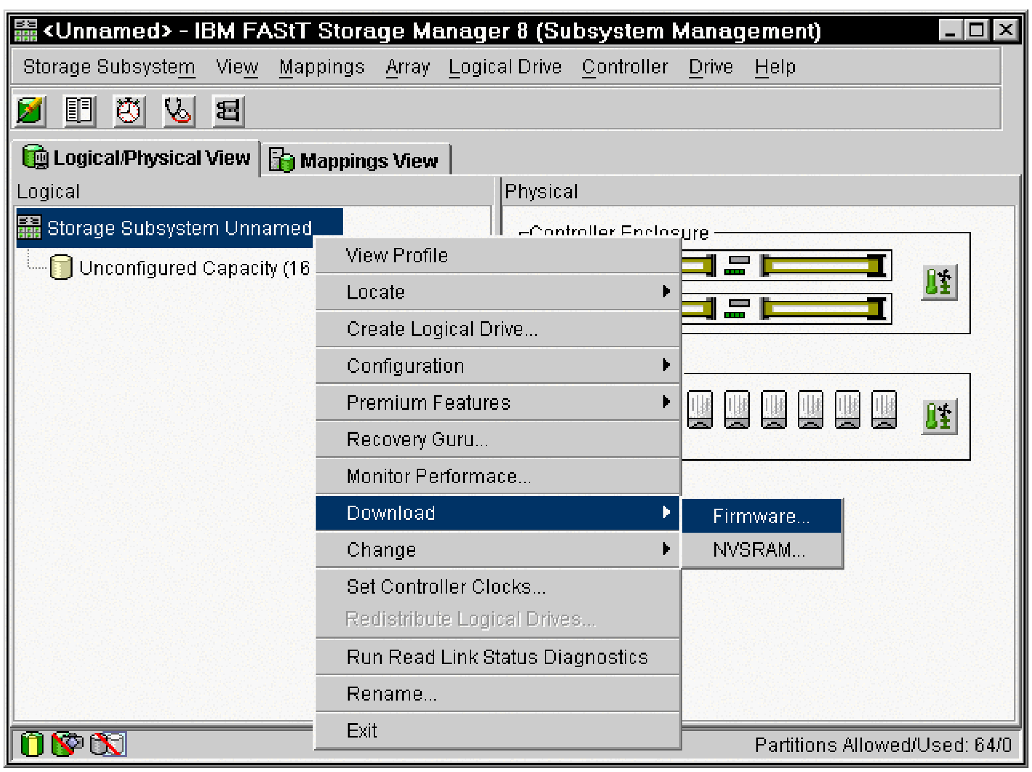

Figure 4-14 Subsystem Management window: Selecting Download-> Firmware

b. You are prompted for the file location where the new firmware you want to download is stored. In the first part of the window, you can see the actual versions of the firmware and NVSRAM installed on the controllers. The second part tells you which firmware revision is included in the file you have chosen as shown in Figure 4-15.

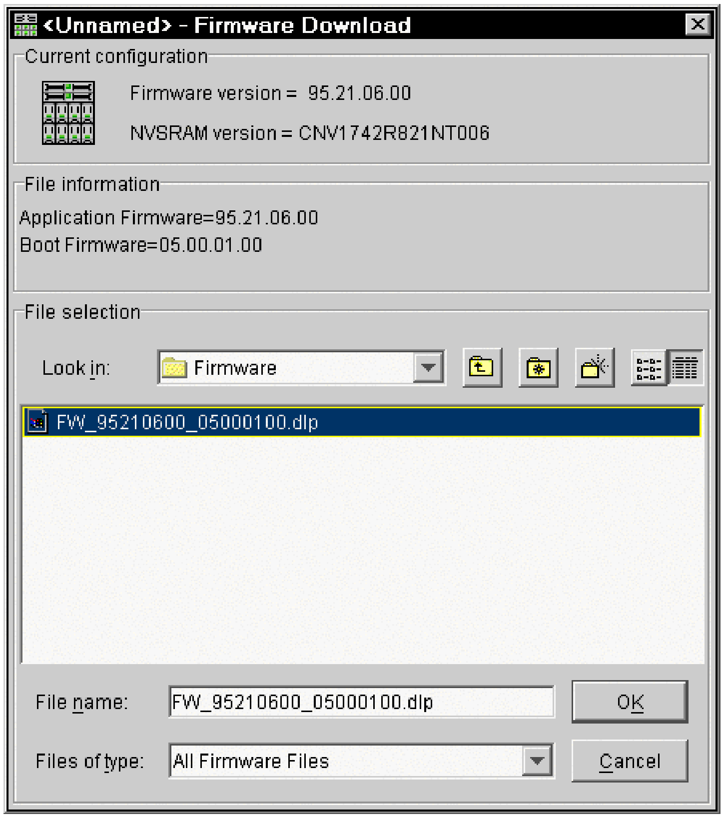

Figure 4-15 Firmware Download window

Click **OK**.

c. A new window (Figure 4-16) appears where you have to confirm that all the requirements are met. Click **Yes** to confirm the download.

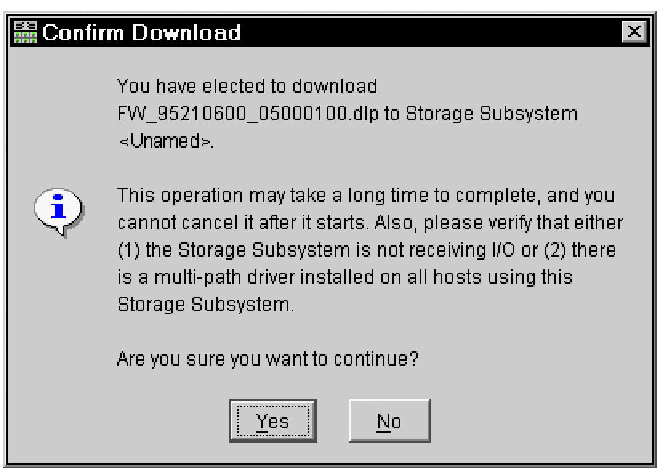

Figure 4-16 Confirming the download of the new firmware

3. The firmware upgrade continues. You can follow the process as shown in the Downloading window (Figure 4-17). The entire download process can take several minutes to finish.

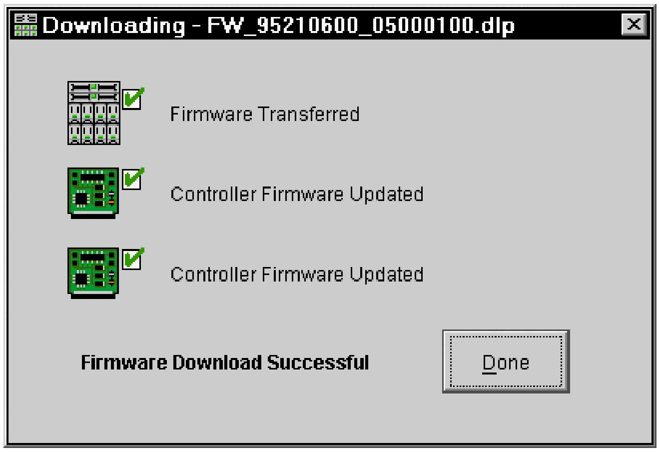

Figure 4-17 Download of firmware completed

4. After you upgrade the firmware, you must also upgrade NVSRAM:
   a. Highlight the storage subsystem again and click **Storage Subsystem-> Download-> NVSRAM** as shown in Figure 4-18.

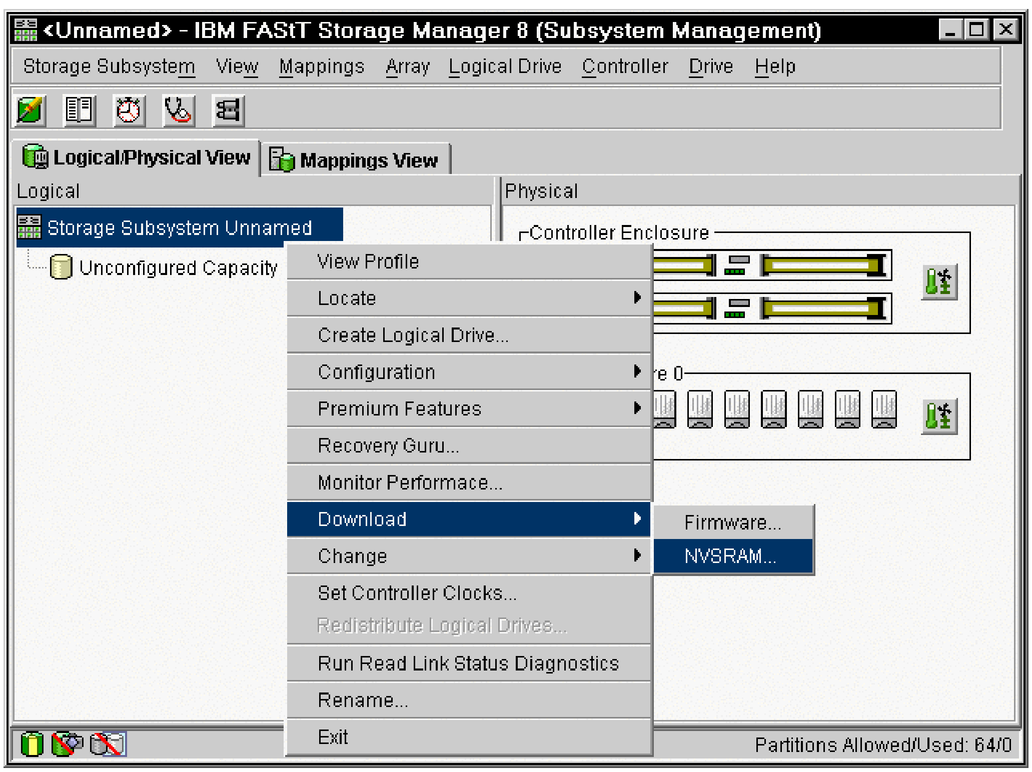

Figure 4-18 Subsystem Management window: Download-> NVSRAM

   b. As for the firmware download, choose the correct file you want to download to the controllers (Figure 4-19). The window has the same structure as the Firmware Download dialog (Figure 4-15). Therefore, you see the new firmware revision and the NVSRAM revision installed on the storage subsystem.

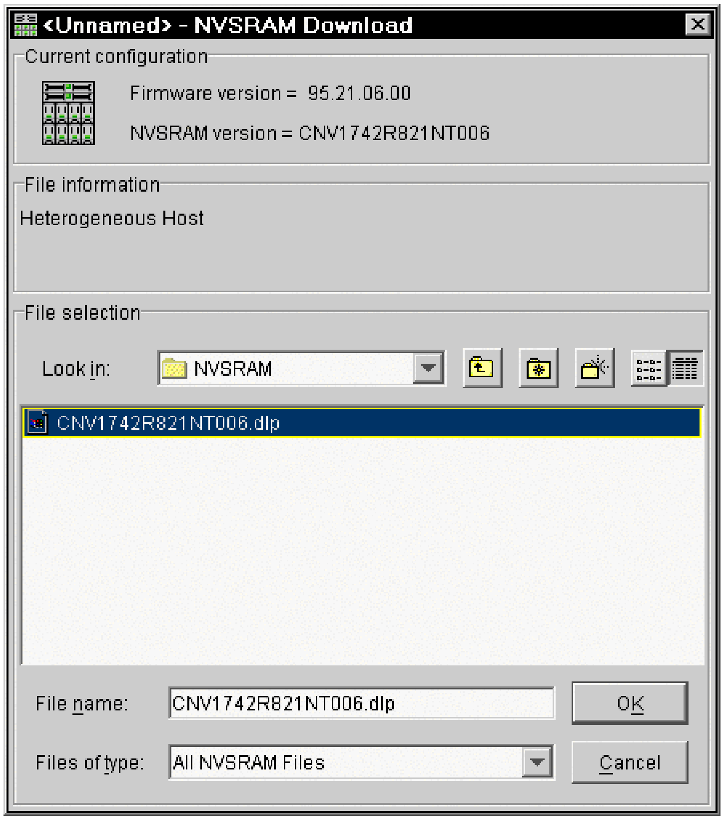

Figure 4-19 NVSRAM Download window

c. The confirmation window opens again to ensure that you are downloading a version compatible to the firmware you just downloaded. Click **Yes** to accept the requirements and continue (see Figure 4-20).

d. If you applied any changes to the NVSRAM settings, for example running a script, you must re-apply them after the download of the new NVSRAM completes. The NVSRAM update resets all settings stored in the NVSRAM to their defaults.

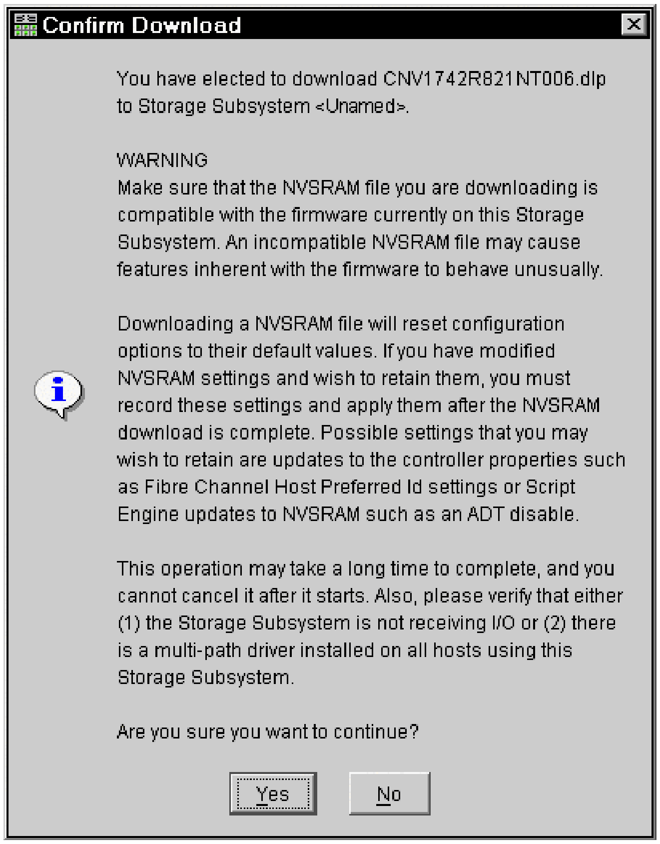

Figure 4-20 Confirm Download of the new NVSRAM

Since the NVSRAM is much smaller than the firmware package, it doesn't take as long as the firmware download. The window in Figure 4-21 shows the status of the process.

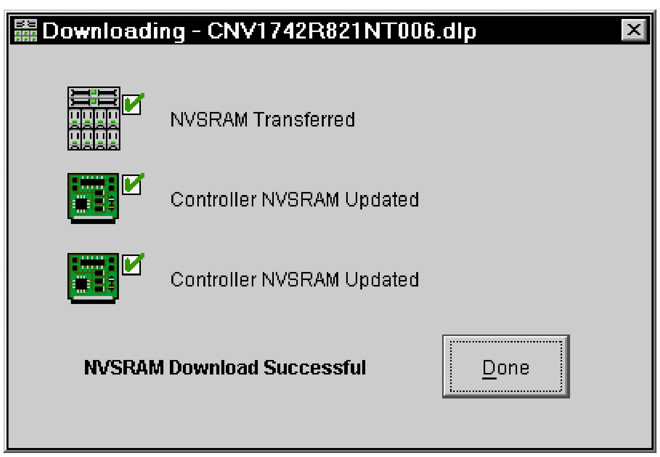

Figure 4-21 Download of NVSRAM completed

After the upgrade procedure, it is not necessary to power cycle FAStT. After the download, the controllers are rebooted automatically one by one and the FAStT Storage Server is online again.

5. If the FAStT Storage Server is not recognized or unresponsive after the upgrade in the Enterprise Management windows, remove the device from the Enterprise Management window and initiate a new discovery. If the FAStT Storage Server is still unresponsive,

reboot the host system and initiate a discovery when the system is up again. This can be caused by the host agent not recognizing properly the updated FAStT.

4.2.4 Updating the drive and ESM board microcode

It may also be necessary to update the microcode of the fibre disk drives and the ESM modules contained in the EXP storage enclosures. These updates cannot be done with the normal FAStT Client. A special version of the FAStT Client, called the *FAStT Field Tool*, is available. The download package contains the microcode for the drives and the ESM boards. You can download this tool from the Web at:

http://www.storage.ibm.com/hardsoft/disk/fastt/index.html

The FAStT Field Tool is only available for Windows NT 4.0 or Windows 2000. If you don't have a system attached to the FAStT Storage Server running a Windows operating system, you can install the FAStT Field Tool on any Windows system that has a TCP/IP connection to the FAStT Storage Server.

Install the Field Tool on your system. The installation is the same as the installation of the FAStT Client package as explained in 4.1.1, "Microsoft Windows NT 4.0 and Windows 2000 installation" on page 66.

Start the FAStT Field Tool. As with the client, you are asked for an initial discovery of the storage subsystems. If the FAStT Storage Server in question is in another TCP/IP network, you need to add the device manually (see also Figure 4-11 on page 87). When the subsystem is added, open the Subsystem Management window for the FAStT Storage Server.

Ensure the proper configuration and status of the system.

Important: All I/O activity must be stopped during the microcode update of the disk drives and the ESM boards.

Updating the disk drive microcode

To update the disk drive microde, follow these steps:

1. Start the download procedure. Click **Diagnostics-> Download Drive Firmware** in the subsystem management window.
2. A dialog appears that allows you to select the drives that must be updated. Click **Product ID**. The drives are listed by type, which makes it easier to select drives of the same type.
3. Highlight all drives of the same Product ID. Click the **Browse** button and choose the correct firmware file. The name of the file must completely match with the model name of the disks you are updating (see Figure 4-22).

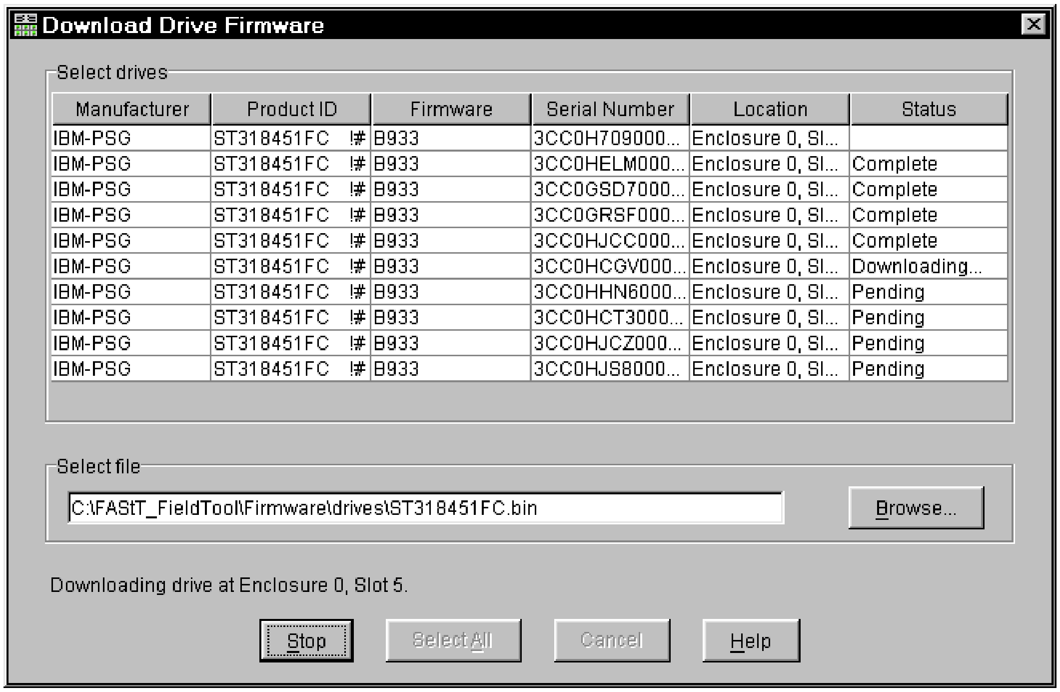

Figure 4-22 Field Tool Select drives to update microcode

4. Double-click the selected file.
5. The Confirm Download window opens as shown in Figure 4-23. Click **Yes** to confirm the download and the process begins.

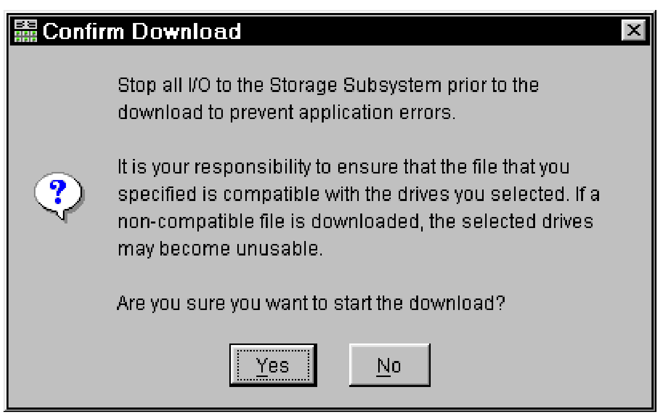

Figure 4-23 Field Tool Confirm Drive Microcode Download

This procedure can take a long time, depending on the number of drives that need to be updated. Be sure to deny all I/O during this time.

In the Status column, you can see the progress for each drive.

6. Repeat the steps for all product types of drives connected to the FAStT Storage Server.

Updating the ESM board microcode

It may be necessary to update the microcode of the ESM boards of the attached EXP 500 or EXP 700. This is a similar procedure as for updating the drive microcode. Follow these steps:

1. In the Field Tool, click **Download-> ESM Firmware** from the Subsystem Management window.

2. A new window opens showing you all attached EXP storage enclosure. As with the drives, highlight all lines with the same product and choose the correct firmware file to be downloaded to the ESM boards (Figure 4-24).

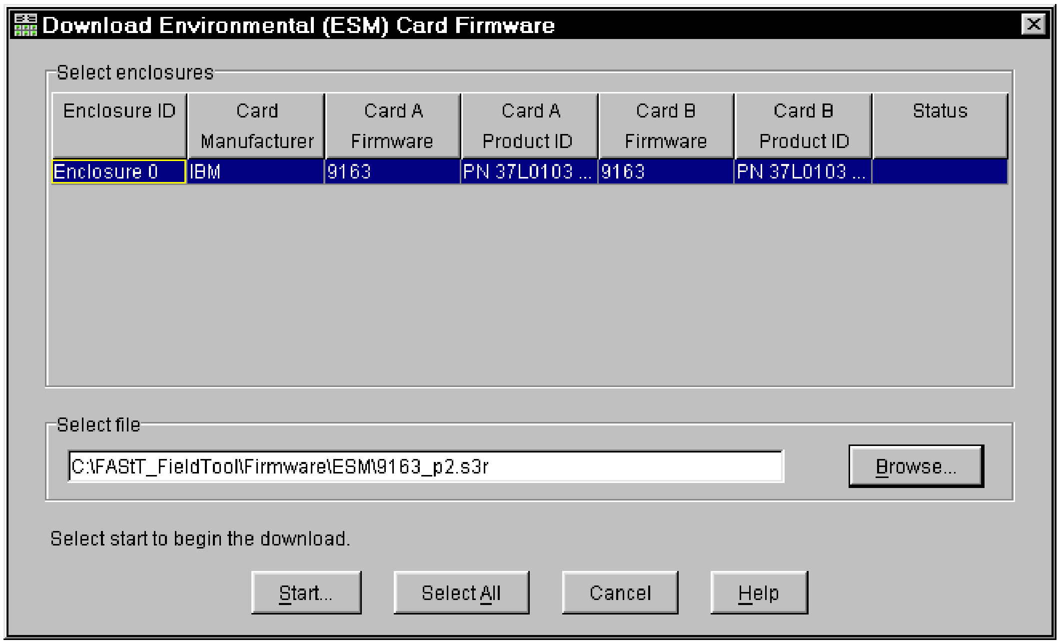

Figure 4-24 Field Tool to select the ESM boards

3. As with the disk drive microcode download, a confirmation window opens before the actual download starts. Be sure that no I/O appears during the update of the ESM boards.

4.2.5 Initial configuration steps

Before defining any arrays or logical drives, you must perform some basic configuration steps. This also applies when you reset the configuration of your FAStT Storage Server, either with the Storage Manager software or through the serial connection. Complete these steps:

1. If you installed more than one FAStT Storage Server, it is important to give it a literal name. To rename the FAStT Storage Server, open the Subsystem Management window. Right-click the subsystem, and click **Storage Subsystem-> Rename** (Figure 4-25).

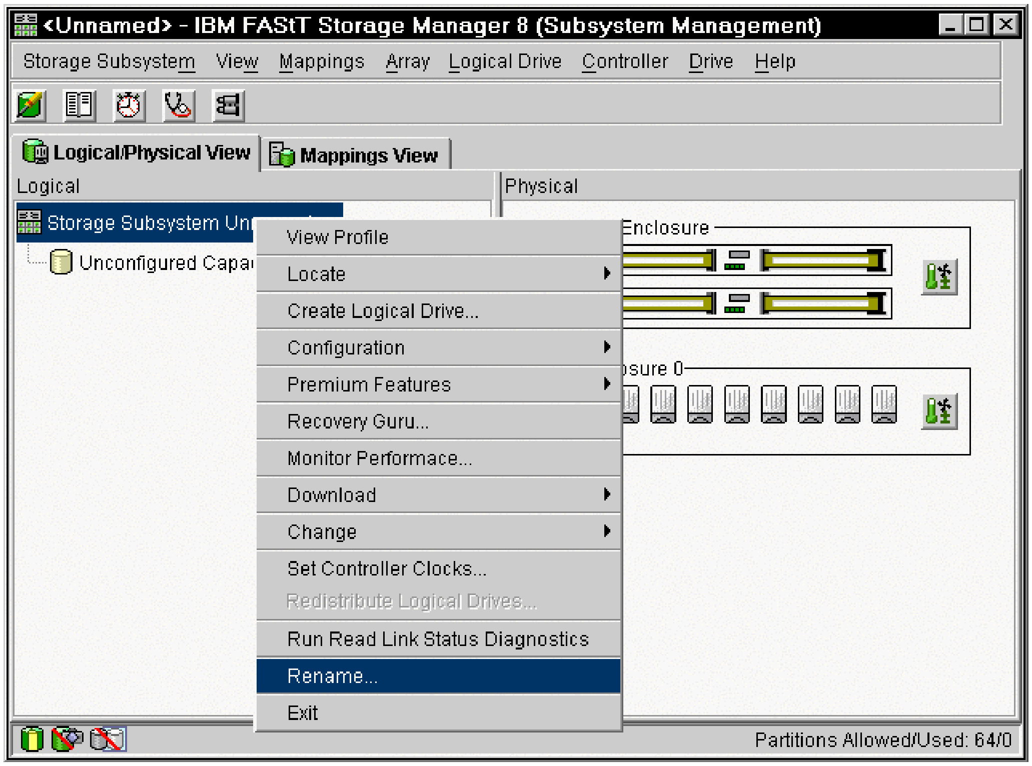

Figure 4-25 Renaming the FAStT Storage Server

2. Enter a new name for the subsystem in the Rename Storage Subsystem dialog (Figure 4-26).

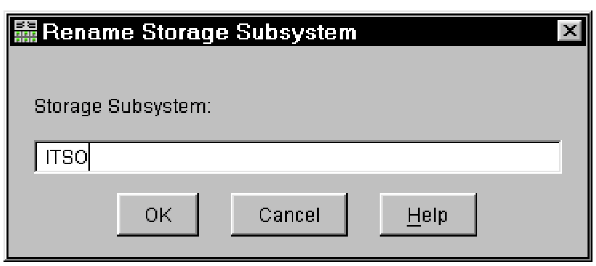

Figure 4-26 Renaming FAStT

3. Since the FAStT Storage Server stores its own event log, synchronize the controller clocks with the time of the host system. Be sure that your local system is working using the correct time. Then click **Storage Subsystem-> Set Controller Clock** (Figure 4-27).

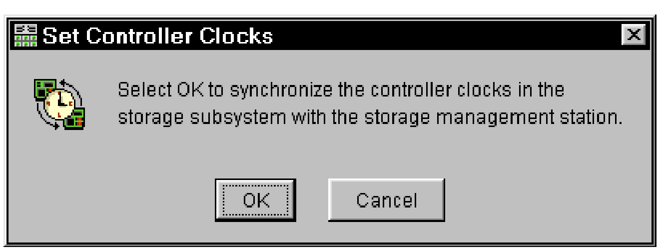

Figure 4-27 Setting the controller clock

Make sure the time of the controllers and the attached systems are synchronized. This simplifies error determination if you start comparing the different event logs.

4. For security reasons, especially if the FAStT Storage Server is directly attached to the network, you should set a password. This password is required for all actions on the FAStT Storage Server that change or update the configuration in any way.

   To set a password, highlight the storage subsystem, right-click, and click **Change-> Password** (Figure 4-28). This password is then stored on the FAStT Storage Server. It is used if you connect through another FAStT Client or the FAStT Field Tool. It is not important whether you are using in-band or out-of-band management.

Figure 4-28 Setting a password on the FAStT Storage Server

4.2.6 Defining hot spare drives

Hot spare drives are special reserved drives that are normally not used to store data. But if a drive in a RAID array with redundancy, such as 1 or 5, fails, the hot spare drive takes on the function of the failed drive and the data is recovered on the hot spare drives, which become part of the array. After this procedure, your data is again fully protected. Even if another drive fails, this cannot affect your data.

If the failed drive is replaced with a new drive, the data stored on the hot spare drive is copied back to the replaced drive, and the original hot spare drive that is now in use becomes a free hot spare drive again. The location of a hot spare drive is fixed and does not wander if it is used.

A hot spare drive defined on the FAStT Storage Server is always used as a so-called *global hot spare*. That is, a hot spare drive can always be used for a failed drive; it is not important in which array or storage enclosure it is situated.

A hot spare drive must be at least of the capacity of the configured space on the failed drive. the FAStT Storage Server can use a larger drive to recover a smaller failed drive to it. Then the remaining capacity is blocked.

If you plan to use several hot spare drives, the FAStT Storage Server uses a certain algorithm to define which hot spare drive is used. The controller first attempts to find a hot spare drive on the same channel as the failed drive. The drive must be at least as large as the configured capacity of the failed drive. If a hot spare drive does not exist on the same channel, or if it is already in use, the controller checks the remaining hot spare drives, beginning with the last hot spare configured. For example, the drive in enclosure 1, slot 4, may fail and the hot spare drives may be configured in the following order:

- HSP 1: enclosure 0, slot 12
- HSP 2: enclosure 2, slot 14
- HSP 3: enclosure 4, slot 1
- HSP 4: enclosure 3, slot 14

In this case, the controller checks the hot spare drives in the following order:

- 3:14
- 4:1
- 2:14
- 0:12

The controller uses a free hot spare drive as soon as it finds one, even if there is another one that may be closer to the failed drive.

To define a hot spare drive, highlight the drive you want to use. Click **Drive-> Hot Spare-> Assign** as shown in Figure 4-29.

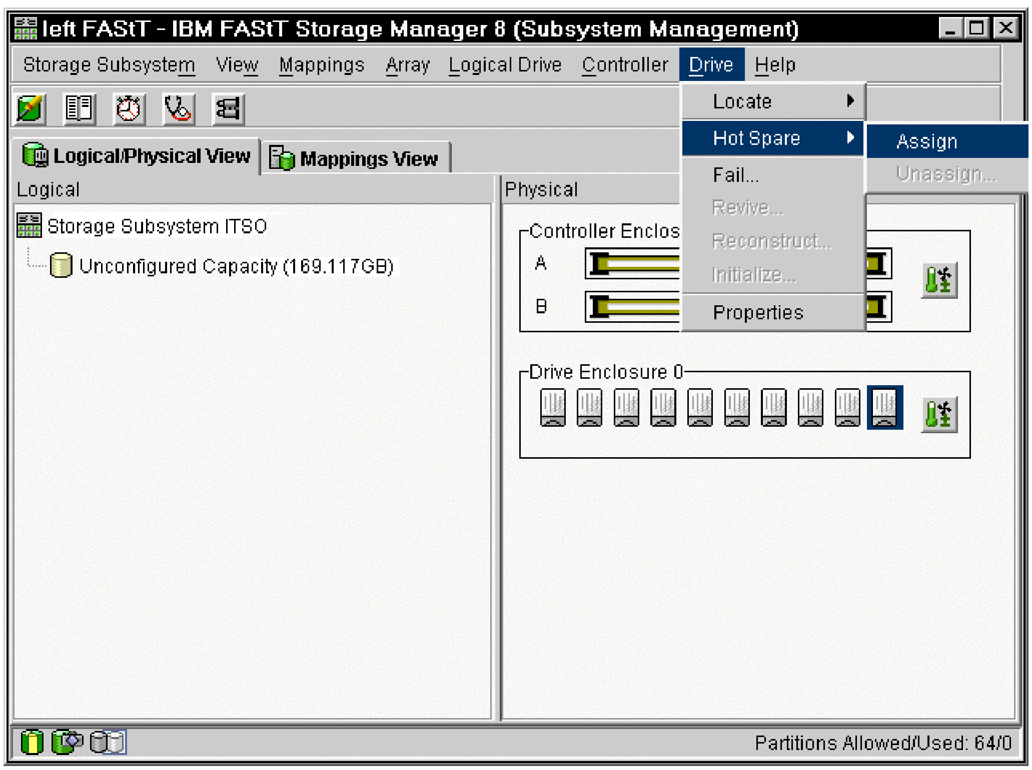

Figure 4-29 Defining a hot spare drive

If there are larger drives defined in any array on the FAStT Storage Server than the drive you have chosen, a warning message appears and notifies you that not all arrays are protected by the hot spare drive.

The new defined hot spare drive then has a small red cross in the lower part of the drive icon.

Especially in large configurations with arrays containing a lot of drives, it may be necessary to define multiple hot spare drives since the reconstruction of a failed drive to a hot spare drive can last for a long time.

To unassign a hot spare drive and have it available again as a free drive, highlight the hot spare drive and select **Drive-> Hot Spare-> Unassign**.

4.3 Creating arrays and logical drives

The storage subsystem is now installed and upgraded to the newest microcode level. Now the arrays and logical drives can be configured. If you are not sure how to divide the available drives into arrays or logical drives, and which restrictions apply, to avoid improper or inefficient configurations of the FAStT Storage Server, refer to 3.3, “Advanced management” on page 51.

The following example starts with a blank FAStT. The same steps are necessary to define additional logical drives on an existing configuration. The main difference is that you have to decide whether you use unconfigured capacity on free disks or free capacity in an already existing array.

1. In the Subsystem Management window (Figure 4-30), right-click the unconfigured capacity and select **Create Logical Drive**.

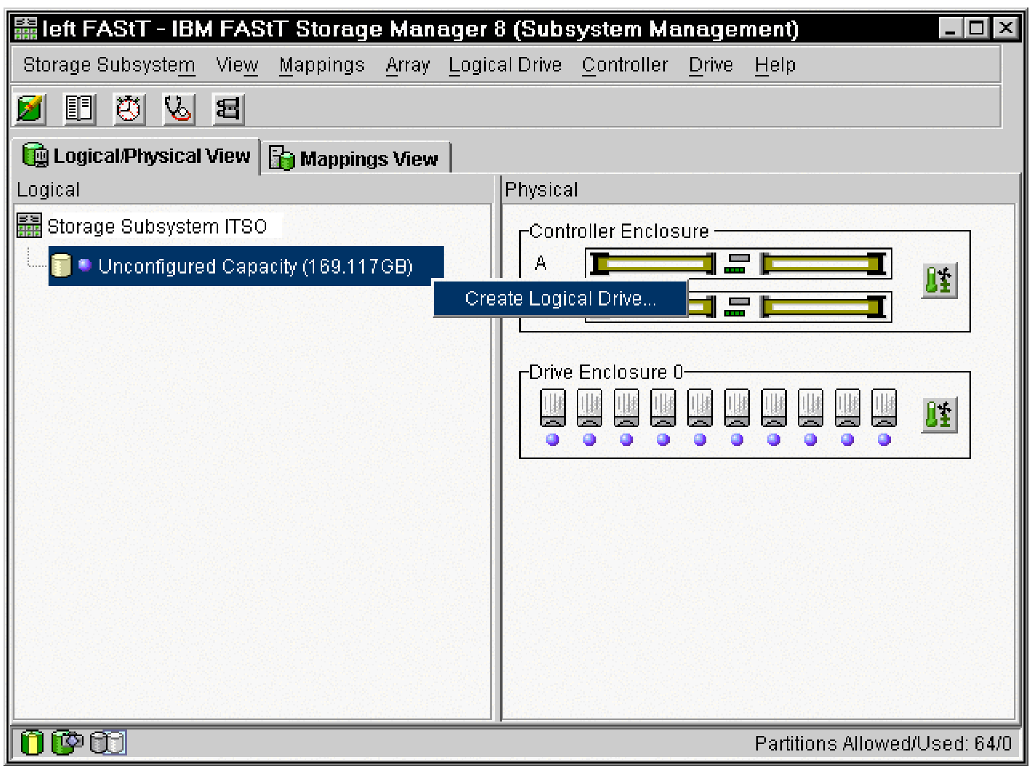

Figure 4-30 Creating a logical drive

   a. If this is the first time you are creating a logical drive from this instance of the FAStT Client, the window shown in Figure 4-31 appears. The Current default host type should be set to **Windows Non-Clustered (SP5 or higher)**. You should only change this type if you attach hosts running the same operating system.

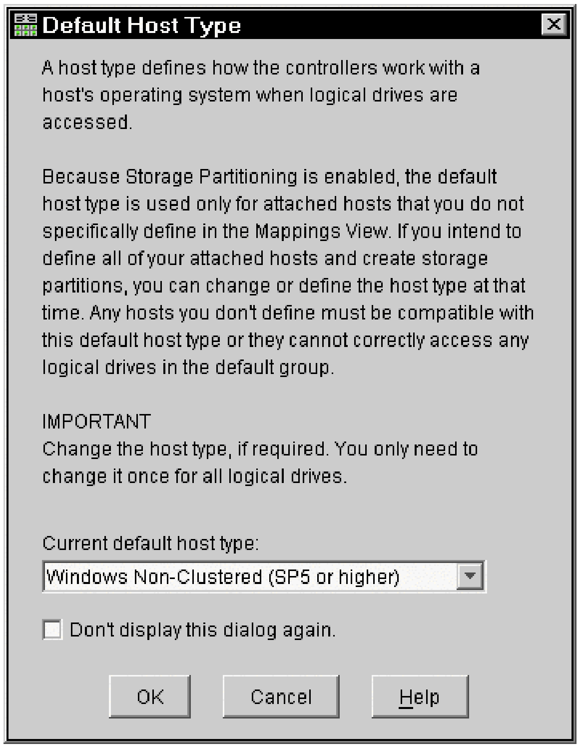

Figure 4-31 Default Host Type window

   b. You can select the check box to disable this dialog now. If you want to change Default Host Type later, click **Storage Subsystem-> Change-> Default Host Type** in the Subsystem Management window to access the window later.

   c. Click **OK**.

2. After you complete this window, the Create Logical Drive Wizard appears. It leads you through the creation of your logical drives. Since we don't have any configured arrays yet, you need to select the **Unconfigured capacity** option to create a new array, as shown in Figure 4-32. Then click **Next**.

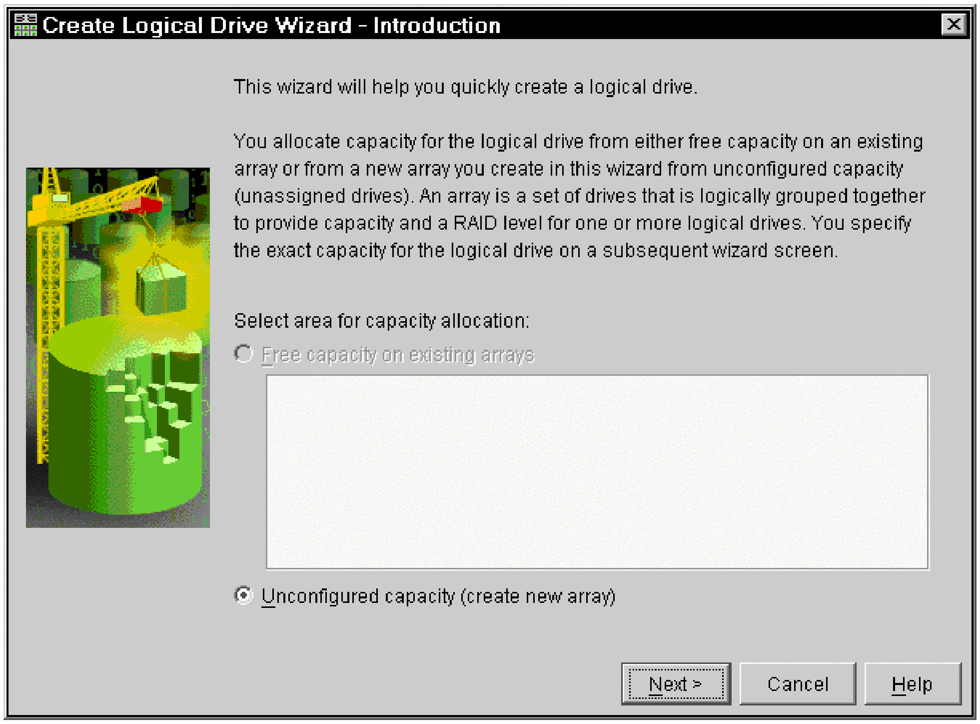

Figure 4-32 Create Logical Drive Wizard

3. Choose the RAID level and number of drives that are used for the array (see Figure 4-33). There are two possibilities for the drive selection. You can choose each drive manually. For this, you need to switch to the *Manual* drive selection option. Be sure that you consider performance and availability concerns when you choose the drives manually.

   Unless you have specific requests and can watch performance and especially availability while choosing the drives manually, we recommend that you keep the default of **Automatic** for the drive selection. Then you only need to choose the number of drives and the capacity you want to use. The storage software shows all possible array combinations for the RAID level you selected. There is only a selectable drive combination, if all drives are of the same size. If you choose a specific array, the software selects the best combination of available drives to keep up with performance and availability.

   You need to define the desired RAID level first, before you choose the number of drives you want to use.

   Click **Next**.

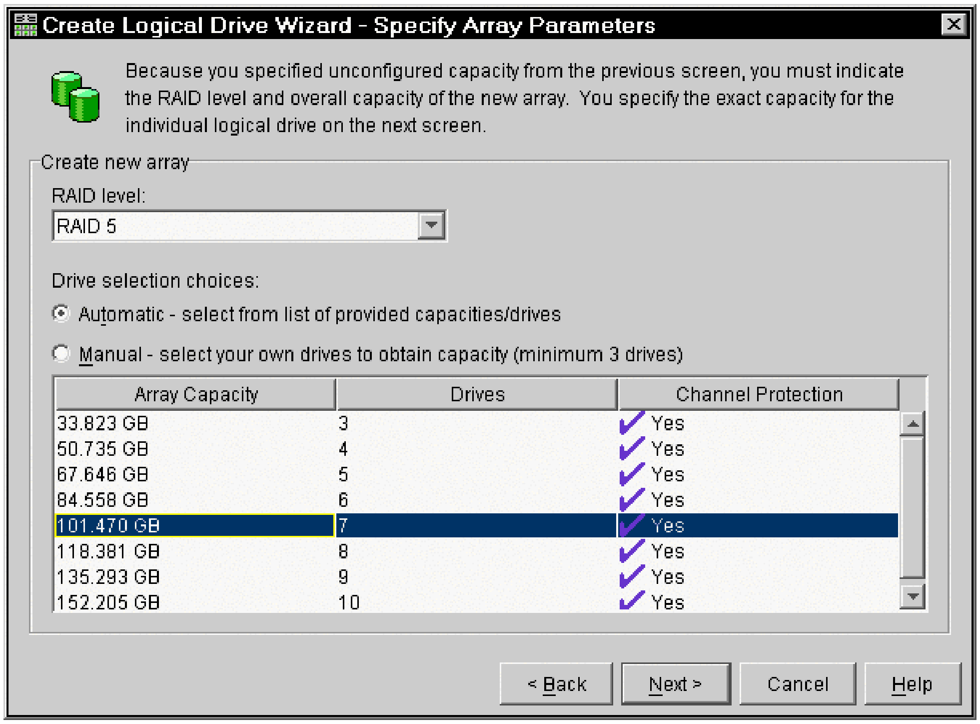

Figure 4-33 Choosing the drives used for the new array

4. Define the logical drive (Figure 4-34). By default, all available space in the array is configured as one logical drive.
   a. If you want to define more than one logical drive in this logical drive, enter the desired size.
   b. Assign a name to the logical drive, which complies to our mapping table (Table 4-2 on page 86).
   c. If you want to change advanced logical drive settings as segment size or cache settings, select the **Customize settings** option.

   Otherwise the creation of the first logical drive is finished. Click **Finish**.

   The newly created logical drive is not mapped automatically but remains unmapped. Otherwise, the drive is immediately seen by the attached hosts. If you change the mapping later, the logical drive, which appears as a physical drive to the operating system, is removed without notifying the hosts. This can cause severe problems.

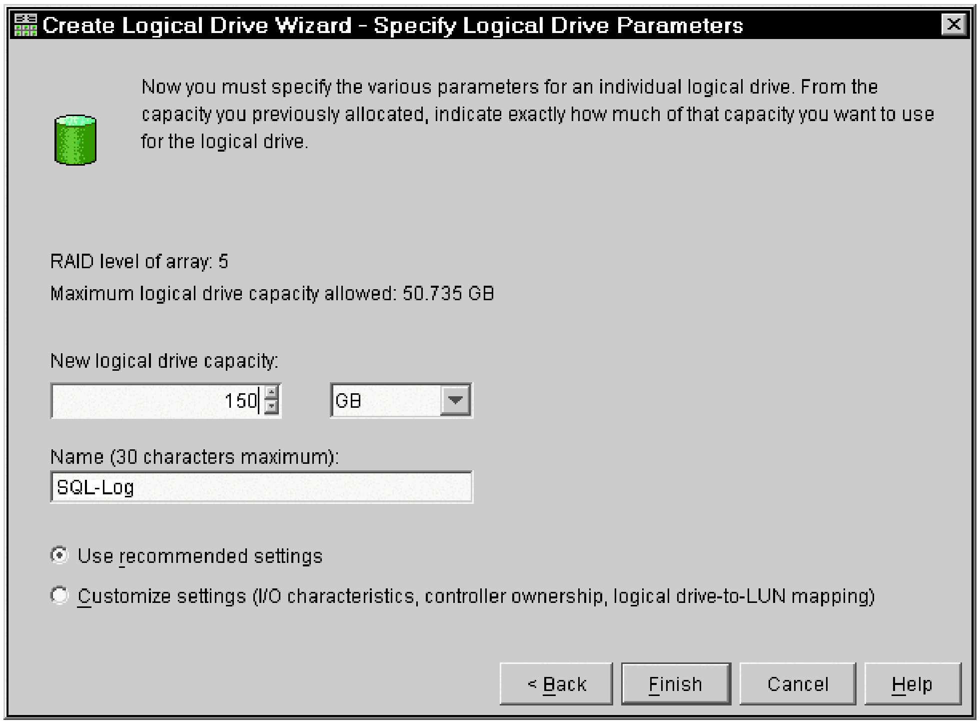

Figure 4-34 Logical drive parameters

5. On the Specify Advanced Logical Drive Parameters panel (Figure 4-35), you define the logical drive exactly to suit your needs:
   a. For Logical Drive I/O characteristics, you can specify file system, database, or multimedia base. Or you can manually set the parameters for the logical drive by choosing *Custom*.
   b. The segment size is chosen according to the usage pattern. For custom settings, you can directly define the segment size.
   c. You can also define the cache read ahead multiplier. Begin by choosing only small values. Otherwise large parts of the cache are filled by read ahead data that may never be used.
   d. The preferred controller handles the logical drive normally if both controllers and I/O paths are online. You can load balance your logical drives throughout both controllers. The default is to alternate the logical drives on the two controllers.

      Obviously it is better to spread the logical drives by the load they cause on the controller. It is possible to monitor the load of each logical drive on the controllers with the performance monitor and move the preferred controller (see 8.1, "Performance monitoring and tuning" on page 250).
   e. You can choose to set the Logical Drive to LUN mapping parameter to run automatically or to be delayed by mapping it later with storage partitioning. See 3.2.2, "Storage partitioning" on page 45. If you choose to map later to the default host group, keep in mind that the logical drive becomes visible immediately after the creation.
   f. Click **Finish**.

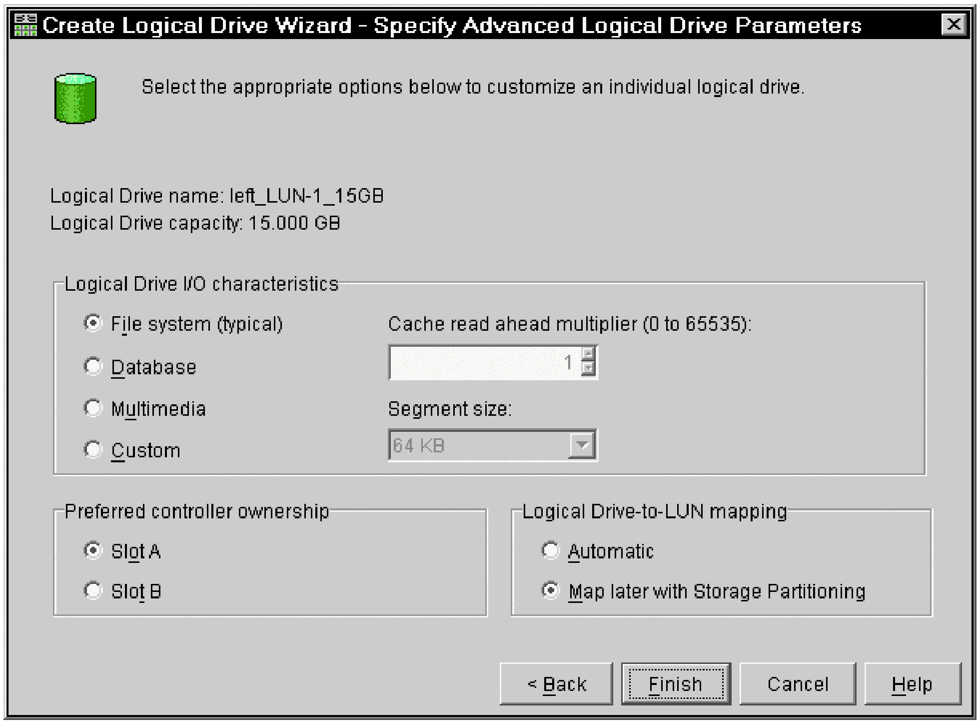

Figure 4-35 Advanced logical drive settings

6. If the logical drive is smaller than the total capacity of the array, a window opens and asks whether you want to define another logical drive on the array. The alternative is to leave the space as unconfigured capacity. For now, leave the capacity unconfigured.

After you define all logical drives on the array, the array is now initialized and immediately accessible.

If you left unconfigured capacity inside the array, you can define another logical drive later in this array. Simply highlight this capacity, right-click, and choose **Create Logical Drive**. Simply follow the steps that we outlined in this section, except for the selection of drives and RAID level. Since you now already defined arrays that contain free capacity, you can choose where to store the new logical drive – on an existing array or on a new one.

4.4 Configuring storage partitioning

Since heterogeneous hosts can be attached to the FAStT Storage Server, you need to configure storage partitioning for two reasons:

- Each host operating system requires slightly different settings on the FAStT Storage Server, so you need to tell the storage subsystem the host type that is attached.
- There is interference between the hosts if every host has access to every logical drive. By using storage partitioning and LUN masking, you ensure that each host or host group only has access to its assigned logical drives.

To configure storage partitioning, follow these steps:

1. Choose the **Mappings View** in the Subsystem Management window. All information, such as host ports and logical drive mappings, are shown and configured here. The right side of the window lists all mappings that are owned by the object you choose in the left side. If

you highlight the storage subsystem, you see a list of all defined mappings. If you highlight a specific host group or host only, its mappings are listed.

2. Define the host groups. Highlight the **Default Group**, right-click, and choose **Define Host Group** as shown in Figure 4-36.

   If no mappings are defined yet, a window with the first steps for defining a storage partition opens.

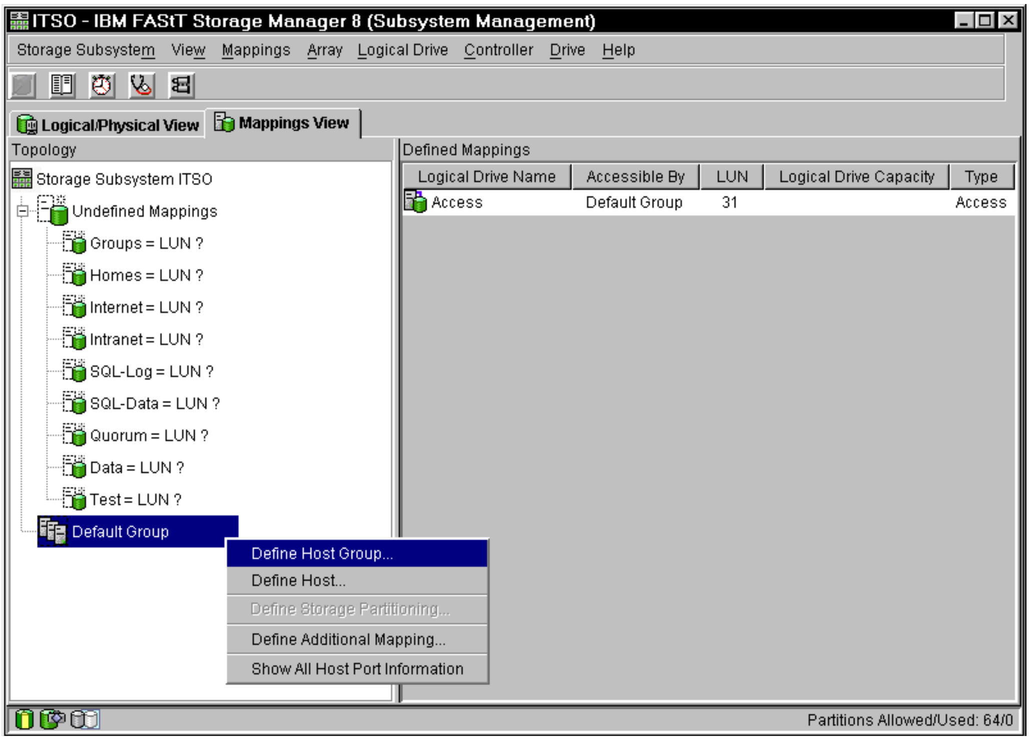

Figure 4-36 Define Host Group

3. The Define Host Group dialog (Figure 4-37) opens. Enter the names of the host groups you need. When you are finished, click **Close** to exit the dialog. Note that if only one server will access the logical disks in a storage partition, then it is not necessary to define a host group because you could use the default host group. However, as requirements are constantly changing, we recommend that you define a host group. Otherwise the addition of new systems is not possible without disrupting the already defined mappings in the default host group.

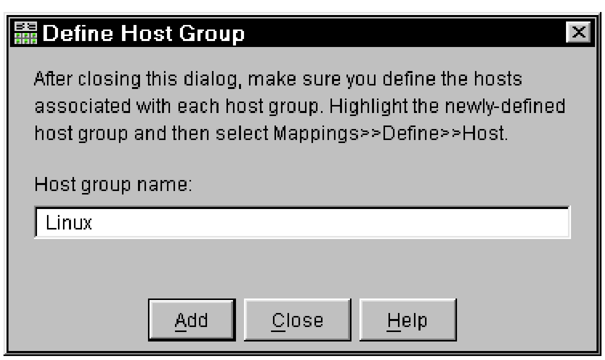

Figure 4-37 Entering the host group name

4. The hosts groups are defined and the hosts in these groups can now be defined according to the example in Table 4-2 on page 86. As shown in Figure 4-38, highlight the group for which you want to add a new host. Right-click and choose **Define Host**.

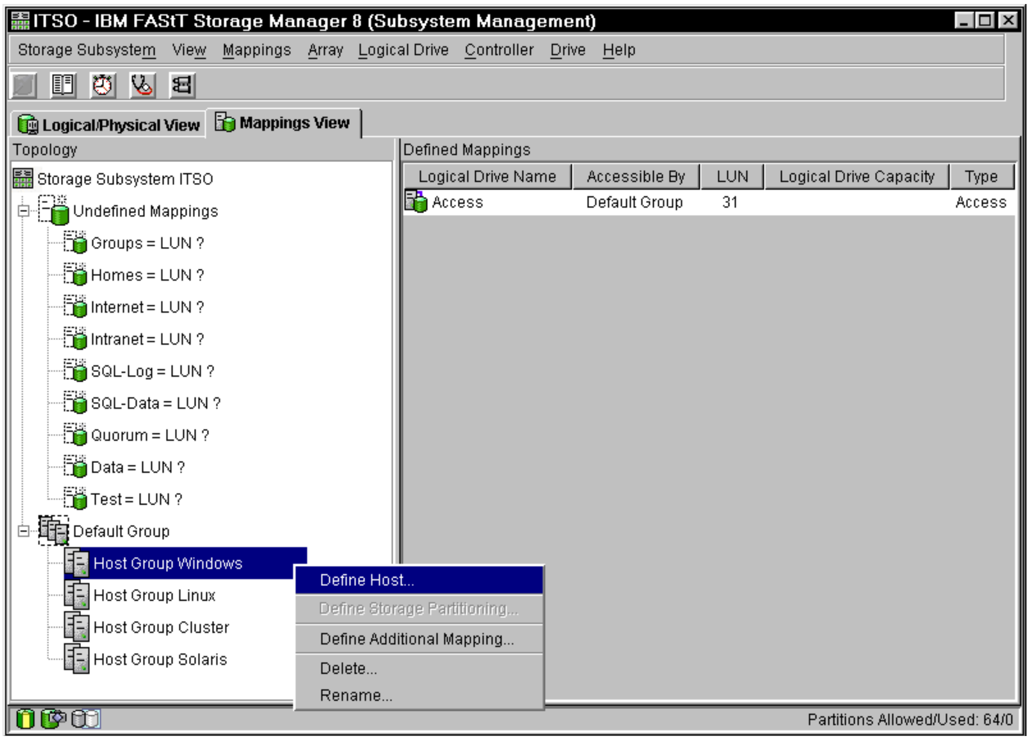

Figure 4-38 Selecting Define Host

5. On the Define Host dialog (Figure 4-39), enter the name of the host you want to define in the selected group. If you finished entering all hosts for this group, click **Close**. Then choose the next host group and define the hosts in that group.

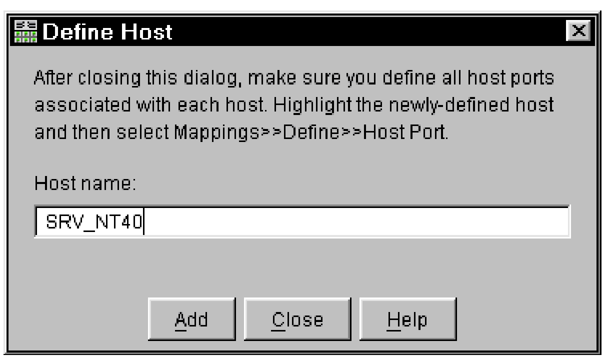

Figure 4-39 Define Host

If you accidentally assigned a host to the wrong host group, you can move the host to another group. Simply right-click the host name and select **Move**. A pop-up window opens and asks you to specify the host group name.

6. Because storage partitioning of the FAStT Storage Server is based on the world-wide names of the host ports, the definitions for the host groups and the hosts only represent a view of the physical and logical setup of your fabric. When this structure is available, it is much easier to identify which host ports are allowed to see the same logical drives and which are in different storage partitions.

Storage partitioning is not the only function of the storage server that uses the definition of the host port. When you define the host ports, the operating system of the attached host is defined as well. Through this information, FAStT can adapt the RDAC or ADT settings for the hosts.

It is important to choose the correct operating system from the list of available operating systems, because this is the part of the configuration where you configure the heterogeneous host support. Each operating system expects slightly different settings and handles SCSI commands a little differently. Therefore, it is important to select the correct value. If you don't, your operating system may not boot anymore or path failover cannot be used if connected to the storage subsystem.

The host port is identified by the world-wide name of the host bus adapter. Highlight the host, right-click, and choose **Define Host Port** as shown in Figure 4-40.

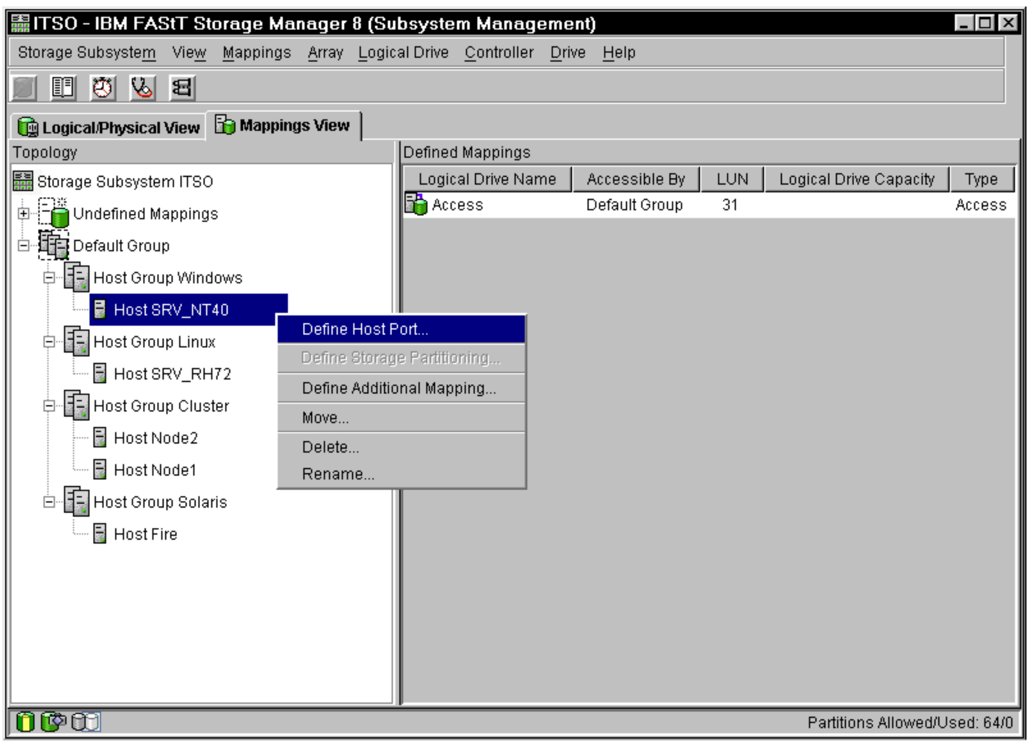

Figure 4-40 Choosing Define Host Port

7. In the Define Host Port dialog (Figure 4-41), enter the port name for this adapter and choose the correct operating system. The host port identifier corresponds to the world-wide name of the adapter port. In the drop-down box, you only see the world-wide names that are currently active. If you want to enter a host port that is not currently active, type the world-wide name in the field. Be sure to check for typing errors. The values should correspond to the example in Table 4-2 on page 86.

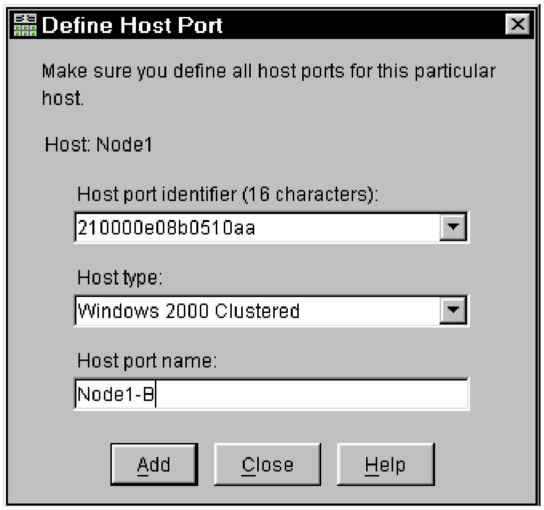

Figure 4-41 Enter the Host Port information

8. Define the mapping for each logical drive created in 4.3, "Creating arrays and logical drives" on page 100. All the information entered on the Define Host Port dialog is needed to ensure a proper operation in a heterogeneous environment with multiple servers attached to the FAStT Storage Server.

   As shown in Figure 4-42, highlight the host group to which you want to map a new logical drive. Right-click and choose **Define Additional Mapping**.

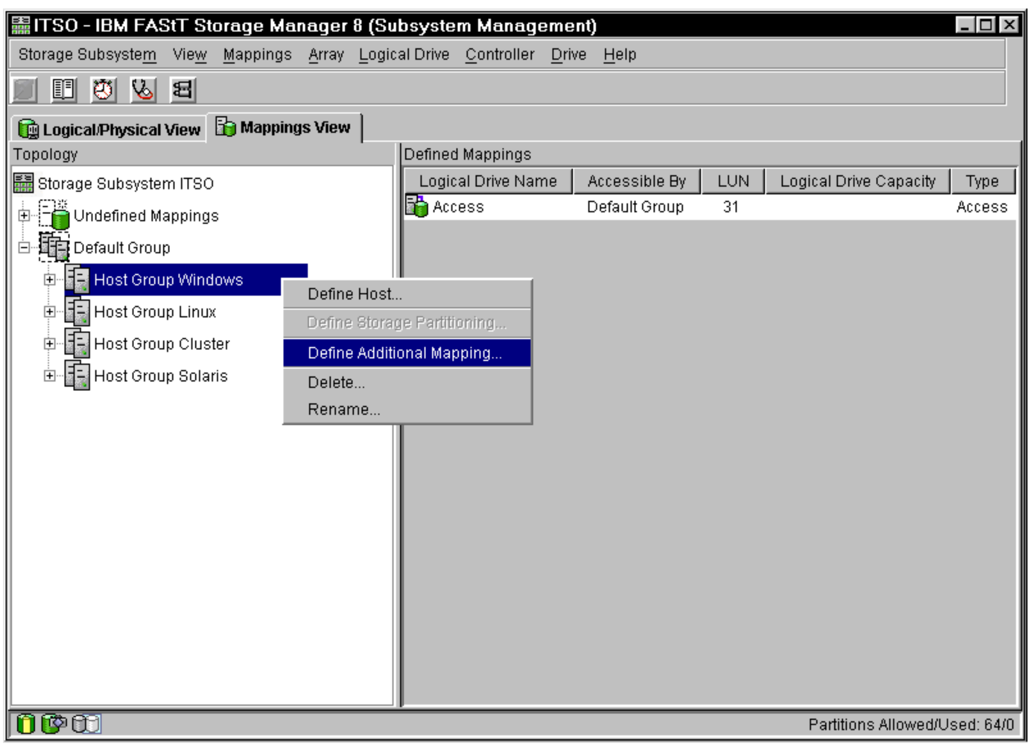

Figure 4-42 Define Additional Mapping

9. In the Define Additional Mapping dialog, select the logical drive you want to map to this host group and assign the correct LUN number (Figure 4-43), according to the mapping table (Table 4-2 on page 86).

a. In the top drop-down list, you can choose the host group or host to which you want to map the logical drive.

b. With the logical unit number, you can influence the order in which the mapped logical drives appear. Starting with LUN 0, the logical drive appears in the operating system.

c. In the list box that follows, you see all unmapped drives. Choose the logical drive you want to map.

If you entered all the information, click **Add** to finish defining this mapping. The first mapping is now defined. In the Subsystem Management window, you see that the number of used storage partitions changed from 0/64 to 1/64.

You can define all other mappings by repeating these steps. You receive an error message after the last logical drive is mapped to a host group or host.

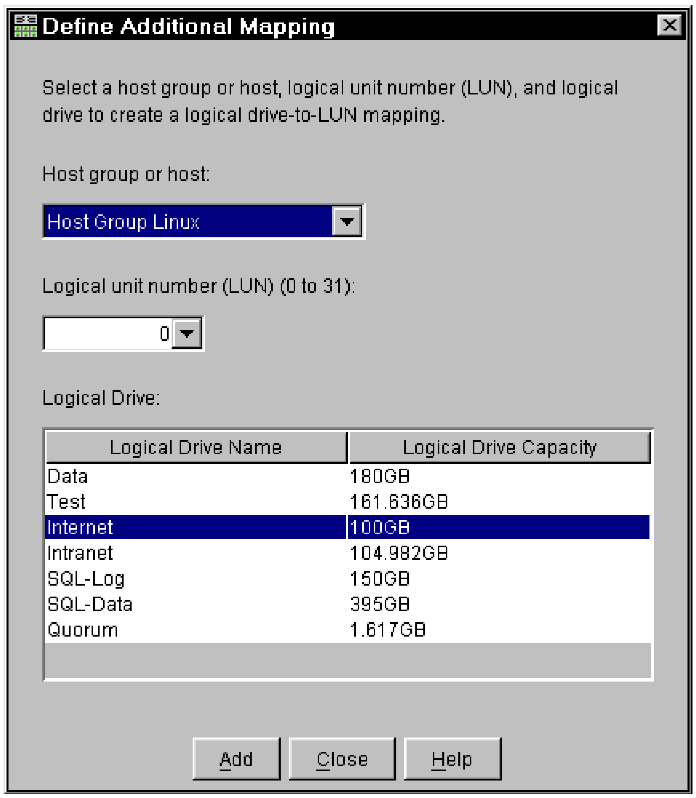

Figure 4-43 Define Additional Mapping

If you create a new mapping or change an existing mapping of a logical drive, the change happens immediately. Therefore, make sure that this logical drive is not in use or even assigned by any of the machines attached to the storage subsystem.

Now all logical drives and their mappings are defined according to our mapping table (Table 4-2 on page 86). All logical drives are now accessible by their mapped host systems.

To make the logical drives available to the host systems without rebooting, the FAStT Utilities package provides the hot_add command line tool for some operating systems. You simply run hot_add, and all host bus adapters are rescanned for new devices and the devices are assigned within the operating system. Linux now requires a new configuration done with the FAStT MSJ.

You may have to take appropriate steps to enable the use of the storage inside the operating system, such as formatting the disks with a file system and mount them.

If you attached a Linux or AIX system to the FAStT Storage Server, you need to delete the mapping of the access logical drive. Highlight the host or host group containing the Linux or AIX system in the Mappings View. In the right part of the window, you see the list of all logical drives mapped to this host or host group. To delete the mapping of the access logical drive, right-click it and choose **Delete**. The mapping of the access logical drive is deleted immediately.

4.5 Finishing the initial configuration

So far we updated the storage subsystem to the latest level, configured logical drives, and mapped them to the attached hosts. There are only two steps left to finish the installation.

The next step is to define the alerting methods in case of failure. The final step is to document the configuration as it is now.

4.5.1 Monitoring and alerting

Included in the FAStT Client package is the Event Monitor service. It enables the host running this monitor to send out alerts via e-mail (SMTP) or traps (SNMP). The Event Monitor can be used to alert problems with any of the FAStT Storage Servers in your environment.

It should be installed and configured on at least two systems that are attached to the storage subsystem and allow in-band management running 24 hours a day. This ensures proper alerting, even if one server is down.

Depending on the setup you choose, different storage subsystems are monitored by the Event Monitor. If you right-click your local system in the Enterprise Management window (at the top of the tree) and choose **Alert Destinations**, this applies to all storage subsystems listed in the Enterprise Management window (Figure 4-44). Also if you see the same storage subsystem through different paths, directly attached and through different hosts running the host agent, you receive multiple alerts. If you right-click a specific storage subsystem, you only define the alerting for this particular FAStT Storage Server.

An icon in the lower left corner of the Enterprise Management window indicates that the Event Monitor is running on this host.

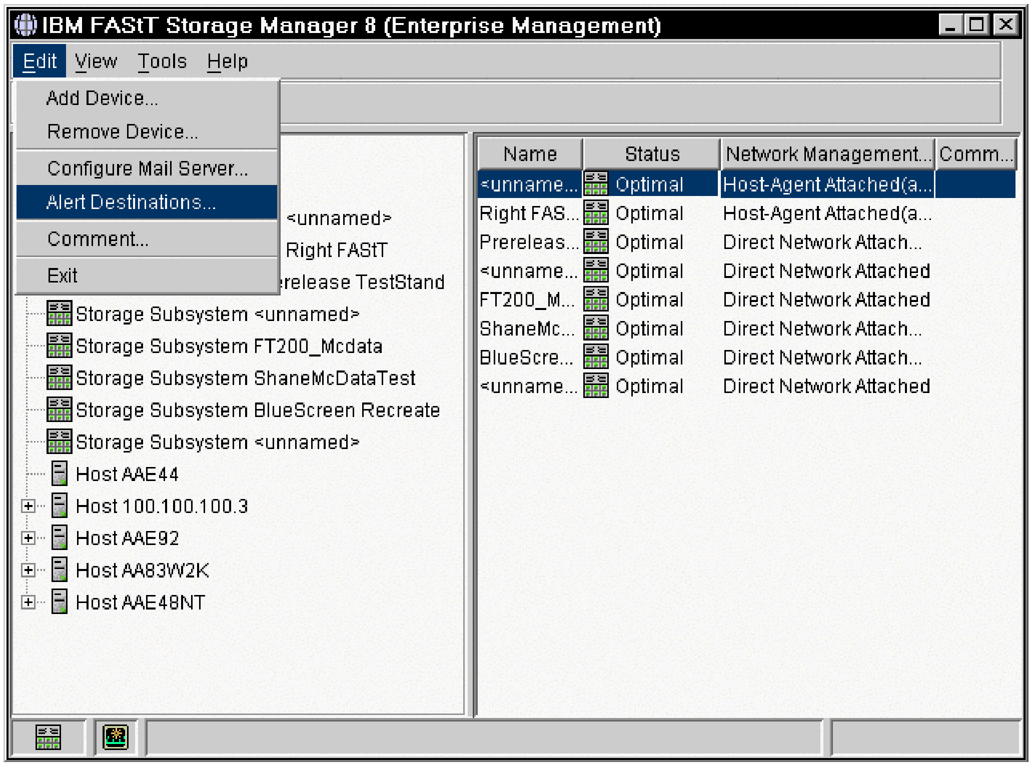

Figure 4-44 Defining Alert Destinations

When you remove or add new devices in the Enterprise Management window, the list of devices is not automatically synchronized with the Event Monitor. If there is a mismatch between the devices listed in the Enterprise Management window and the devices known to the Event Monitor, the icon of the Event Monitor in the menu bar is highlighted. To synchronize the Event Monitor list with the Enterprise Management list of storage subsystems, click the icon shown in Figure 4-45.

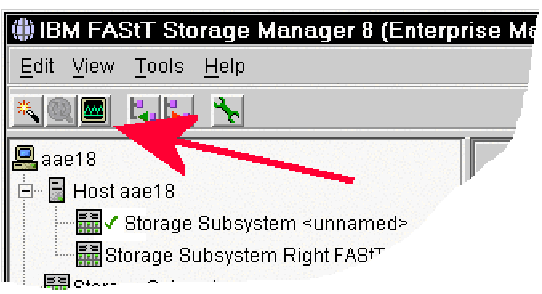

Figure 4-45 Unsynchronized Event Monitor

If you want to send e-mail alerts, you have to define an SMTP server first. Click **Edit-> Configure Mail Server**. Enter the IP address or the name of your mail server and the sender address.

If you open the Alert Destination dialog, you define the e-mail addresses to which alerts are sent. If you do not define an address, no SMTP alerts are sent. You also can validate the e-mail addresses to ensure a correct delivery and test your setup.

If you choose the SNMP tab (Figure 4-46), you can define the settings for SNMP alerts: the IP address of your SNMP console and the community name. As with the e-mail addresses, you can define several trap destinations.

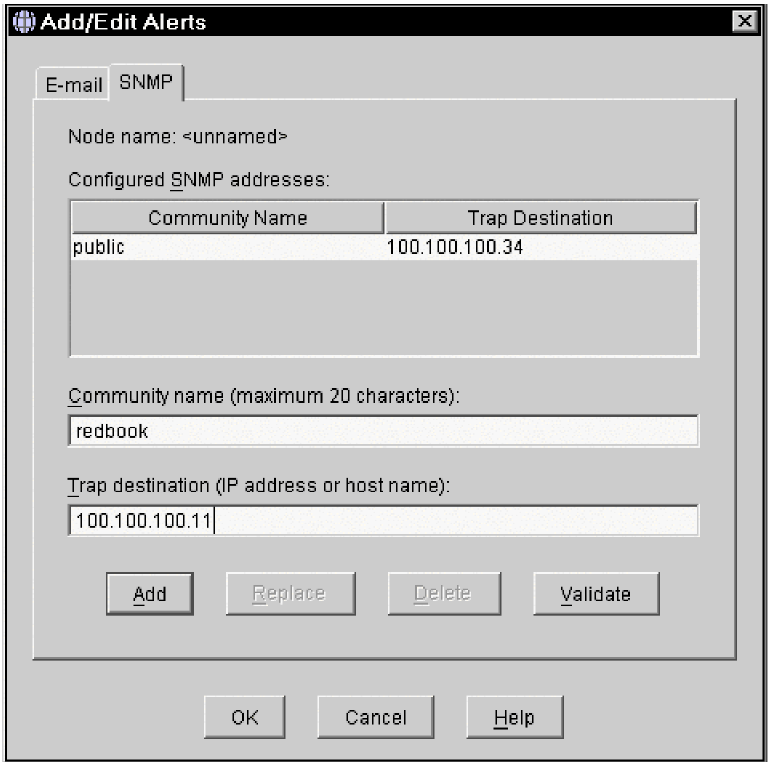

Figure 4-46 Adding or editing alert destinations

You need an SNMP console for receiving and handling the traps sent by the service. There is an MIB file included in the Storage Manager software, which should be compiled into the SNMP console to allow proper display of the traps. Refer to the documentation of the SNMP console you are using to learn how to compile a new MIB.

Notice the green check mark in front of the subsystem. The location of the check mark shows on which level the alert is defined. All FAStT Storage Servers beneath this level are monitored and alerts are processed. In Figure 4-45, only alerts for the unnamed storage subsystem, managed through the host agent running on system AAE18, are sent.

4.5.2 Saving the subsystem profile

Configuring a FAStT Storage Server is a complex task. Therefore, the so-called subsystem profile is a single location where all the information on the configuration is stored. The profile includes information on the controllers, attached drives and enclosures, their microcode levels, arrays, logical drives, and storage partitioning.

Tip: You should save a new profile each time you change the configuration of the FAStT storage subsystem even for minor changes. This applies to all changes regardless of how minor they may be. The profile should be stored in a location where it is available even after a complete configuration loss, for example after a site loss.

To obtain the profile, open the Subsystem Management window and click **View-> Storage Subsystem Profile**.

All information in the profile is gathered from the various components when you request the profile. The profile can be saved locally and included in the documentation to maintain a change history for the storage subsystem. We recommend that you save a new version of the profile and store it securely whenever a configuration change takes place. Even in the case of a complete configuration loss, you can restore the array and logical drive configuration as well as the mappings for the storage partitioning. This is particularly interesting for scenarios that use Remote Volume Mirroring. The profile window is shown in Figure 4-47.

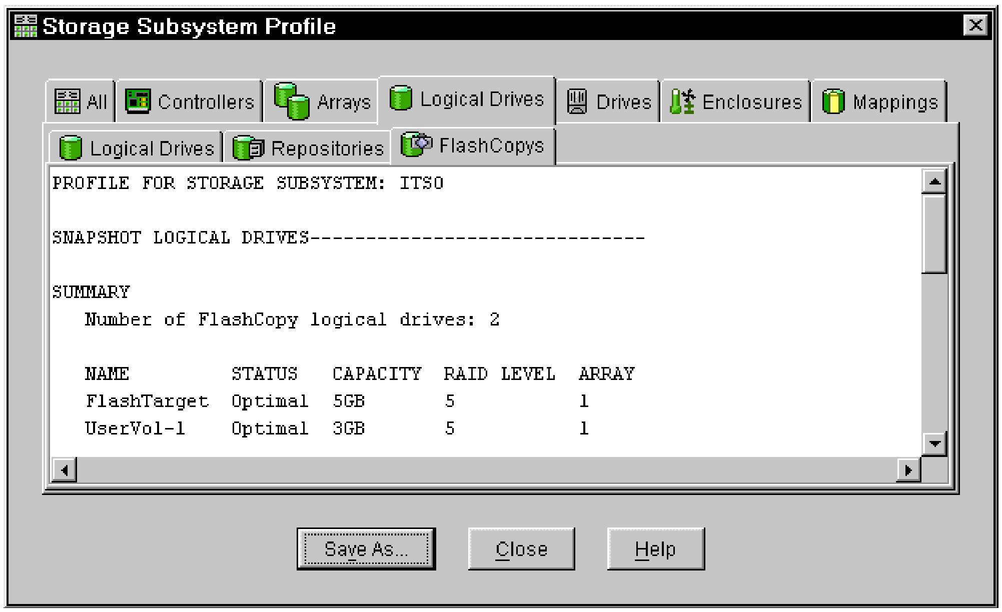

Figure 4-47 Storage Subsystem Profile showing logical drives

4.6 Maintenance, tuning, diagnostics, and troubleshooting

The previous sections describe how to install the drivers and the host software, as well as how to set up a basic configuration of the FAStT Storage Server. This section takes you through the different maintenance and tuning options available, for example, to tune the storage server for performance.

4.6.1 Component properties

All information received in the profile from the previous section is gathered from the various components of the storage subsystems as controllers, disk drives, enclosures. The information collected is also available for each component. This may be useful to track or locate a failure in a component.

Highlight the component you want to examine, right-click, and choose **Properties**. The properties appear for this component, for example, for one of the attached drive enclosures (see Figure 4-48).

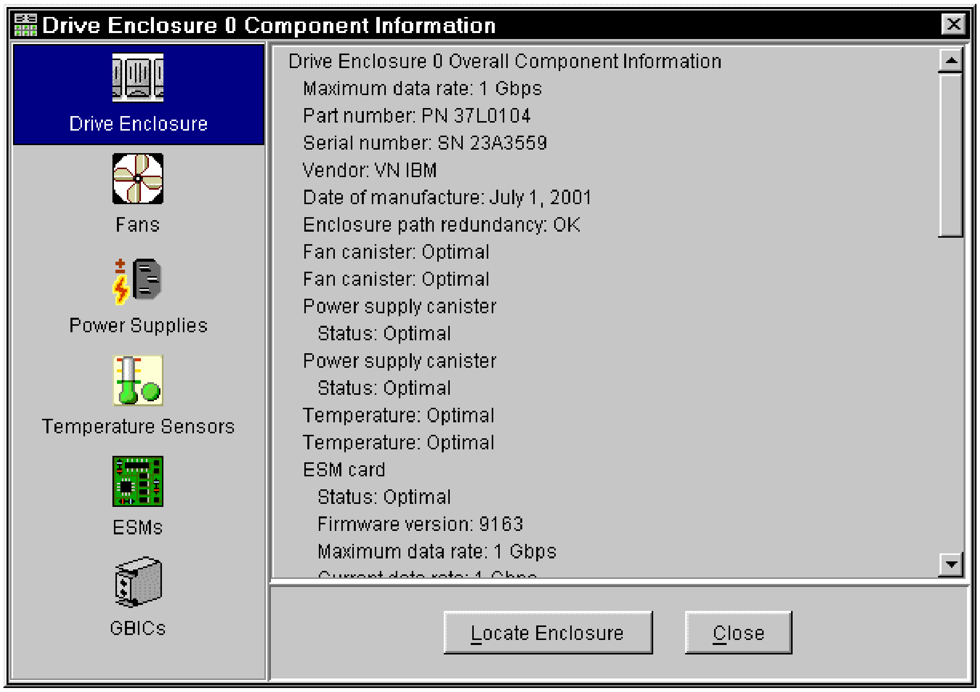

Figure 4-48 Properties of an attached drive enclosure

You can easily locate all components with the help of the Storage Manager software. Simply highlight the component, for example, an array. Then right-click and select **Locate**. The Locate Array dialog (Figure 4-49) opens. All drives of this array start flashing the yellow LED. When you finish locating the components, click **OK** to stop the flashing of the LEDs.

Figure 4-49 Locating a disk array

Sometimes it is useful to see all components associated with one object. For example, you want to know which arrays and logical and physical drives are managed by Controller A. Right-click **Controller A** and choose **View Associated Components**. In the window, you see a list of all disk drives contained in arrays managed by this controller. You also see the logical drives and arrays as shown in Figure 4-50.

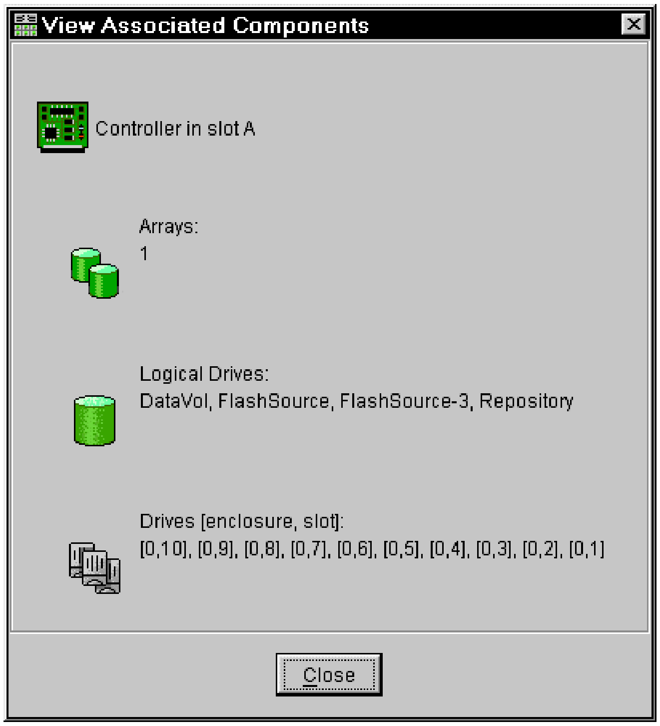

Figure 4-50 Associated components of Controller A

4.6.2 Modification priority

The modification priority defines how much processing time is used for operations modifying the logical drive relative to the system performance. Operations that cause a logical drive modification are:

- Initializing a logical drive
- Reconstructing after a disk failure
- Copying back from a hot spare drive
- Changing the segment size of a logical drive
- Dynamic logical drive expansion
- Adding free capacity to an array
- Defragmenting an array
- Changing the RAID level of an array

If the logical drive contains critical data, you may prefer a high modification priority to keep the time of a critical state, for example after losing a disk, as short as possible, even if this affects the system performance during the modification process.

To change the modification priority, click **Logical Drive-> Change-> Modification Priority** and enable the logical drives to be included in the media scan (see Figure 4-51).

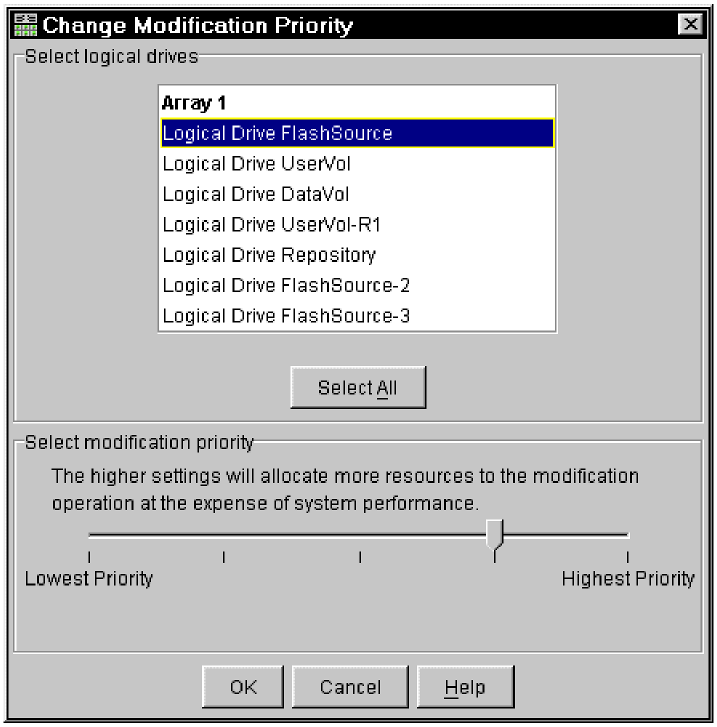

Figure 4-51 Modification priority for a logical drive

If a logical drive modification is in progress, a status bar appears at the bottom of the dialog.

4.6.3 Cache settings

If you want to tune performance, it may be necessary to change the settings for the use of the cache. There are two global settings that affect all logical drives on the FAStT Storage Server. The other settings that are available affect only a single logical drive. This allows you a very granular optimization of the FAStT Storage Server.

Whenever you change cache settings, you should monitor the FAStT Storage Server with the built-in performance monitor to see if changes really improve performance or degrade it.

Global cache settings

The global settings for the cache usage can be modified to reflect special needs. Usually there is no need to change these values, because they are efficient in most cases. However, if you want to tune performance, it is more efficient to change the values for individual logical drives. This allows a much more granular optimization of the storage subsystem.

In the Subsystem Management window, click **Storage Subsystem-> Cache Settings**. The Change Cache Settings dialog (Figure 4-52) shows you the global cache settings. The values for Start and Stop flushing are discussed in 8.1.2, “Cache parameters” on page 252, as well as the cache block size value. The changes occur immediately.

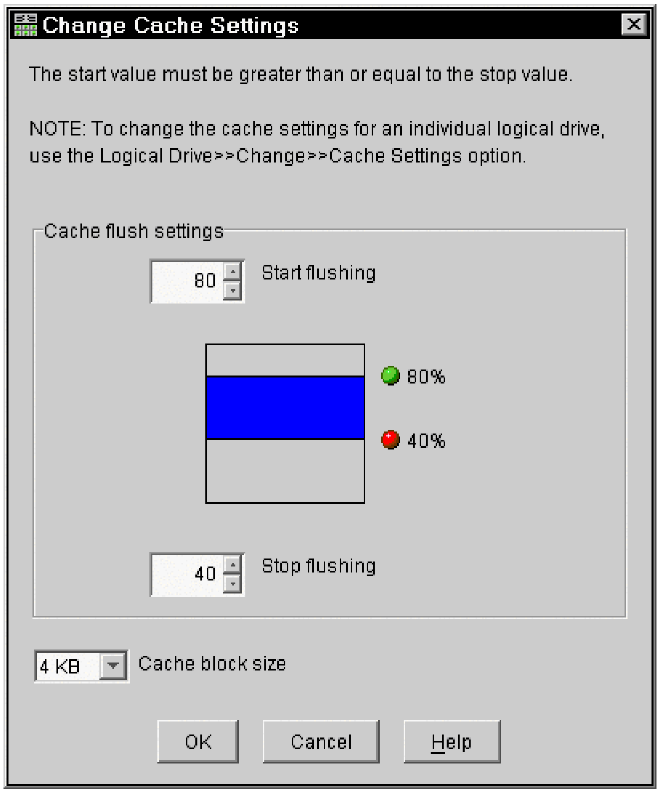

Figure 4-52 Global Cache Settings

In our example, we already changed the cache flush settings.

Logical drive cache settings

For each logical drive, there are several parameters that you can change. The parameters may be changed for all drives at once or for each drive separately.

The default settings are read and write cache for all logical drives, with cache mirroring to the alternate controller for all write data. The write cache is only used if the battery for the controller is fully charged. Read ahead is not normally used on the logical drives.

The *Read Caching* parameter can be safely enabled without risking data loss. There are only rare conditions where it is useful to disable this parameter, which then provides more cache for the other logical drives.

The *Write Caching* parameter allows the storage subsystem to cache write data instead of writing it directly to the disks. This can improve performance significantly especially for environments with random writes such as databases. For sequential writes, the performance gain varies with the size of the data written. If the logical drive is only used for read access, it may improve overall performance to disable the write cache for this logical drive. Then no cache memory is reserved for this logical drive.

By default, a write cache is always mirrored to the other controller to ensure proper contents, even if the logical drive moves to the other controller. Otherwise the data of the logical drive can be corrupted if the logical drive is shifted to the other controller and the cache still contains unwritten data. If you turn off this parameter, you risk data loss in the case of a controller failover, which may also be caused by a path failure in your fabric.

The cache of the FAStT Storage Server is protected, by a battery, against power loss. If the batteries are not fully charged, for example, just after powering on, the controllers automatically disable the write cache. If you enable the parameter, the write cache is used, even if no battery backup is available, resulting in a higher risk of data loss.

The read-ahead multiplier defines the number of data blocks that should be read ahead. The default value of zero does not read ahead any data. Depending on the usage pattern for this logical drive, for example sequential reads, it may increase performance to change to a higher value. A value that is too high can cause an overall performance decrease because the cache is filled with read ahead data that is never used. Use the performance monitor to watch the cache hit rate for this logical drive to find a proper value. See also 8.1, “Performance monitoring and tuning” on page 250.

To change any settings, right-click a logical drive and choose **Change-> Cache Settings**. The Change Cache Settings dialog (Figure 4-53) allows you to change the parameters that we described.

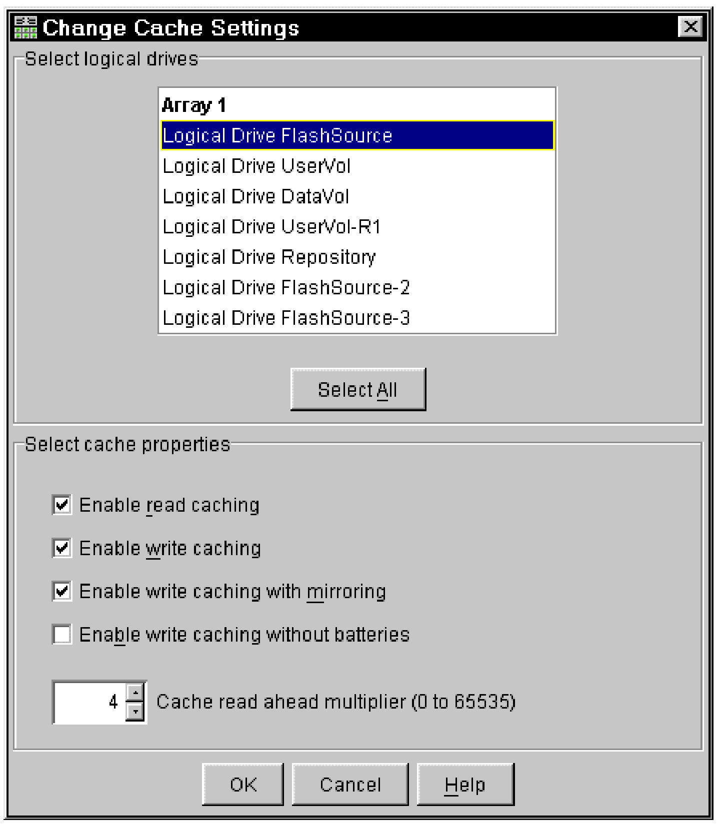

Figure 4-53 Logical drive cache settings

4.6.4 Media scanning

Media scan enables the background media scan, which can provide a higher availability of the data. It checks, as a background operation, the physical disks for defects by reading the raw data from the disk and writing it back. This detects possible problems caused by bad sectors of the physical disks before they disrupt normal data reads or writes.

Depending on the global media scan rate, this can impact performance but improve data integrity. Before you can enable the media scan for a logical drive, you must enable the global media scan. Click **Storage Subsystem-> Change-> Media Scan Settings** and enable the media scan on the FAStT Storage Server. On the Change Media Scan Settings dialog (Figure 4-54), you can also define the duration of one complete scan of all selected logical drives.

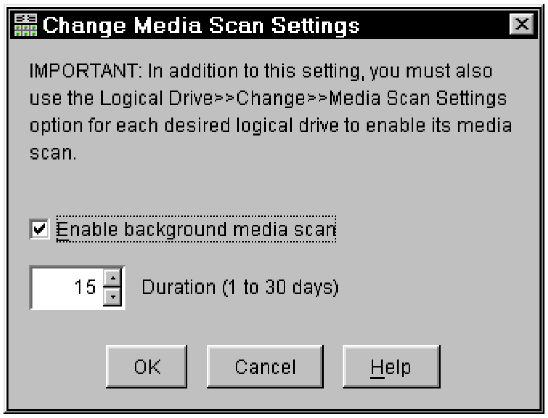

Figure 4-54 Enabling media scan

If you enabled the media scan globally, define which logical drive to include in the media scan. Right-click a logical drive and choose **Change-> Media Scan Settings**. In the Change Media Scan Settings window (Figure 4-55), select the logical drives to include in the media scan.

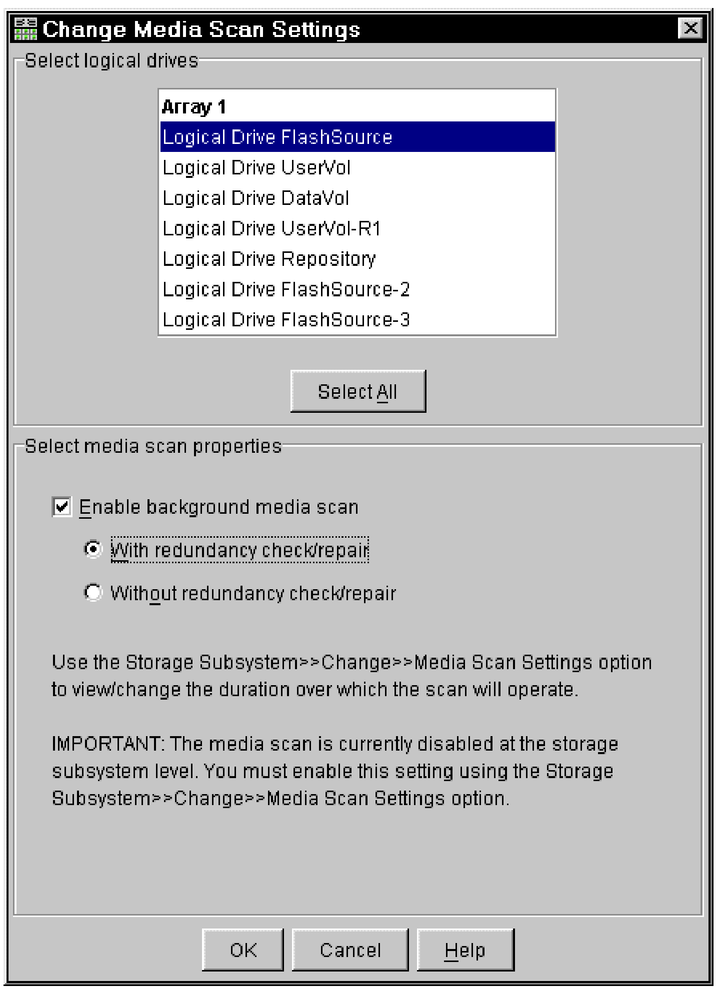

Figure 4-55 Enabling a logical drive media scan

If you also enable the redundancy check/repair option, the media scan also checks redundancy information on a logical drive with RAID 3 or 5 or it compares the data blocks of RAID 1 mirrored devices.

4.6.5 Loading and saving the configuration

If you want to set up multiple storage subsystems with the same configuration, you can save the configuration to a file and apply it to another subsystem. The subsystem must have the same hardware layout, number of enclosures and drives, and drive capacities.

The saved configuration only includes the array and logical drive configuration, the name of the subsystem, its cache settings, and media scan rate. It does not include any information regarding storage partitioning. This information is included only in the storage subsystem profile.

All information is stored in a file that contains a script for the script editor. To save the configuration of the subsystem, open the Subsystem Management window, highlight the subsystem, and click **Storage Subsystem-> Configuration-> Save**. The configuration is saved to a file.

To load the configuration on a subsystem, open the Enterprise Management window and select the subsystem. Click **Tools-> Load Configuration** from the menu. Point to the file containing the configuration and load it. The script editor and a warning message appear. To load the configuration onto the FAStT Storage Server, choose **Execute**. You may also edit the script before executing it.

Attention: This procedure replaces any configuration on the storage subsystem. All data stored on the FAStT Storage Server will be lost because all logical drives will be initialized.

The procedure can take a long time, depending on the number of arrays and logical drives defined. When the procedure finishes, the subsystem contains the same configuration as the source subsystem.

4.6.6 The major event log

The FAStT Storage Server logs all major events in the storage subsystem to a reserved part on the attached disks. Even when no server is attached, or attached but down, the FAStT Storage Server still monitors itself and log errors.

To view the log file, open the Subsystem Management window and click **View-> Event Log**. The event log is transferred from the storage subsystem and displayed. If you choose to view the details, you can directly see the exact error description and location. By default, only the last 50 critical entries are displayed as shown in Figure 4-56. In the top right corner, you can select the number of entries to be displayed. On the left side, you can choose to display non-critical entries.

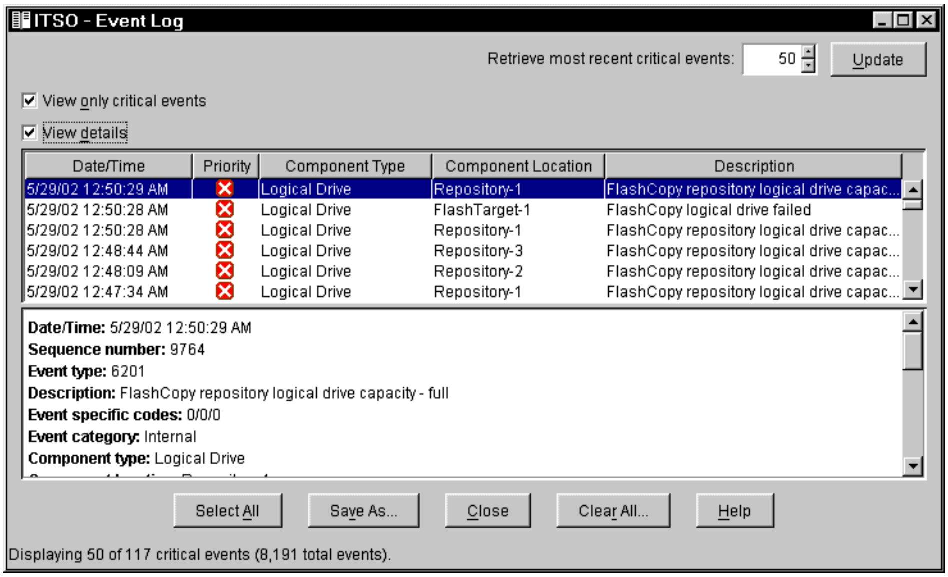

Figure 4-56 Major event log

You can save the entries from the event log. Select the entries you want to save or click the **Select All** button. Then save the log file locally on the host system.

After the initial setup or an error situation, it may be useful to clear the complete event log to have a fixed starting point.

The log can store 8192 entries and then overwrite the oldest entries as needed. Since an intermittent failure can cause a large number of error messages, it is useful to save the log as soon as the failure is detected to ensure that no event is lost.

4.6.7 Changing ownership of a logical drive

Each logical drive has a preferred controller of ownership. This controller normally handles all I/O requests for this particular logical drive. The alternate controller only takes over and handles the I/O requests in the case of a failure along the I/O path, for example a defect host bus adapter or switch. Normally the logical drive’s ownership is alternated between the two controllers while you define them.

All heavily stressed logical drives may reside on only one controller and the other one handles only a small amount of all I/O requests. To balance the workload between the controllers, you can change the preferred ownership of a logical drive to the other controller, and normally the storage subsystem is balanced better regarding the workload. To change the preferred ownership of a logical drive, it must reside on this one. You cannot change the ownership if the logical drive is handled by its alternate controller.

Important: Be sure that the operating system using the logical drive uses a multipath I/O driver. Otherwise, it loses access to the logical drive.

Highlight the logical drive and click **Logical Drive-> Change-> Ownership/Preferred Path**. Then select and the controller to which you want the logical drive to move. Depending on the current workload, the operation can take while to finish.

4.6.8 Maintaining arrays

Even though the FAStT Storage Server always tries to optimize the layout of the disk arrays, you may want to change some settings to optimize the disk usage or the performance. Three options are available that each affect the entire array and all logical drives residing on it.

Changing RAID level

The RAID level of the array is normally defined when you create the first logical drive on this array. There are several points to watch for which RAID level is optimal for your usage pattern. See also 8.1.5, “RAID levels” on page 254.

The FAStT Storage Server allows you to dynamically change the RAID level of an array, so you can change the RAID level during your normal operation. You should consider that you can cause a huge amount of additional load during the process, which may affect the overall performance. We recommend that you carry out this operation during periods of minimal I/O activity.

There are no restrictions as to which RAID level can be migrated. Depending on the RAID level you want to migrate to, you need enough free space in the array to perform this operation. If there is not enough free space, you receive an error message. In this case, add more free capacity first as explained in “Expanding an array” on page 123.

If the operation frees up space on the array, it is available as free space to either define a new logical drive or expand already existing ones.

To change the RAID level of an array, highlight the array, right-click and click **Change-> RAID Level** and the desired RAID level. The operation starts and cannot be stopped once it is started. The data remains accessible during this operation, which can take a long time.

Expanding an array

To increase the size of an array, you add new physical drives to it. This process can be done concurrently with access to the data on the logical drives. After the migration process, the new drives are included in the array and provide new free space within the array. This free space can then be used to define new logical drives or expand existing ones.

To add new drives to an array, highlight the array, right-click, and choose **Add free Capacity (Drives)**. On the Add Drives window (Figure 4-57), you need at least one unassigned drive that can be added to the array.

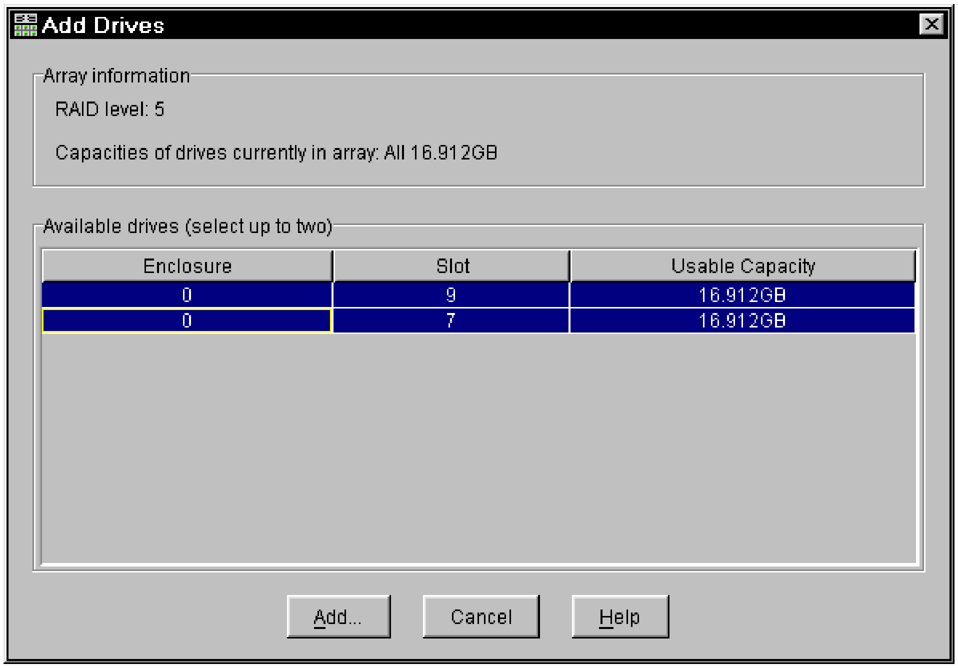

Figure 4-57 Adding new drives to an array

For RAID levels 3 and 5, you must select at least one drive. For RAID levels 1 and 10, you must choose an even number of drives.

Once the procedure is started, you cannot stop it. Because the subsystem needs to redistribute the data contained in the array to all drives including the new ones, there is a performance impact during this operation. However, the logical drives of the array remain available to the host systems.

Defragmenting an array

A logical drive can be deleted anytime to free the space in the array. The free space may be fragmented within the array in different free space nodes.

Because new logical drives cannot spread across several free space nodes, the logical drive size is limited to the greatest free space node available, even if there is more free space in the logical drive. Compare this with array 1 in Figure 4-58.

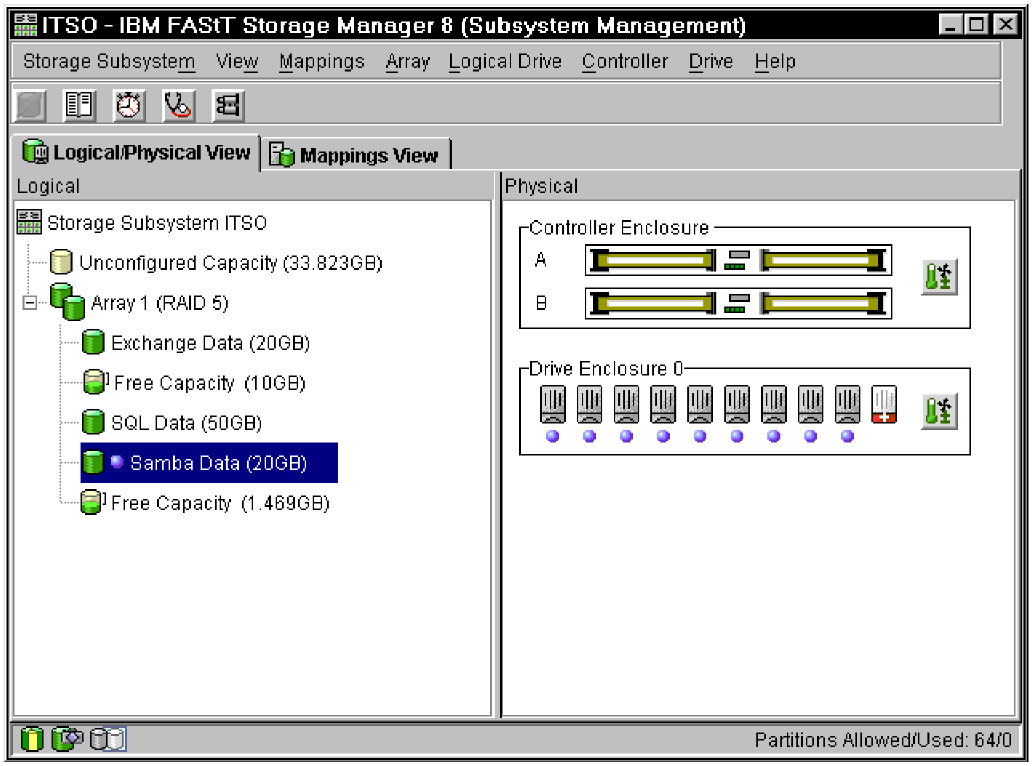

Figure 4-58 Defragmented array

The array needs to be defragmented first to consolidate all free space nodes to one free space node for the array. Then a new logical drives can use the whole available free space.

Open the Subsystem Management window. Highlight the array to defragment and click **Array-> Defragment** to start the procedure. The defragmentation can run concurrently with normal I/O, but it impacts performance because data of the logical drives must be moved within the array. Depending on the array configuration, this process continues to run for a long period of time. Once the procedure is started, it cannot be stopped again. During this time, no configuration changes can be performed on the array.

The defragmentation done on the FAStT Storage Server only applies to the free space nodes on the array. It is not connected to a defragmentation of the file system used by the host operating systems in any way.

4.6.9 Maintaining logical drives

There are several properties of the logical drive that you can change. This section explains the properties.

Changing the segment size

The segment size of a logical drive can be changed to tune the drive for performance. You can find an explanation about segment size in 8.1.4, “Segment size” on page 254. Normally a small segment size is used for databases, normal sizes for file server, and large segment sizes for multimedia applications.

With the segment size, you can influence the throughput, because the throughput is the number of I/O operations per second multiplied by the segment size. If you increase the segment size, you gain more throughput.

To change the segment size of a logical drive, highlight the drive and click **Logical Drive-> Change-> Segment Size**. The available segment sizes are:

- 8 KB
- 16 KB
- 32 KB
- 64 KB
- 128 KB
- 256 KB

You can only change the segment size by one in one step. For example, if you want to change the segment size from 16 KB to 128 KB, this requires three consecutive changes of the segment size.

Once this process started, it cannot be stopped, but all data on the logical drive is available throughout the operation. As with all operations on logical drives or arrays, the overall performance may be affected.

Dynamic logical drive expansion

With FAStT Storage Manager v8, a new feature is introduced – the dynamic logical drive expansion. Now you can expand an array, define a new logical drive in this array, and enlarge a logical drive. You can run this process during normal operation since the logical drive remains accessible all the time.

Because the logical drive appears as a disk device in the host operating system and a file system is used on this disk, some restrictions apply depending on the type of operating system and file system you are using. See also 3.2.1, “Arrays and logical drives” on page 43.

To enlarge a logical drive, free space must be available, either as free capacity in the array or as unconfigured capacity in the form of free disk drives. The disk doesn’t need to be assigned to the array. The dynamic logical drive expansion invokes the array expansion if needed. If the array needs to be expanded, a maximum of two disk drives can be added to the array at any time, but the outlined procedure can be repeated.

Highlight the logical drive to be expanded and click **Logical Drive-> Increase Capacity**. In the Increase Logical Drive Capacity dialog (Figure 4-59), you can enter the amount of space by which the logical drive will be enlarged.

In the top part, you see the current size of the logical drive and available free capacity, either free space in the array or newly added drives as in our example. If no free space is available in the array, click **Add Drives** before you continue to enlarge the logical drive. This takes you to the same dialog as described in “Expanding an array” on page 123, but you can only add two drives at a time here, because you are combining the two operations.

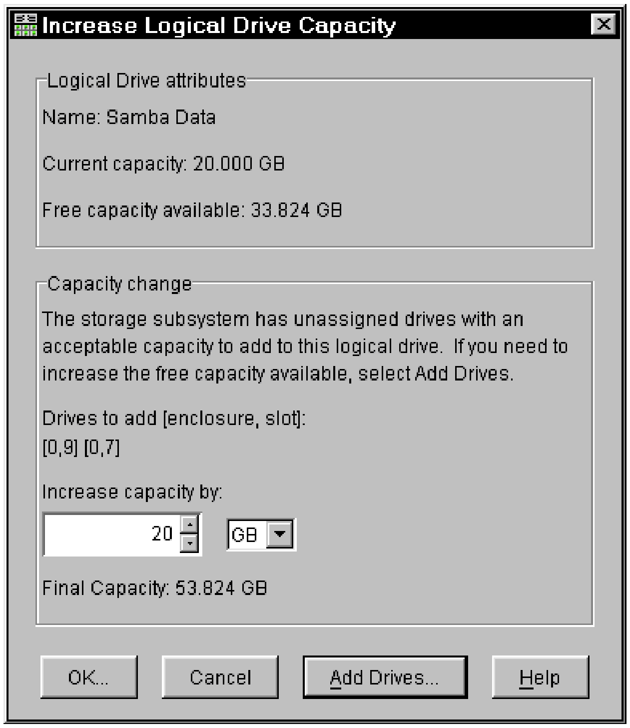

Figure 4-59 Dynamic Logical Drive Expansion

Click **OK**. A warning message appears indicating that this operation cannot be stopped once it is started and that it may take a long time to finish. However, the data on the selected logical drive and all other logical drives on this array (if you added new drives) remains accessible during this time. As with all operations requiring a redistribution of the data on the physical disks, the procedure may affect the performance.

4.6.10 Handling premium features

Some features available for the FAStT Storage Server are not part of the base firmware. They have to be enabled with a premium feature. The storage partitioning premium feature is activated by default. This allows you to attach multiple heterogeneous hosts to the same FAStT Storage Server. Such functions as FlashCopy and Remote Volume Mirroring must be enabled first before they can be used.

An error condition, called *Out of Compliance*, may occur and make it necessary to disable the premium feature and re-enable it again. See 8.4.2, “FAStT Client cannot connect to the FAStT Storage Server” on page 263.

Listing premium features

To list the currently enabled premium features, click **Storage Subsystem-> Premium Features-> List** in the Subsystem Management window. A dialog (Figure 4-60) appears that shows you all enabled premium features.

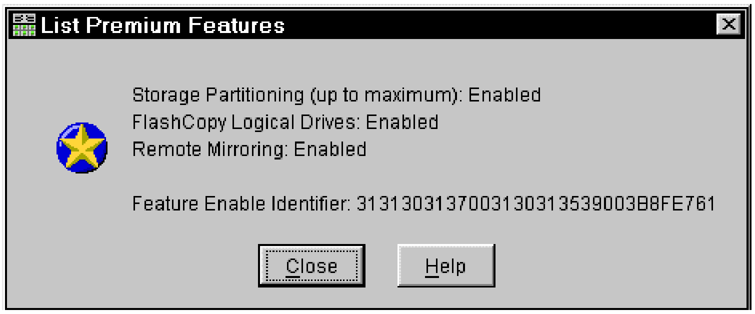

Figure 4-60 List of enabled premium features

In our example, the storage partitioning for the maximum number of hosts is enabled, as well as the FlashCopy and Remote Volume Mirroring feature. The long number, called the *Feature Enable Identifier*, is necessary if you need a new feature key, since this key is unique for each FAStT Storage Server.

If you bought the FlashCopy or Remote Volume Mirroring feature, you receive a CD-ROM that contains a feature key generator. If you invoke the key generator, you need to enter this Feature Enable Identifier. The key file created by the key generator is then used to enable the feature as described next.

Enabling a premium feature

To enable a premium feature, you need the appropriate key file, provided either by IBM or generated yourself. When you have the key, click **Storage Subsystem-> Premium Features-> Enable** in the Subsystem Management window. In the dialog window, point to the location where the key file is stored. You need to confirm whether to enable the premium feature selected as shown in Figure 4-61. If the Feature Enable Identifier doesn't match the FAStT Storage Server, the key is not installed.

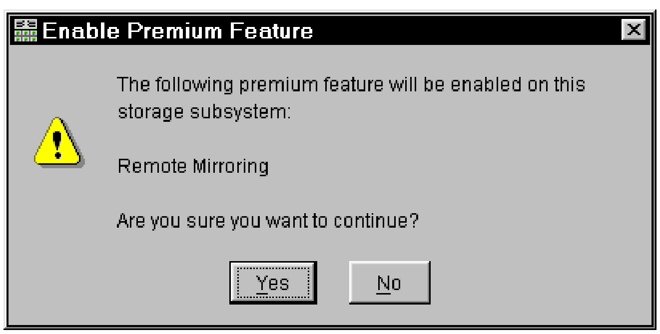

Figure 4-61 Enabling a premium feature

The change happens immediately, and you can use the new functions directly.

Disabling a premium feature

To disable a premium feature, click **Storage Subsystem-> Premium Features-> List** in the Subsystem Management window. Choose the feature you want to disable from the list and confirm.

Keep in mind that the change happens immediately. If you use storage partitioning and disable the premium feature, then you cannot create new partitions. Any existing partitions remain.

4.6.11 Network setup of the controllers

With version 8.2 of the FAStT Storage Manager and the assigned firmware version of 5.2xxx for the controllers, two major changes are incorporated into the network driver. The default behavior is to change an IP address.

Now, by default, FAStT tries to use BOOTP to request an IP address. If no BOOTP server can be contacted, the controllers fall back to the fixed IP addresses. These fixed addresses, by default, are:

- Controller A: 192.168.128.101
- Controller B: 192.168.128.102

In previous versions of the controller firmware, there was no fall back to fixed IP addresses. There are two other possible configuration settings:

- Use fixed IP addresses
- Use BOOTP and not fall back to fixed IP addresses if a BOOTP server is unavailable

The other major change is that the network driver of the FAStT controllers is reset every time a configuration change is made or whenever a physical connection is made to the Ethernet ports. So there is no longer any need to power cycle FAStT after a configuration change. Resetting the network driver by inserting the network cable is also useful, for example when the FAStT Storage Server cannot contact the BOOTP server because of a network problem. After fixing this problem, you only need to unplug and plug the cable back into the FAStT Storage Server. It assigns an IP address through BOOTP to the controllers. In older versions of the firmware, resetting the controllers was necessary.

To use the network ports of the controllers, you need to attach both controllers to an Ethernet switch or hub. The built-in Ethernet controller supports either 100 Mbps or 10 Mbps.

To manage storage subsystems through a firewall, configure the firewall to open port 2463 for TCP data.

To change the default network setting (BOOTP with fallback to a fixed IP address), you need a serial connection to the controllers in the FAStT Storage Server.

> **Attention:** Follow the procedure outlined here exactly as it is presented, because some commands that can be issued from the serial console can cause data loss.

1. Connect to the FAStT Storage Server with a null modem cable to the serial port of your system. For the serial connection, choose the correct port and the following settings:
   - 19200 Baud
   - 8 Data Bits
   - 1 Stop Bit
   - No Parity
   - Xon/Xoff Flow Control
2. Send a break signal to the controller. This varies depending on the terminal emulation. For most terminal emulations, as HyperTerm, which is included in the Microsoft Windows products, press Ctrl+Break.

   There is a known issue with early version of HyperTerm shipped with Windows NT 4.0. You need at least version 3 to work properly. To download the latest version, go to:

   `http://www.hilgraeve.com`
3. If you only receive unreadable characters, press Ctrl+Break again, until the following message appears:

```
Press <SPACE> for baud rate within 5 seconds.
```

4. Press the Space bar to ensure the correct baud rate setting. If the baud rate was set, a confirmation appears.
5. Press Ctrl+Break to log on to the controller. The following message appears:

```
Press within 5 seconds: <ESC> for SHELL, <BREAK> for baud rate.
```

6. Press the Esc key to access the controller shell. The password you are prompted for is `infiniti`.
7. Run the **`netCfgShow`** command to see the current network configuration. The controller dumps a list similar to the output shown in Example 4-5.

Example 4-5 Current network settings

```
-> netCfgShow

==== NETWORK CONFIGURATION: ALL INTERFACES ====
Network Init Flags    :  0x00
Network Mgmt Timeout  :  30
Startup Script        :
Shell Password        :

==== NETWORK CONFIGURATION: dse0 ====
Interface Name        :  dse0
My MAC Address        :  00:a0:b8:0c:d4:66
My Host Name          :  flute17a
My IP Address         :  9.11.200.155
Server Host Name      :  host
Server IP Address     :  0.0.0.0
Gateway IP Address    :  9.11.200.1
Subnet Mask           :  255.255.252.0
User Name             :  guest
User Password         :
NFS Root Path         :
NFS Group ID Number   :  0
NFS User ID Number    :  0
value = 0 = 0x0
->
```

8. To change the above values, enter the **`netCfgSet`** command. You are asked for each entry to keep, clear, or change the value. See Example 4-6.

Example 4-6 Changing the network settings

```
-> netCfgSet


    '.' = clear field;  '-' = to previous field;
    '+' = next interface;  ^D = quit (keep changes)

==== NETWORK CONFIGURATION: ALL INTERFACES ====
Network Init Flags    :  0x00
Network Mgmt Timeout  :  30
Network Route #1      :  dest=0.0.0.0
RAIDMGR Server #1     :  0.0.0.0
Network Manager #1    :  0.0.0.0
```

```
Startup Script       :
Shell Password       :

==== NETWORK CONFIGURATION: dse0 ====
My MAC Address       :  00:a0:b8:0c:d4:66
My Host Name         :  flute17a
My IP Address        :  9.11.200.155           0.0.0.0
Server Host Name     :  host
Server IP Address    :  0.0.0.0
Gateway IP Address   :  9.11.200.1
Subnet Mask          :  255.255.252.0
User Name            :  guest
User Password        :
NFS Root Path        :
NFS Group ID Number  :  0
NFS User ID Number   :  0

Network Configuration successfully written to NVSRAM.
value = 0 = 0x0
->
```

The *Network Init Flag* defines the mode the controllers are using for the network setup:

– `Network Init Flag : 0x00` is the default. Try BOOTP first. If this fails, use the fixed IP address defined in `My IP Address`, which defaults to 192.168.128.10x.

– `Network Init Flag : 0x01` uses only the fixed IP address defined in `My IP Address`.

– `Network Init Flag : 0x02` uses only BOOTP without fallback to a fixed IP address.

9. After you assign a fixed IP address to Controller A, disconnect from Controller A and repeat the procedure for Controller B. Remember to assign a different IP address.
10. Because the configuration changed, the network driver is reset and uses the new network configuration.
11. If you want to use a BOOTP server to assign the controllers a specific IP address, you need to define the MAC addresses of the controllers in your BOOTP server. The MAC addresses of the controllers are printed on each controller. On a FAStT500 or FAStT700 controller, you can find the label with the MAC address just beneath the latch on the front. On FAStT200, you can find the label on the back.

4.6.12 Resetting the controllers

This section explains how to reset the controllers in the FAStT Storage Server. You lose all data. This includes configuration data (as well as the array and logical drive configuration) and data stored in the logical drive.

Back up any data before you continue the procedure.

There are two ways to reset the controllers in the FAStT Storage Server. The preferred way is to reset the controllers through the Storage Manager software, either directly or host-agent attached. Normal problem determination does not require this procedure to be done.

Attention: Both procedures cause a loss of all data, including array and logical drive configurations. Back up all data before you proceed.

Resetting with FAStT Storage Manager

The FAStT Storage Server can be completely reset through the client software.

1. Open the Subsystem Management window, highlight the storage subsystem, and click **Storage Subsystem-> Configuration-> Reset**.
2. Before you can confirm the start of the procedure, you must enter `yes` in the dialog as an extra confirmation since you lose all data.
3. Once the storage subsystem is reset, rescan for the device in the Enterprise Management window.
4. If you are managing through an in-band connection, it may be necessary to reboot the system running the host agent to recognize the storage subsystem.

Resetting through the serial connection

You need a serial connection to the controllers in the FAStT Storage Server. Because you have to reboot the FAStT Storage Server during this procedure, you have to stop all I/O activity.

1. Connect to the FAStT Storage Server with a null modem cable to the serial port of your system. For the serial connection, choose the correct port and the following settings:
   - 19200 Baud
   - 8 Data Bits
   - 1 Stop Bit
   - No Parity
   - Xon/Xoff Flow Control
2. Send a break signal to the controller. This varies depending on the terminal emulation. For most terminal emulations, such as HyperTerm, which is included in the Microsoft Windows products, press Ctrl+Break.

   There is a known issue with early version of HyperTerm, you need at least version 3 to work properly. To download the latest version, go to:

   http://www.hilgraeve.com
3. If you only receive unreadable characters, press Ctrl+Break again, until the following message appears:

   ```
   Press <SPACE> for baud rate within 5 seconds.
   ```
4. Press the Space bar to ensure the correct baud rate setting. If the baud rate was set, a confirmation appears. Press Ctrl+Break to log on to the controller. The following message appears:

   ```
   Press within 5 seconds: <ESC> for SHELL, <BREAK> for baud rate.
   ```
5. Press the Esc key to access the shell of the controller. The password you are prompted for is `infiniti`.
6. At the shell, enter **`sysWipe`**. This resets the controller completely. The command is executed in the background. Wait until the procedure finishes and a message appears on the shell. Disconnect the serial cable.
7. If a second controller is installed, connect to the serial port of Controller B and repeat the procedure to ensure that you reset both controllers.
8. After this procedure, reboot the entire FAStT Storage Server by power cycling the controller unit. Be sure to switch off both power supplies.
9. Because all configuration data is destroyed on the storage subsystem, you have to initiate a discovery in the Enterprise Management window. The reset FAStT Storage Server appears as a new storage subsystem.

4.7 The FAStT Management Suite Java

The FAStT Management Suite Java is a tool used to manage the IBM FAStT Host Bus Adapters in the Intel environment. It is particularly useful for obtaining the WWN without having to reboot.

4.7.1 Overview of the FAStT MSJ

The FAStT MSJ consists of two components:

- An agent that is installed on all servers that have one or more host bus adapters
- A GUI that can be installed on any machine and that is used to manage the host bus adapters installed in the servers

The agent is supported on the following operating systems:

- Windows NT 4.0 with service pack 6a
- Windows 2000 with service pack 2
- Red Hat Linux 7.1 upgraded to Kernel 2.4.9-21
- Novell NetWare 5.1 with NW5SP2

The GUI is supported on the following operating systems:

- Windows NT with service pack 6 a
- Windows 2000 with service pack 2
- Red Hat Linux 7.1 upgraded to Kernel 2.4.9-21

On a Windows or NetWare system, FAStT MSJ supports the following host bus adapters:

- IBM Netfinity Fibre Channel Host Bus Adapters (2100)
- IBM FAStT Fibre Channel Host Bus Adapters (2200)
- IBM FAStT FC-2 Host Bus Adapter (2300)

On a Linux system, FAStT MSJ only supports the following adapters:

- IBM FAStT Fibre Channel Host Bus Adapters (2200)
- IBM FAStT FC-2 Host Bus Adapter (2300)

Some of the supported features include:

- Timely and accurate detection of I/O failures
- Local and remote management of adapters
- Performance statistics
- Central control point in a network environment
- Diagnostics and utilities
- Device\Port Configuration for failover capability (Linux only)

FAStT MSJ is downloadable from the IBM PC Support Web site. The latest version, at the time this redbook was written, is version 2.0, release 32, for Windows NT4, Windows 2000, NetWare, and Linux. It is available from the following Web site:

http://www.pc.ibm.com/qtechinfo/MIGR-39194.html

4.7.2 Installing the FAStT MSJ

This section explains how to install FAStT MSJ. FAStT MSJ supports the following two configurations:

- **Standalone**

  If you want to manage the Host Bus Adapters in a stand-alone system, install both the GUI and Agent on the local system.

- **Networked**

  If you want to manage the host bus adapters on a remote system, install only the agent on the remote system and the GUI on a management system. A management system may also contain host bus adapters.

Follow these steps to install FAStT MSJ in a Windows or NetWare environment:

1. Unpack the downloaded package file to a path that you specify by running the self-extracting executable.
2. Install the FAStT MSJ by running the installation file unpacked in step 1.
3. The InstallAnywhere status window opens while the application prepares to install. Then the FAStT MSJ splash screen appears. When the Introduction window opens, click **Next** to continue. Accept the licence conditions and click **Next** again.
4. The Choose Product Features window opens as shown in Figure 4-62. Select the options you want to install on this host and click **Next**.

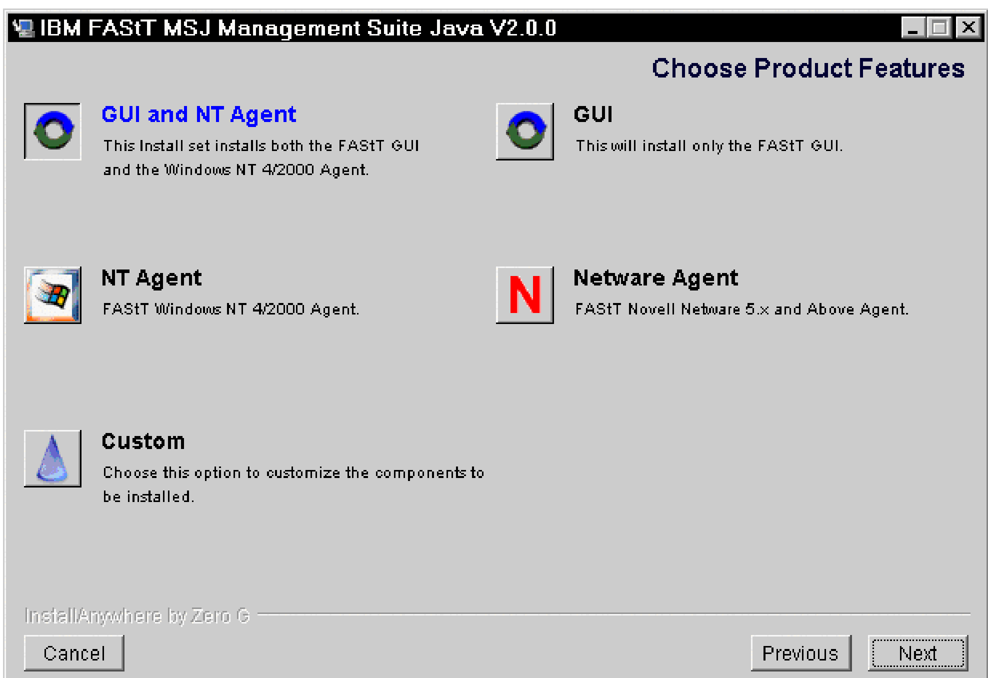

Figure 4-62 FAStT MSJ installer

5. The Information window opens. Click **Next**.
6. The Choose Install Folder windows appears. Specify your installation path and click **Next**.
7. On Windows, the Select Shortcut Profile window opens. Select **All User Profiles** and click **Next**.

8. On Windows, the Create Desktop Icon window opens. Select the **Create desktop icon** check box and click **Install**.
9. The Installing window appears and shows the ongoing status of the installation process. When everything is installed, the Install Complete window appears.
10. Click **Done** to end the installation.

For information on installing on Linux, see 4.1.2, “Linux (RedHat 7.1 and 7.2, SuSE 7.3, and TurboLinux 7.0)” on page 70.

4.7.3 Using the FAStT MSJ GUI

The FAStT Management Suite Java GUI is used to manage the host bus adapters in servers running the FAStT MSJ agent. With this tool, you can view and make changes to the configuration of host bus adapters, perform diagnostic tests, and check the performance of the host bus adapters. This section briefly explains how you can use this software. Refer to the product documentation and online help for more information.

Follow these steps for a brief tour of the FAStT MSJ GUI:

1. Start the GUI by double-clicking the **FAStT MSJ shortcut** located on the desktop.
2. The IBM FAStT Management Suite Java - HBA View window opens as shown in Figure 4-63.

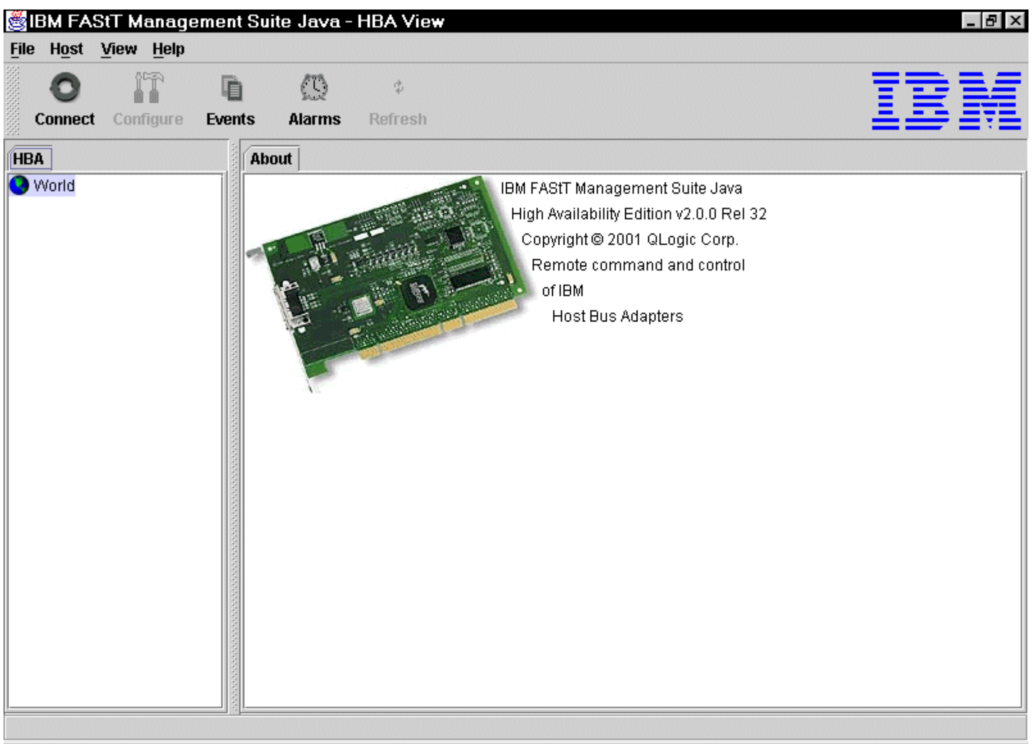

Figure 4-63 FAStT HBA view

This is the main window in the GUI. To connect to a server agent, click the **Connect** icon and the Connect to Host window opens as shown in Figure 4-64.

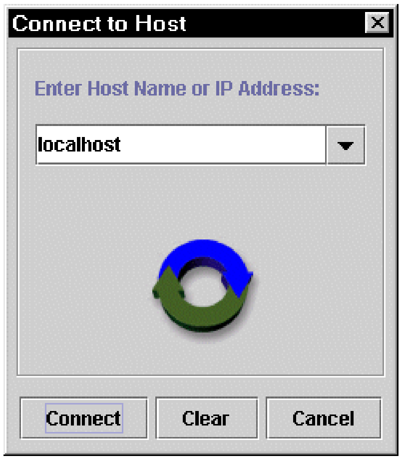

Figure 4-64   FAStT MSJ Connect to Host

3. Enter the name or the IP address of the server to which you want to connect, and click **Connect**.
4. The server is added to the left-hand side pane of the main window on the **HBA** tab. The host bus adapters that are installed in the server appear below the tab.
5. Click **Host aae19** to see the two tabs displayed in the right-hand pane (Figure 4-65). They are:
   – **Information**: Contains basic information about the currently connected server
   – **Security**: Contains security settings for the connected agent

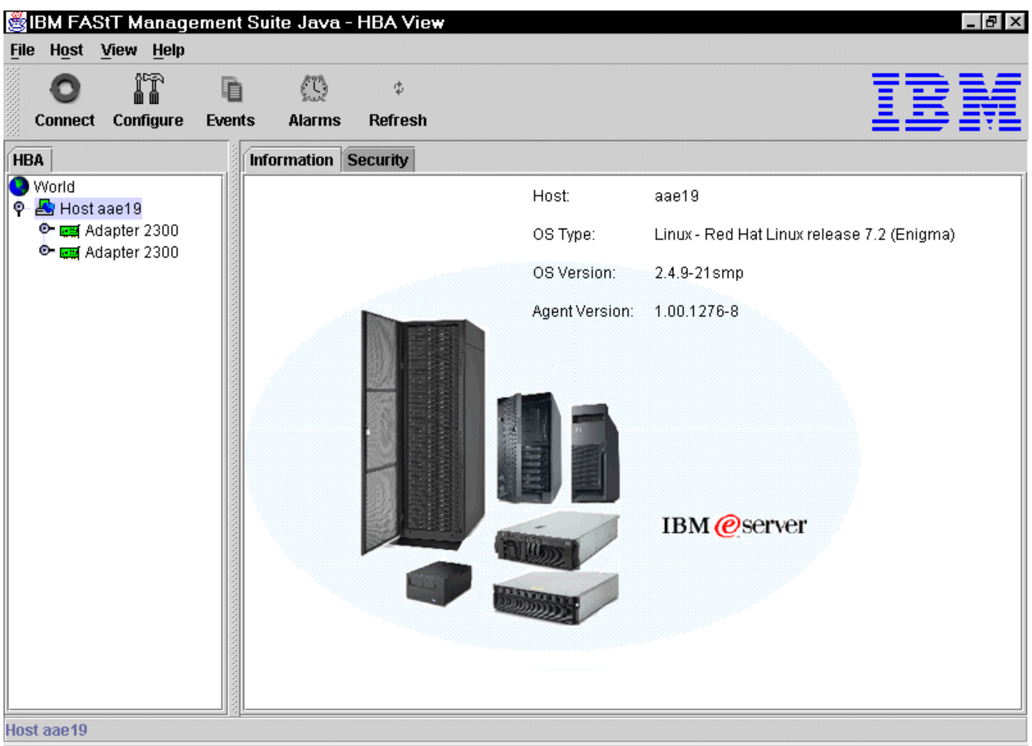

Figure 4-65   FAStT MSJ Host view

6. Click one of the host bus adapters in the left-hand pane and the right-hand pane changes. There are now seven tabs as shown in Figure 4-66.

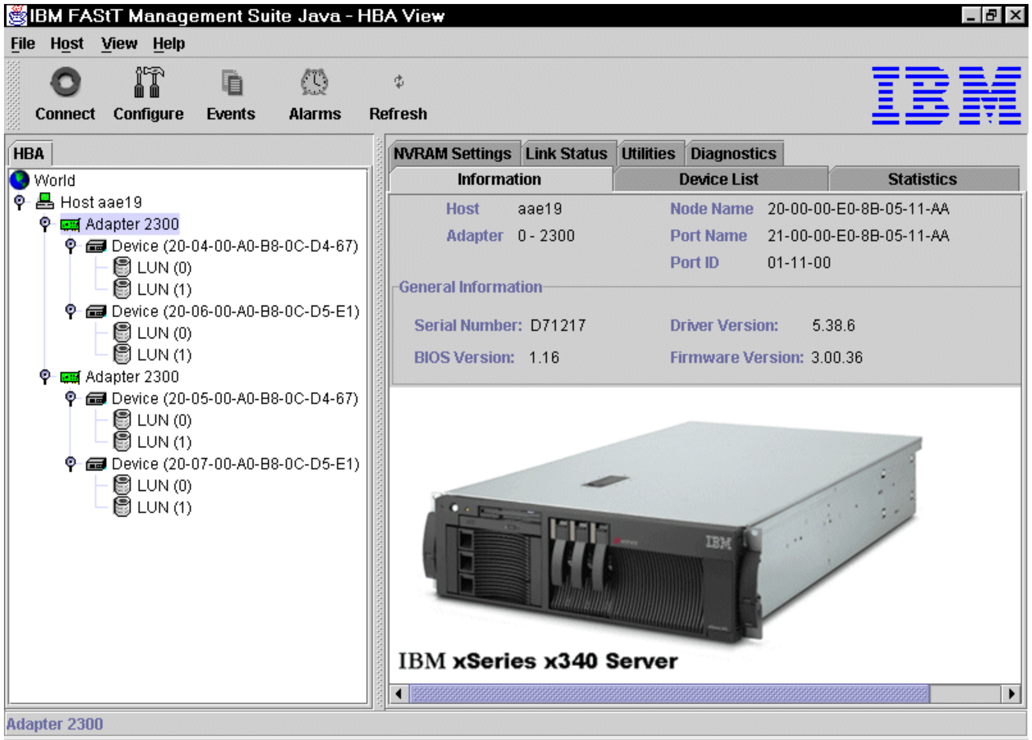

Figure 4-66 FAStT MSJ HBA view

The seven tabs are:

- **Information**: Displays general information about the server and host bus adapter, such as world-wide name, BIOS version, driver version, etc.
- **Device List**: Displays the devices currently available to the host bus adapter.
- **Statistics**: Displays a graph of the performance and errors on the host bus adapters over a period.
- **Link Status**: Displays the current link status.
- **NVSRAM Settings**: Displays the current settings and allows you to make remote configuration changes to the NVSRAM.
- **Utilities**: Allows you to update the flash and NVSRAM remotely (not on Linux).
- **Diagnostics**: Allow you to run diagnostic tests remotely.

For more detailed information on all of these functions, refer to the users guide, which is part of the download package.

5

FlashCopy and the FAStT Storage Server

This chapter introduces the FlashCopy option, which is a premium feature, of the FAStT Storage Manager 8.2. It introduces the various components of FlashCopy, what they do, and how to set them up. This is followed by a step-by-step guide to using the features and some examples of automated scripts and where they can be used.

5.1 FlashCopy: How it works

A FlashCopy logical drive is a point-in-time (PIT) image of a logical drive. The FlashCopy option is a premium feature that you must purchase separately from IBM or your IBM Business Partner. FlashCopy only works with IBM FAStT Storage Manager v8.0 and above. To learn more about the basic functions of FlashCopy, see 3.2.3, “FlashCopy” on page 48.

This section describes the various parameters for FlashCopy and what they mean.

5.1.1 Creating a FlashCopy logical drive

You can create FlashCopy logical drives either through the Create FlashCopy Logical Drive Wizard or by using the command line interface (CLI) with the `create` command. The latter can be scripted to support automatic operations. See 3.3.4, “Script Editor and command line interface” on page 58, for more information.

Note the following points:

- Refer to Appendix A, “Additional host-specific instructions for FlashCopy logical drives” on page 269. Failure to complete the additional steps required for your host operating system may result in an inaccurate point-in-time image of the base logical drive.
- You cannot create a FlashCopy logical drive of a base logical drive that is the secondary logical drive in a Remote Volume Mirror.
- If the FlashCopy logical drive is to be based on the root disk of the host operating system, the final point-in-time image may not be completely consistent with the base logical drive.

To create a FlashCopy logical drive, follow these steps:

1. Stop the host application that is accessing the base logical drive and unmount the base logical drive.

   Important: Unmounting the base logical drive does not apply when the base logical drive is the root disk of the host operating system.

2. Select a base logical drive from the Logical View. Then, click **Logical Drive-> FlashCopy-> Create**. Or you can right-click the logical drive and select **Create FlashCopy Logical Drive**.

   The Create FlashCopy Logical Drive Wizard begins.
3. Using the wizard, follow the instructions on each panel. Click the **Next** button when you're ready to move to the next panel. Each panel has context-sensitive help. Click the **Help** button to receive help for a particular screen.
4. After you create one or more FlashCopy logical drives, mount the base logical drive and restart the host application using that base logical drive.
5. Assign logical drive-to-LUN mappings between the FlashCopy logical drive and the host that will access the FlashCopy logical drive, using the Mappings View of the Subsystem Management window.

   Note: In some cases, depending on the host operating system and any logical drive manager software in use, mapping the same host to both a base logical drive and its associated FlashCopy logical drive may result in conflicts. See Appendix A, “Additional host-specific instructions for FlashCopy logical drives” on page 269, for more information.

6. Run the host-based hot_add utility to register the FlashCopy logical drive with the host operating system. Then run the host-based SMdevices utility to associate the mapping between the physical device name and the logical drive name.

Important: If you use this FlashCopy on a regular basis (for example, for backup purposes), use the Disable FlashCopy and Re-create FlashCopy options to reuse the FlashCopy. This alleviates the need to stop the host application and unmount the base logical drive while the FlashCopy is being created again. Using the Disable FlashCopy and Re-create FlashCopy options also preserves the existing mappings to the FlashCopy logical drive. For more information, see 5.1.2, "Disabling and recreating a FlashCopy logical drive" on page 141.

5.1.2 Disabling and recreating a FlashCopy logical drive

This section explains how you can disable the FlashCopy logical drive and then recreate it later.

Disabling a FlashCopy logical drive

If you no longer need a FlashCopy logical drive, you may want to disable it. As long as a FlashCopy logical drive is enabled, your storage subsystem performance is impacted by the copy-on-write activity to the associated FlashCopy repository logical drive. When you disable a FlashCopy logical drive, the copy-on-write activity stops.

If you disable the FlashCopy logical drive instead of deleting it, you can retain it and its associated repository. Then, when you need to create a different FlashCopy of the same base logical drive, you can use the re-create option to reuse a disabled FlashCopy. This takes less time than to create a new one.

When you disable a FlashCopy logical drive, note that:

- You cannot use that FlashCopy logical drive again until you use the re-create option on that logical drive.
- Only that FlashCopy logical drive is disabled. All other FlashCopy logical drives remain functional.

If you do not intend to re-create a FlashCopy, you can delete that FlashCopy logical drive instead of disabling it.

Re-creating a FlashCopy logical drive

Re-creating a FlashCopy logical drive takes less time than to create a new one. If you have a FlashCopy logical drive that you no longer need, instead of deleting it, you can reuse it (and its associated FlashCopy repository logical drive) to create a different FlashCopy logical drive of the same base logical drive.

When you re-create a FlashCopy logical drive, note that:

- The FlashCopy logical drive must be in either an optimal or a disabled state.
- All copy-on-write data on the FlashCopy repository logical drive is deleted.
- FlashCopy and FlashCopy repository logical drive parameters remain the same as the previously disabled FlashCopy logical drive and its associated FlashCopy repository logical drive. After the FlashCopy logical drive is re-created, you can change parameters on the FlashCopy repository logical drive through the appropriate menu options.
- The original names for the FlashCopy and FlashCopy repository logical drives are retained. You can change these names after the re-create option completes.

5.1.3 FlashCopy parameters

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive, its capacity, threshold level warning, and other parameters. The FlashCopy repository logical drive capacity is created as a percentage of the base logical drive and contains the copy-on-write data. The Storage Management software provides a warning message when your FlashCopy repository logical drive exceeds the threshold level.

The FlashCopy repository logical drive's failure policy determines what happens when the FlashCopy repository logical drive becomes full (that is, all of its capacity has been used). The failure policy can be set to either fail the FlashCopy (default setting) or fail incoming I/O to the FlashCopy's base logical drive.

Do not ignore the FlashCopy repository logical drive “threshold exceeded” notification. This is the last and only warning you receive before the FlashCopy repository logical drive becomes full. You have the option to increase the capacity of the FlashCopy repository logical drive or increasing the FlashCopy repository logical drive threshold capacity warning level. Increasing the warning threshold reduces the time you have to respond the next time you receive a threshold exceeded notification.

If a FlashCopy logical drive or FlashCopy repository logical drive is displayed as a missing logical drive, the storage subsystem has detected drives associated with the FlashCopy, or FlashCopy repository logical drive are no longer accessible. Missing logical drives, in most cases, are recoverable. See 8.2.4, “Missing logical drives” on page 259.

FlashCopy logical drive maintenance

The default FlashCopy repository logical drive capacity is set to 20% of the base logical drive, if enough free capacity exists to create a FlashCopy repository logical drive of this size. The default threshold level for the FlashCopy repository logical drive is set to 50%.

If you are not sure how large to make the FlashCopy repository logical drive or how high to set the FlashCopy repository logical drive full warning, accept the default settings. You can estimate later how quickly the FlashCopy's repository capacity is being used. For more information, see 5.1.5, “Estimating FlashCopy repository life” on page 146.

Viewing FlashCopy logical drive failure settings

To see the current failure settings, complete these steps:

1. Select a FlashCopy repository logical drive in the Logical View of the Subsystem Management window.
2. Click **Logical Drive-> Properties**, or right-click and select **Properties**.

   The FlashCopy Repository Logical Drive - Properties window (Figure 5-1) opens.

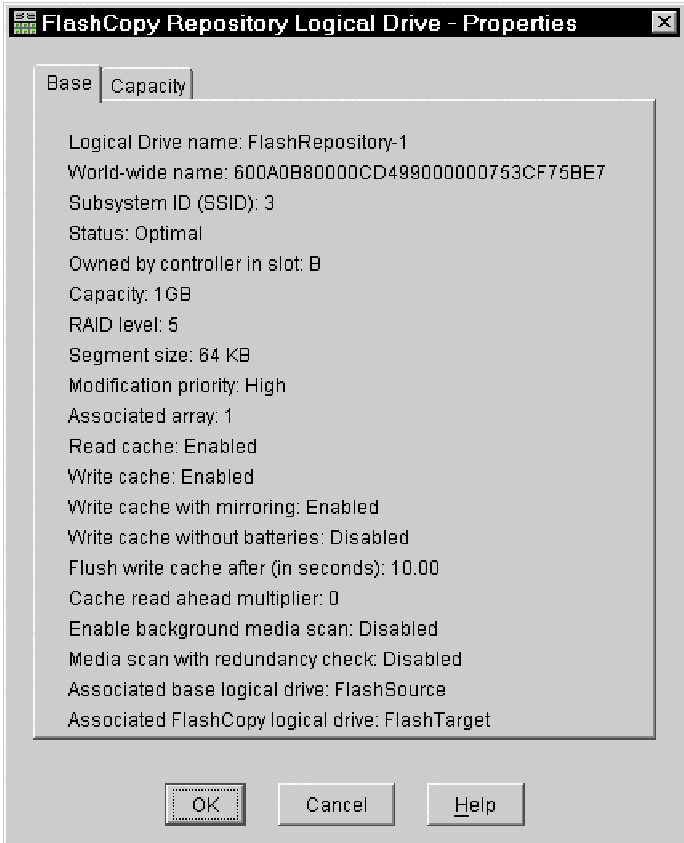

Figure 5-1 FlashCopy Repository logical drive properties

3. Select the **Capacity** tab (Figure 5-2) to view the currently defined settings.

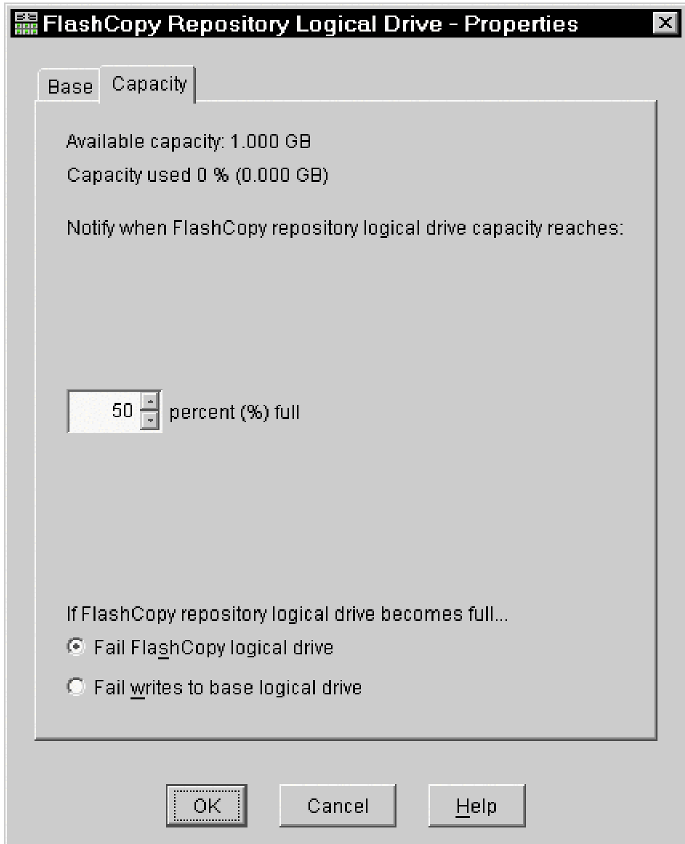

Figure 5-2 Repository Capacity settings

If the FlashCopy repository logical drive is set to fail, the data will not be recoverable when the drive becomes full and FlashCopy cannot be accessed. The only available option (if this situation occurs) is to delete the FlashCopy logical drive or re-create the FlashCopy logical drive to create a new point-in-time image.

If the FlashCopy repository logical drive is set to fail writes to the base logical drive, the data is recoverable. But, the FlashCopy repository logical drive capacity must be increased before writes to the base logical drive are not rejected. See 5.1.6, “Increasing the capacity of a FlashCopy repository logical drive” on page 147.

Note the following points:

- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy repository logical drive.
- Deleting a FlashCopy repository logical drive automatically deletes the associated FlashCopy logical drive.
- Deleting a FlashCopy logical drive and then creating it again forces you to stop the host application and unmount the base logical drive while the FlashCopy is being created again.
- Re-creating a FlashCopy logical drive alleviates the need to create a FlashCopy repository logical drive, as well as re-map the assigned logical drive-to-LUN mappings between the FlashCopy logical drive and the host.
- After the FlashCopy logical drive is re-created, you can change parameters on the FlashCopy repository logical drive through the appropriate menu options.

- To avoid another “FlashCopy repository logical drive capacity full” failure, increase the capacity of the FlashCopy repository logical drive. See 5.1.6, “Increasing the capacity of a FlashCopy repository logical drive” on page 147.

5.1.4 Estimating FlashCopy repository logical drive capacity

During the creation of a FlashCopy logical drive, a physical logical drive called the *FlashCopy repository logical drive* is created to store FlashCopy data and copy-on-write data. The default setting for the FlashCopy repository logical drive capacity is 20% of the base logical drive's capacity. In general, this capacity should be sufficient. However, use the following information to help determine the appropriate capacity of the FlashCopy repository logical drive:

- A FlashCopy repository logical drive may be no smaller than 8 MB.
- The amount of write activity to the base logical drive after the FlashCopy logical drive has been created dictates how large the FlashCopy repository logical drive needs to be. As the amount of write activity to the base logical drive increases, the number of original data blocks that need to be copied from the base logical drive to the FlashCopy repository logical drive also increases.
- The estimated life expectancy of the FlashCopy logical drive contributes to determining the capacity of the FlashCopy repository logical drive. If the FlashCopy logical drive is created and remains enabled for a long period of time, the FlashCopy repository logical drive runs the risk of reaching its maximum capacity. For more information, see 5.1.5, “Estimating FlashCopy repository life” on page 146.
- The amount of management overhead required on the FlashCopy repository logical drive to store FlashCopy logical drive data contributes to determining the FlashCopy repository logical drive's capacity. The amount of management overhead actually required is fairly small and can be calculated using the simple formulas detailed in the following section.
- There is not necessarily a one-to-one correlation between the number of data blocks that change on the base logical drive and the amount of copy-on-write data stored on the FlashCopy repository logical drive. Depending on the location of data blocks that need to be copied, for performance reasons, the controller may copy over a full set of 32 blocks, even if only one set of blocks has changed. Keep this in mind when determining the percentage of the base logical drive's capacity that may be copied to the FlashCopy repository logical drive.

Calculating expected overhead

Use the following formula to calculate the amount of management overhead required to store FlashCopy data on the FlashCopy repository logical drive. This formula should be used merely as a guide, and FlashCopy repository logical drive capacity should be re-estimated periodically.

> **Note:** Conversion from bytes to kilobytes, and then to megabytes, is required for this formula.

The formula to calculate the amount of management overhead required is:

```
192 KB + (X/2000)
```

Here *X* is the capacity of the base logical drive in bytes.

Example

For a 5 GB base logical drive, where 30% of the data blocks are expected to change on the base logical drive, the estimated FlashCopy repository logical drive capacity can be calculated as follows:

1. Convert the base logical drive's capacity to bytes.

   When converted, 5 GB equals 5,368,709,120 bytes.

2. Divide the base logical drive's capacity (in bytes) by 2000.

   When divided, the result is 2,684,354.56 bytes.

3. Convert the result from step 2 (in bytes) to kilobytes (KB).

   When converted, the result is 2621.44 KB.

4. Add 192 KB to the results from step 3.

   ```
   192 KB + 2621.44 KB = 2813.44 KB
   ```

5. Convert the result from step 4 to megabytes (MB).

   When converted, the amount of management overhead required is calculated to be 2.75 MB (or 0.002686 GB).

6. In this example, 30% of the data blocks on the base logical drive are expected to change. To accurately calculate the FlashCopy repository logical drive capacity, sufficient space needs to be allowed for the copy-on-write data as well as the management overhead (calculated in step 5).

   To calculate the copy-on-write space required, calculate the percentage of the base logical drive expected change:

   ```
   30% of 5 GB = 1.5 GB
   ```

   The final estimated FlashCopy repository logical drive capacity for this example is:

   ```
   1.5 GB + 0.002686 GB = 1.502686 GB
   ```

7. In the Create FlashCopy Logical Drive Wizard: Specify Repository Capacity window, use the percentage (%) full box of base logical drive to set the estimated FlashCopy repository logical drive capacity (Figure 5-2 on page 144).

   > **Note:** The percentage (%) full box box sets the FlashCopy repository logical drive capacity as a percentage of the base logical drive. Using the percentage (%) full box, increase or decrease the percentage until the FlashCopy Repository Logical Drive Capacity value matches the estimated capacity calculated in step 6. (Some rounding up may be required.)

5.1.5 Estimating FlashCopy repository life

During the creation of a FlashCopy logical drive, you are asked to define various properties for the FlashCopy repository logical drive, including the FlashCopy repository logical drive's name, capacity, a logical drive-to-LUN mapping, and the repository full condition. When defining the FlashCopy repository logical drive's properties, keep in mind the kind of usage you have planned for the FlashCopy logical drive. Understanding how the FlashCopy logical drive will be used can help you to estimate the life expectancy of the FlashCopy repository logical drive.

If numerous I/O requests are written to the base logical drive, the FlashCopy repository logical drive, which contains the FlashCopy data (information about the FlashCopy) and copy-on-write data, could eventually exceed the base logical drive capacity if all the original data blocks are changed. The default setting suggests 20% of the base logical drive capacity, but this is a setting that can be fine tuned after some usage data becomes available.

The following procedure describes how to estimate the life expectancy of a FlashCopy repository logical drive. Use it merely as a guide. Note that a FlashCopy repository logical drive's life expectancy should be carefully re-estimated periodically.

1. Highlight the FlashCopy repository logical drive in the Logical View. Click **Logical Drive-> Properties** and choose the **Capacity** tab. Or you can right-click the FlashCopy repository logical drive, select **Properties**, and click the **Capacity** tab.
2. Highlight the FlashCopy repository logical drive in the Logical View. Click **View-> Go To-> FlashCopy Logical Drive**. Or right-click the FlashCopy repository logical drive and select **Go To FlashCopy Logical Drive**.
3. Record the creation timestamp day and time.
4. Record the capacity used (GB) and the available capacity (GB).
5. Determine the elapsed time (t) by subtracting the creation time from the current time, expressing the elapsed time in either minutes, hours, or days.
6. The total time (Tr) that the repository is available for copy-on-write data can now be estimated (based on the current usage) by multiplying the elapsed time (t) by the available capacity (Ct), and then dividing the resultant number by the capacity used (Cu).

   In summary, note that:

   – Tr = Total Time Available
   – t = Elapsed Time
   – Ct = Available Capacity
   – Cu = Used Capacity

   Therefore, the formula you need to use is:

   ```
   Tr = (t * Ct) / Cu
   ```

   > **Note:** The Total Time Available (Tr) indicates the total usage time for the FlashCopy repository logical drive.

Based on the results, you should now be able to make an informed decision about whether you should increase the capacity of the FlashCopy repository logical drive. If the repository capacity becomes 100% full during the FlashCopy's expected lifetime, then you should increase the FlashCopy repository logical drive's capacity. For more information, see 5.1.6, "Increasing the capacity of a FlashCopy repository logical drive" on page 147.

5.1.6 Increasing the capacity of a FlashCopy repository logical drive

This option is used to increase the storage capacity of an existing FlashCopy repository logical drive. Typically, this option is used when a warning is received that the FlashCopy repository logical drive is in danger of becoming full.

You can achieve an increase in storage capacity by:

- Using free capacity available on the array of the FlashCopy repository logical drive.
- Adding unconfigured capacity (in the form of unused drives) to the array of the FlashCopy repository logical drive. Use this option when no free capacity exists on the array.

> **Important:** A maximum of two drives may be added at one time to increase FlashCopy repository logical drive capacity.

The storage capacity of a FlashCopy repository logical drive *cannot* be increased if:

- One or more hot spare drives are in use in the logical drive.
- The logical drive has a *Non-Optimal* status.
- Any logical drive in the array is in any state of modification.
- The controller that owns this logical drive is in the process of adding capacity to another logical drive (each controller can add capacity to only one logical drive at a time).
- No free capacity exists in the array.
- No unconfigured capacity (in the form of drives) is available to add to the array.

To increase the capacity, highlight a FlashCopy repository logical drive in the Logical View of the Subsystem Management window. Then, click **Logical Drive-> Increase Capacity**. Or you can right-click and select **Increase Capacity**.

Important: If no free capacity or unconfigured capacity is available, the Increase Capacity option is not available.

The Increase Repository Capacity dialog opens. The FlashCopy repository logical drive name, the associated FlashCopy logical drive name, the associated base logical drive name, current capacity, and amount of free capacity available for the selected repository are displayed. If free capacity is available, the maximum free space is shown in the Increase capacity by box.

If there is no free capacity on the array, the free space that is shown in the Increase capacity by box is 0. Drives must be added to create free capacity on the array of the standard logical drive. See “Dynamic logical drive expansion” on page 126 for more information.

5.2 Step-by-step guide to FlashCopy

This section presents an easy-to-follow, step-by-step procedural guide to common administration tasks for a FAStT FlashCopy solution. The tasks that are covered include:

- Checking the status of the Flash Copy premium feature
- Creating a FlashCopy drive
- Mapping a FlashCopy drive to a host
- Viewing FlashCopy drive status
- Disable FlashCopy a drive
- Re-creating FlashCopy drive
- Resizing FlashCopy repository drive
- Deleting a FlashCopy drive

5.2.1 Checking the status of the Flash Copy premium feature

Use this procedure to view a list of premium features on the storage subsystem and to verify that the FlashCopy feature has been enabled. It requires you to:

1. Check the premium options list.
2. View the FlashCopy icon to verify that the feature has been enabled.

Checking the premium option

To check the premium option, follow these steps:

1. From the Subsystem Management window, click **Storage Subsystem-> Premium Features-> List**. Or you can right-click and click **Premium Features-> List** (see Figure 5-3).

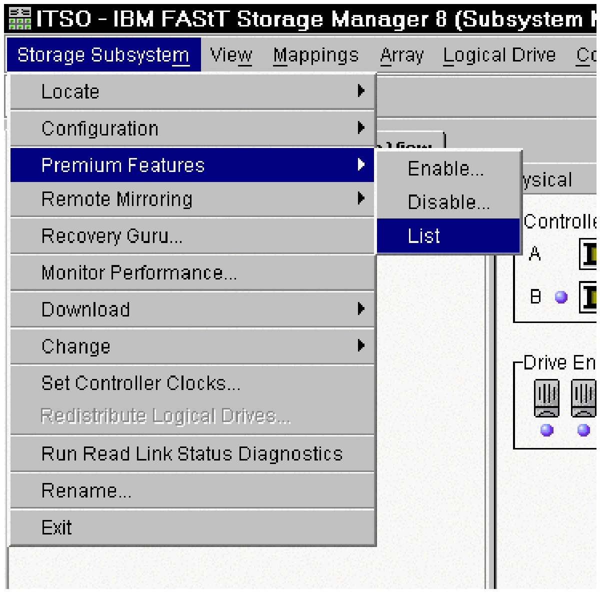

Figure 5-3 Listing premium features

The List Premium Features dialog opens. It lists the following items:

– Premium features enabled on the storage subsystem
– Feature Enable Identifier

2. Verify that “FlashCopy Logical Drives:” indicates *Enabled* as shown in Figure 5-4.

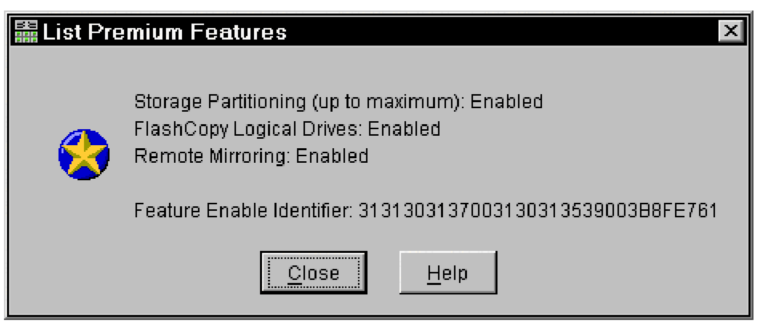

Figure 5-4 Premium features

Figure 5-5 demonstrates an example where the FlashCopy copy feature is not enabled.

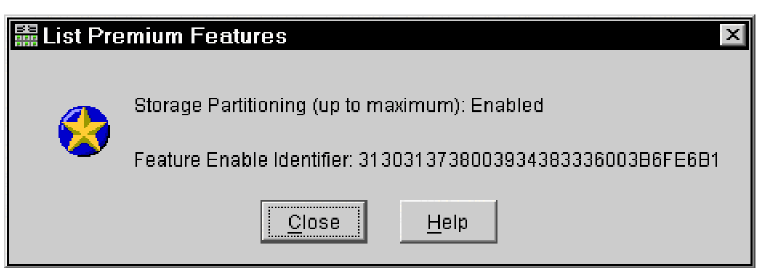

Figure 5-5 Feature list: No FlashCopy

If you identify that the FlashCopy feature is not enabled, see 5.2.2, “Enabling the FlashCopy premium feature” on page 150, for the procedure to enable it.

3. Click **Close** to close the dialog.

Note: If you receive a “Premium Features - Out of Compliance” error message during a management session, use the Recovery Guru to resolve the problem.

Viewing the FlashCopy icon

Check the status of the FlashCopy feature icon at the bottom left of the device management GUI, as shown in Figure 5-6.

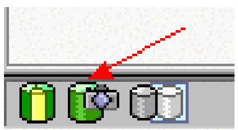

Figure 5-6 FlashCopy feature icon

The example in Figure 5-7 shows a disabled FlashCopy feature.

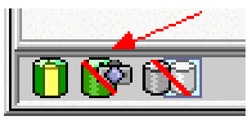

Figure 5-7 FlashCopy feature icon disabled

5.2.2 Enabling the FlashCopy premium feature

If the current status of the FlashCopy Option is not *Enabled*, complete the following steps:

1. From the Subsystem Management window, click **Storage subsystem-> Features ->Enable**. The Select Feature Key File window opens.

   Note: The Select Feature Key File program filters files with the *.key* extension.

2. Select the folder in which you placed the generated key file.
3. Select the appropriate key file, and then click **OK**.
4. The Enable Feature window opens. Click **Yes**.
5. The FlashCopy Option is now enabled. The icon in the Premium Feature status area no longer displays a red slash. To further verify the status of the option, click **Storage subsystem-> Features-> List**.

5.2.3 Creating a FlashCopy drive

This section takes you through the procedure to create a FlashCopy drive from a base volume.

Important:

- Refer to the Create FlashCopy Logical Drive Wizard online help for additional instructions on the operating system-specific procedures. Failure to complete the additional steps required for your host operating system may result in an inaccurate point-in-time image of the base logical drive.
- You *cannot* create a FlashCopy logical drive of a base logical drive that is the secondary logical drive in a Remote Volume Mirror.

To create a FlashCopy drive, complete these steps:

1. Place the application into backup mode:

   a. Stop the host application that is accessing the base logical drive, sync file system. Unmount the base logical drive if possible. This enables a valid consistent copy to be taken.

      It may not be convenient to stop database applications. However, in this case, it is required to place the application into a backup mode or place it in an acquiesced state for the duration of the FlashCopy creation.

   b. It is also important to back up application recovery files, such as role back and redo logs, because these may be located on different physical disk storage or logical drives.

2. Launch the Create FlashCopy Logical Drive Wizard:

   a. Select a base logical drive from the Logical View.

   b. Click **Logical Drive-> FlashCopy-> Create**. Or you can right-click and select **Create FlashCopy Logical Drive**. See Figure 5-8.

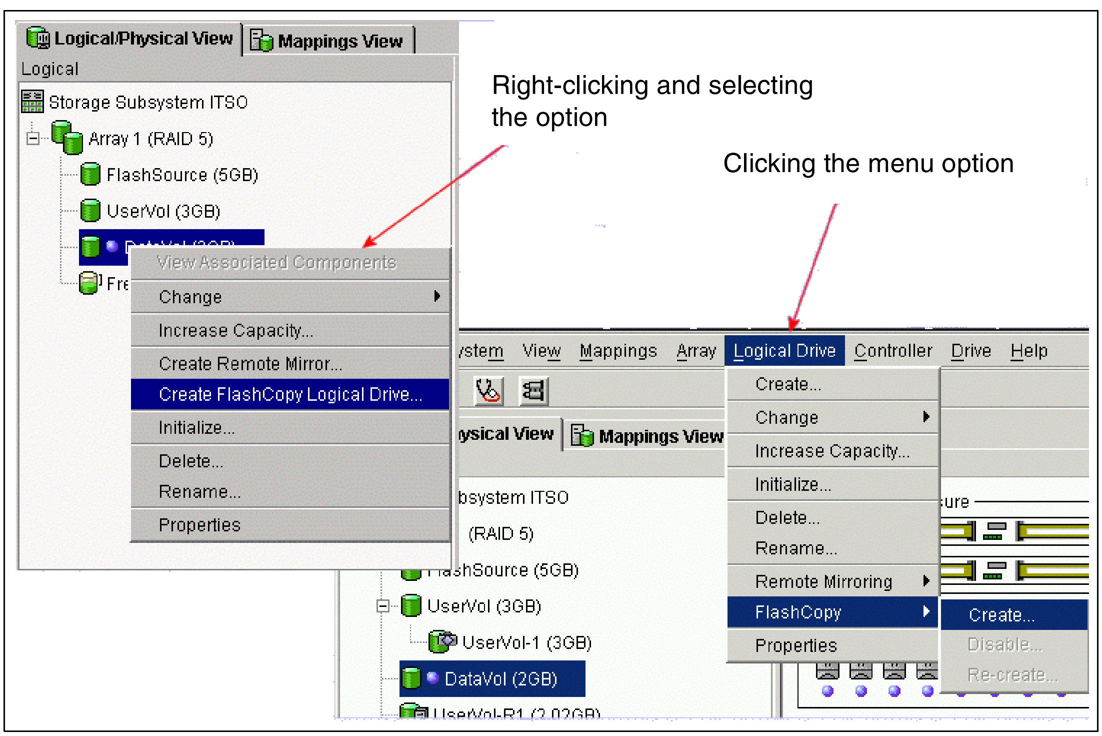

Figure 5-8 Create FlashCopy menu options

The Create FlashCopy Logical Drive Wizard begins as shown in Figure 5-9.

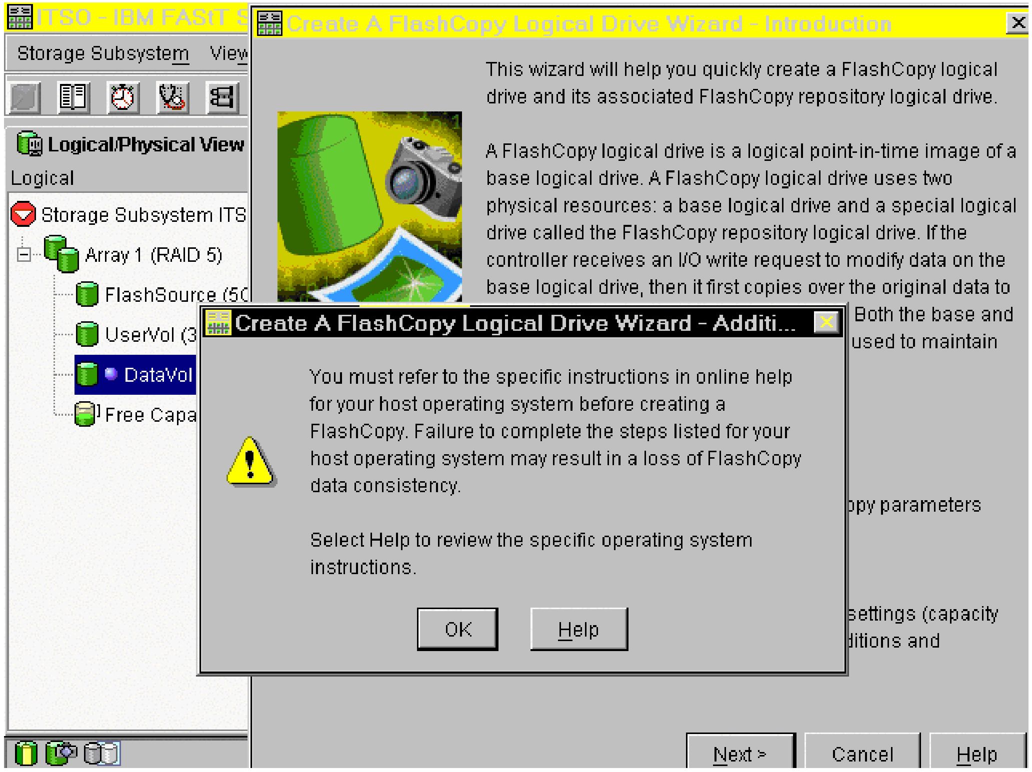

Figure 5-9 Create FlashCopy Logical Drive Wizard startup

> **Attention:** If the FlashCopy logical drive is to be based on the root disk of the host operating system, the final point-in-time image may not be completely consistent with the base logical drive.

3. Create the FlashCopy logical drive:

   a. Review the information on the initial window as shown in Figure 5-9. Click **OK** to proceed to the wizard introduction window.

   Follow the instructions on each wizard panel, and click the **Next** button when you're ready to continue to the next panel.

   > **Note:** Each wizard panel has context-sensitive help. Click the **Help** button on a particular panel to receive help for that panel.

   b. The Introduction window (Figure 5-10) defines what a FlashCopy logical drive is and the physical components associated with a FlashCopy logical drive. It enables you to select either the Simple or Advanced path through the Create FlashCopy Logical Drive Wizard:

   - **Simple Path**: Proceeds to the *Specify Name* panel, which provides a preview of the FlashCopy and repository default names. You can also change the defaults on this panel.
   - **Advanced Path**: Proceeds to the *Allocate Capacity* panel, on which you select the Free Capacity or Unconfigured Capacity node on which to place the FlashCopy repository logical drive.

If no free capacity exists or the available free capacity is unusable, a warning message appears.

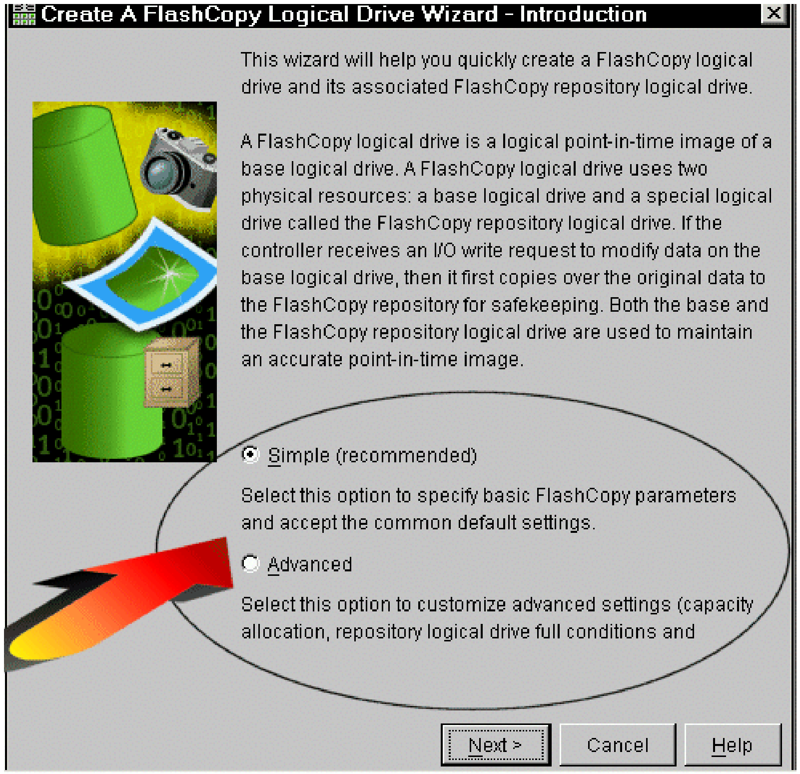

Figure 5-10 Wizard Introduction window

c. If you want to use the Advanced path, select **Advanced** and click **Next**. The Allocate Capacity window (Figure 5-11) appears on which you can choose:

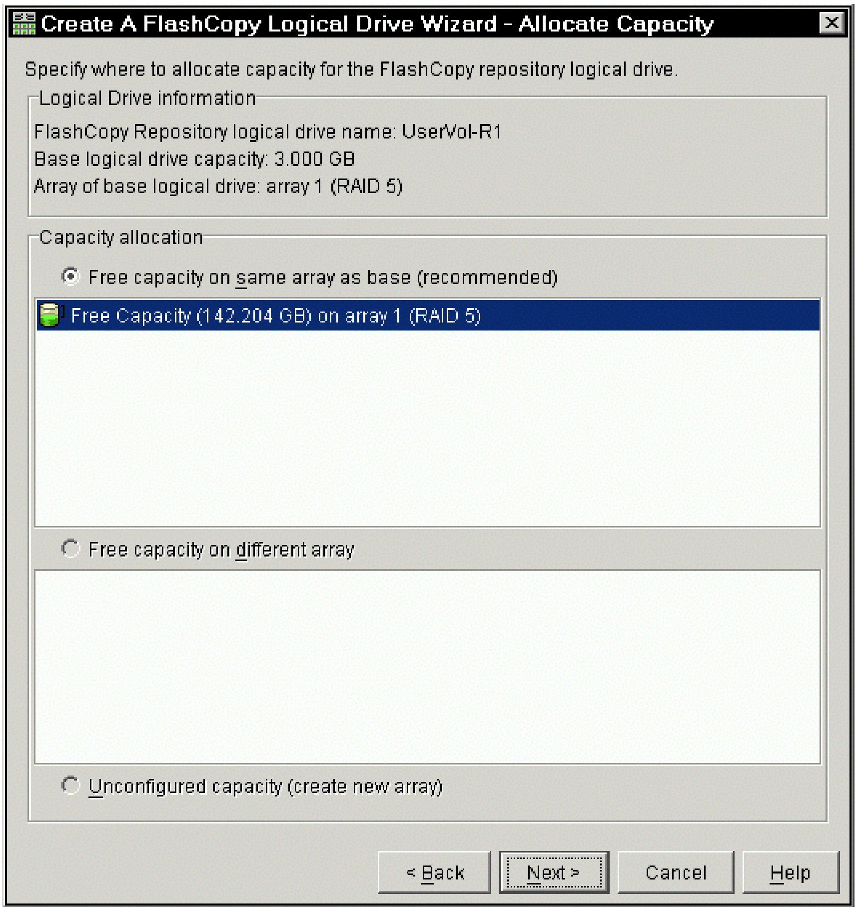

Figure 5-11 Allocate Capacity window

- **Free Capacity**:

  If you select **Free Capacity** and click **Next**, the Specify Logical Drive Parameters window (Figure 5-12) opens. Define the FlashCopy logical drive-to-LUN mapping parameter and the FlashCopy repository logical drive full conditions. Click the **Next** button.

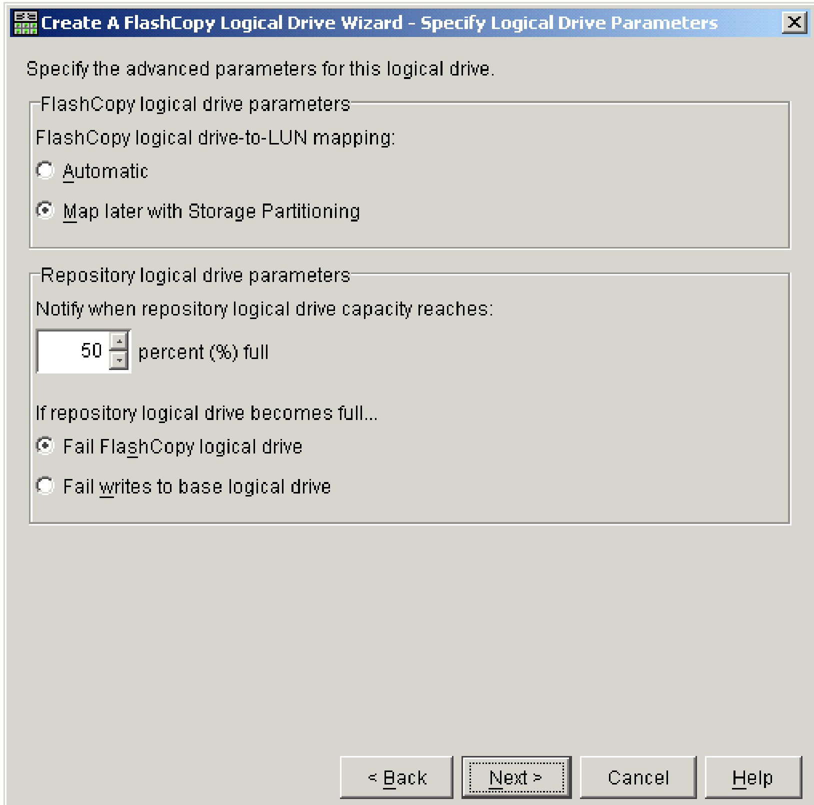

Figure 5-12 Specify Logical Drive Parameters

Then the Preview window (Figure 5-13) opens and shows the FlashCopy and repository parameters. It allows you to click **Back** to return to the previous windows and edit the parameters or to click **Finish** to continue.

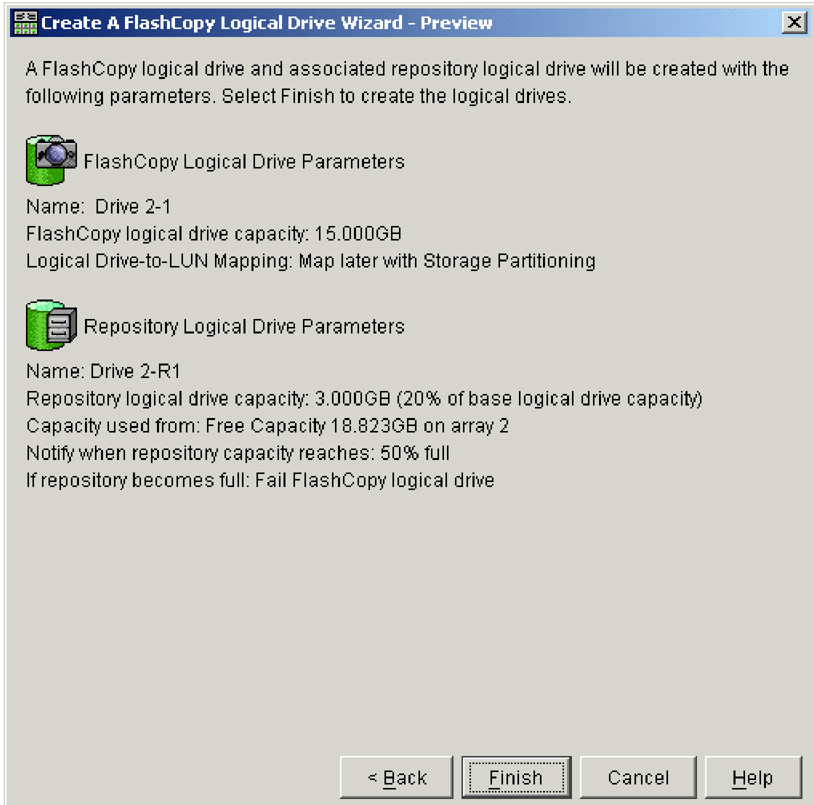

Figure 5-13 Preview window

- **Unconfigured Capacity**:

  On the Allocate Capacity window (Figure 5-11 on page 154), you may select **Unconfigured Capacity** and click **Next**. Then on the Specify Array Parameters window (Figure 5-14), you specify a new array where the repository resides. You also specify the RAID level of the array that meets the FlashCopy repository logical drive data storage and protection requirements.

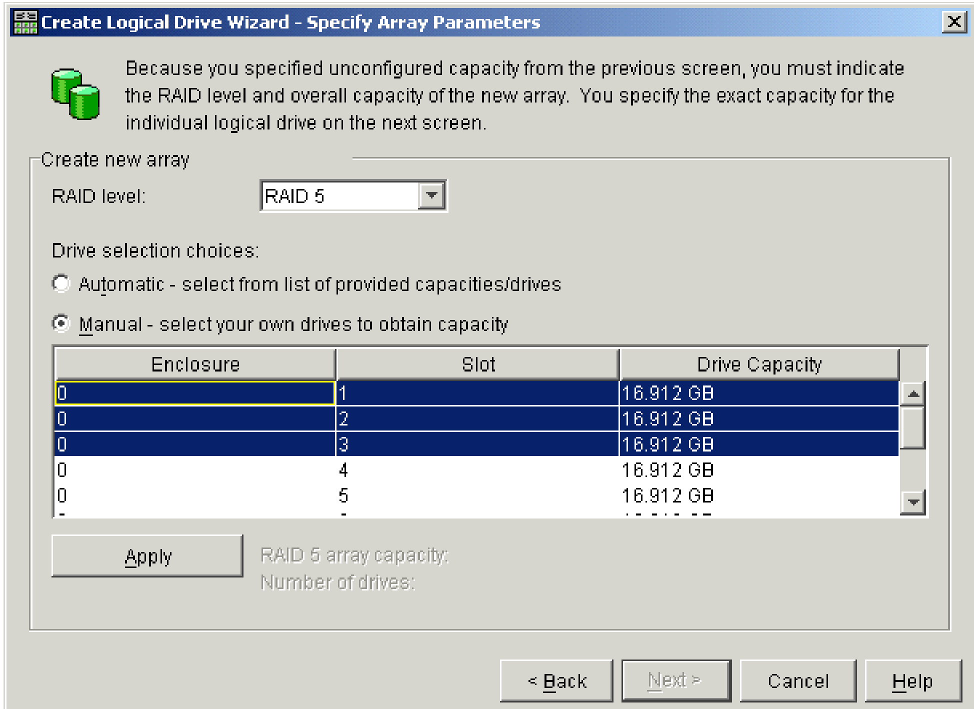

Figure 5-14 Specify Array Parameters window

Once you select where you are placing the volume, same array, and free capacity on another array, or you create a new array, the actual steps for defining the FlashCopy logical drive parameters are the same as for the Free Capacity option as explained here.

On the Specify Logical Drive Parameters window (Figure 5-12), you define the FlashCopy logical drive-to-LUN mapping parameter, the threshold percentage full parameter, and the FlashCopy repository logical drive full conditions.

The Preview window (Figure 5-13) provides a preview of the FlashCopy and repository parameters. It allows you to click **Back** to return to the previous windows to edit the parameters or to click **Finish** and continue.

d. On the Specify Names window (Figure 5-15), you define the FlashCopy logical drive name and the name of its associated FlashCopy repository logical drive.

   The default naming convention for the first FlashCopy uses the base volume name and adds a suffix of “-1” for the logical drive and “-R1” for the repository drive. The second FlashCopy uses 2 instead of 1. This is repeated up to the four volumes.

   For example, if you are creating the first FlashCopy logical drive for a base logical drive called *DataVol*, then the default FlashCopy logical drive name is DataVol-1, and the associated FlashCopy repository logical drive default name is DataVol-R1. The default name of the next FlashCopy logical drive you create based on DataVol is DataVol-2, with the corresponding FlashCopy repository logical drive named DataVol-R2 by default.

   Change the default names if required.

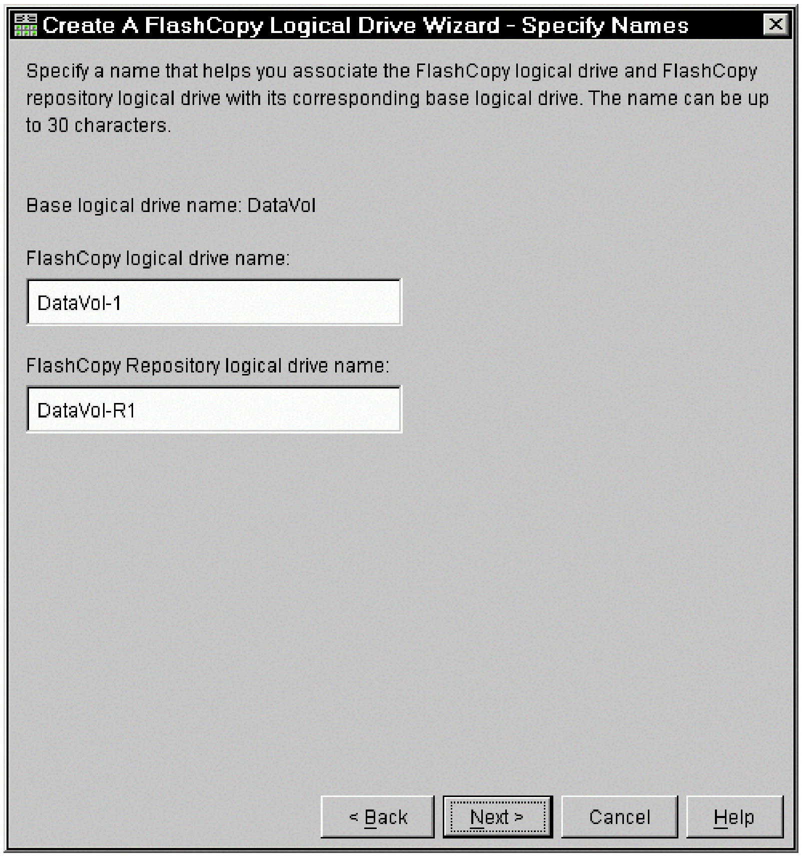

Figure 5-15 Specify Names

Tips:

- Regardless of whether you use the software-supplied sequence number that (by default) populates the *FlashCopy logical drive name* or *FlashCopy repository logical drive name* field, the next default name for a FlashCopy or FlashCopy repository logical drive still uses the sequence number determined by the software. For example, you may name the first FlashCopy of base logical drive DataVol *DataVolMay28*, and don't use the software-supplied sequence number of 1. Then the default name for the next FlashCopy of Accounting is still *DataVol-2*.
- The next available sequence number is based on the number of existing FlashCopies of a base logical drive. If you delete a FlashCopy logical drive, its sequence number becomes available again.
- You must choose a unique name for the FlashCopy and FlashCopy repository logical drives. Otherwise an error message is displayed.
- There is a 30-character limit. After you reach this limit in either the FlashCopy logical drive name or FlashCopy repository logical drive name fields, you can no longer type in the field. If the base logical drive is 30 characters, then the default names for the FlashCopy and its associated FlashCopy repository logical drive use the base logical drive name truncated enough to add the sequence string. For example, for "Host Software Engineering Group GR-1", the default FlashCopy name would be "Host Software Engineering GR-1". The default repository name would be "Host Software Engineering G-R1."

Click **Next** to proceed to continue.

e. On the Specify FlashCopy Repository Logical Drive's Capacity window (Figure 5-16), set the repository drive capacity as a percentage of the base logical drive's capacity.

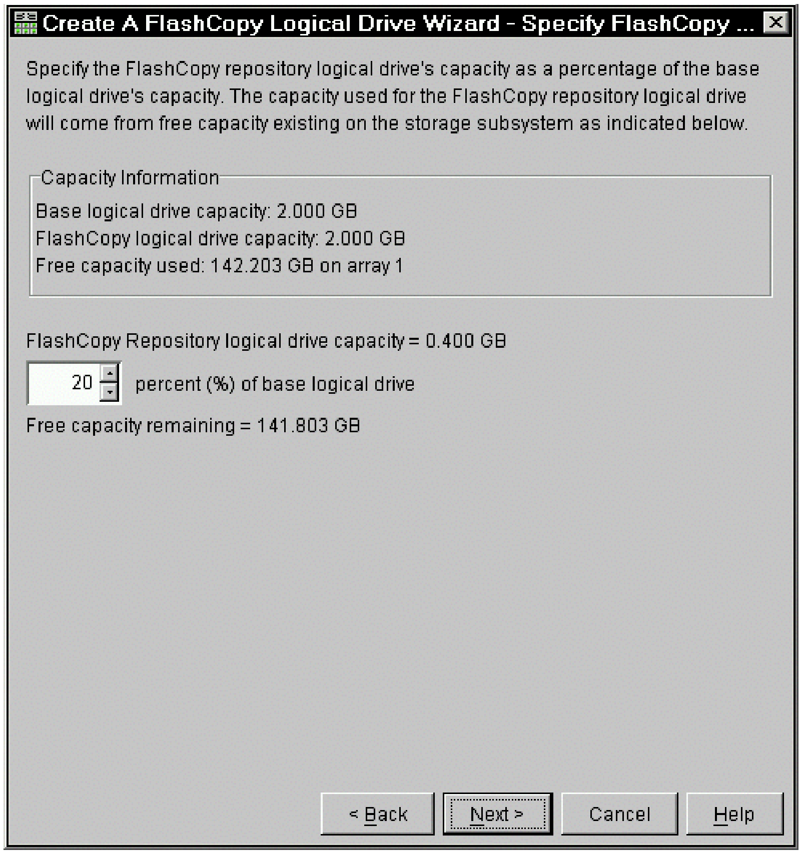

Figure 5-16 Specify Repository Drive Capacity

Use the percent (%) of base logical drive box to set the desired capacity. Click **Next** to continue.

The capacity needed varies, depending on the frequency and size of I/O writes to the base logical drive and how long you need to keep the FlashCopy logical drive. In general, you should choose a larger capacity for the repository if you intend to keep the FlashCopy logical drive for a long period of time or if a large percentage of data blocks will change on the base logical drive during the life of the FlashCopy logical drive due to heavy I/O activity. Use historical performance monitor data or other operating system utilities to help you determine typical I/O activity to the base logical drive.

Important: In most situations, the 20% default value should be ample capacity for your FlashCopy repository logical drive. For information on determining the size, see 5.1.4, “Estimating FlashCopy repository logical drive capacity” on page 145, and 5.1.5, “Estimating FlashCopy repository life” on page 146.

f. The Create FlashCopy Logical Drive Wizard - Preview window opens. It displays components associated with the FlashCopy. Review the information and click **Next**.

g. The Create FlashCopy Logical Drive Wizard - Completed window (Figure 5-17) appears. It display the associated logical drives and components that make up the FlashCopy relationship. Click **Finish** to continue.

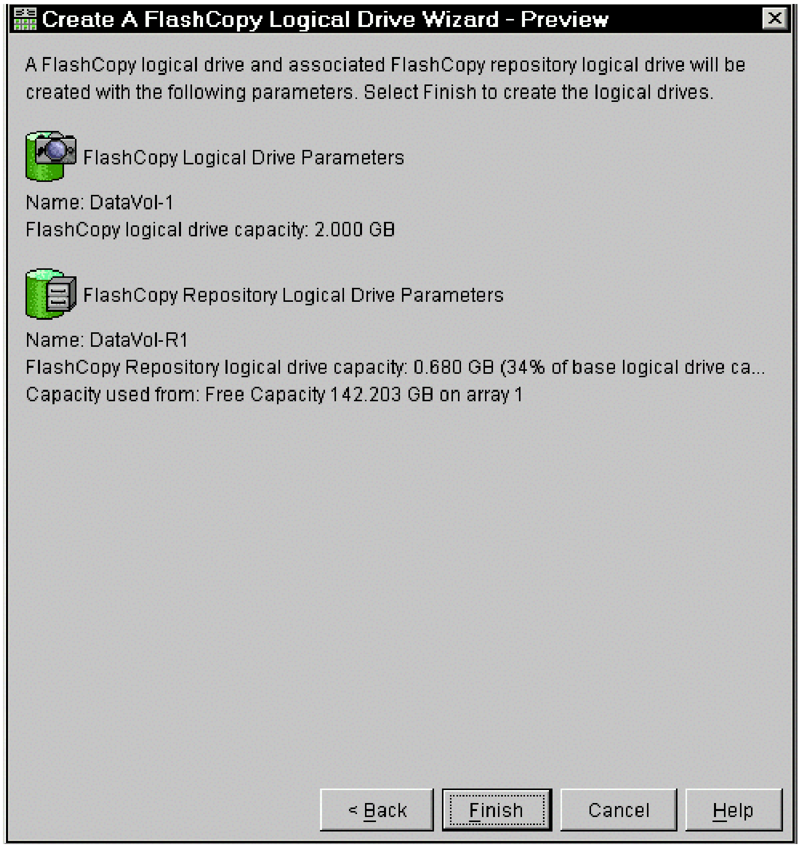

Figure 5-17 Preview window

The FlashCopy drives are now displayed in the device management GUI, as shown in Figure 5-18. See the section on viewing FlashCopy status for information about interpreting the display icons.

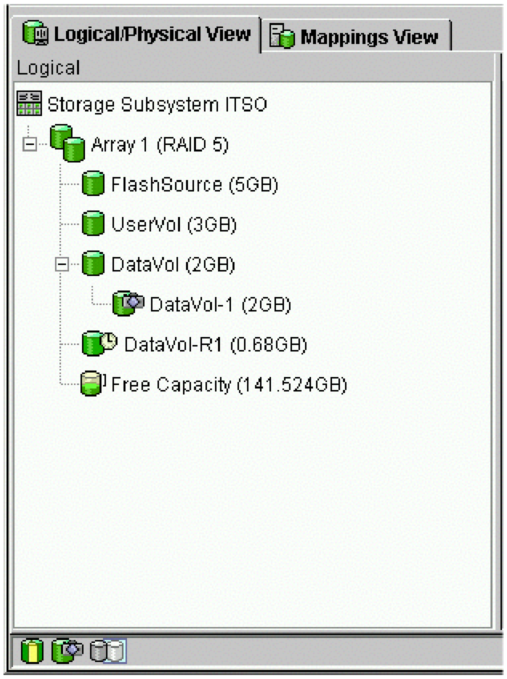

Figure 5-18 FlashCopy volumes

4. Restart the host application. After you create one or more FlashCopy logical drives, mount the base logical drive and restart the host application using that base logical drive.

5.2.4 Mapping a FlashCopy drive to a host

Assign logical drive-to-LUN mappings between the FlashCopy logical drive and the host that will access the FlashCopy logical drive using the Mappings View of the Subsystem Management window. In some cases, depending on the host operating system and if any logical drive manager software in use, mapping the same host to both a base logical drive and its associated FlashCopy logical drive may result in conflicts. For operating-system specific instructions, see Appendix A, "Additional host-specific instructions for FlashCopy logical drives" on page 269.

To map the FlashCopy logical drive to a host, follow these instructions:

1. Open the Mappings View (Figure 5-19) of the Management window. The newly create FlashCopy logical drive appears in the undefined mapping section.

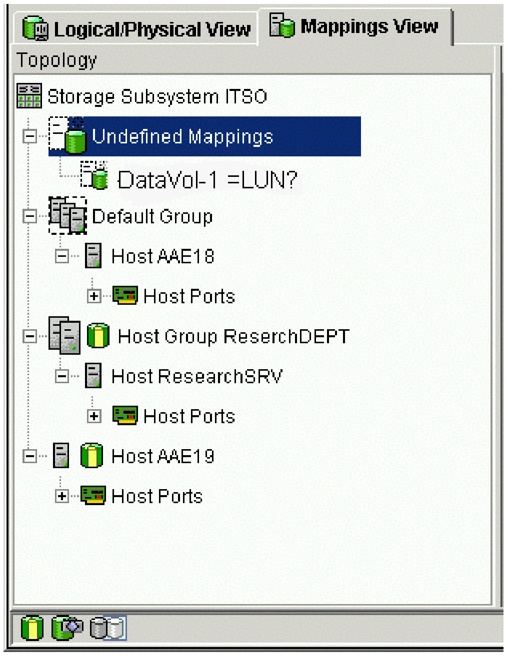

Figure 5-19 Undefined FlashCopy Disk

2. Select the host or host group you want to map the drive.
3. Right-click and select **Define Additional Mapping**, as shown in Figure 5-20.

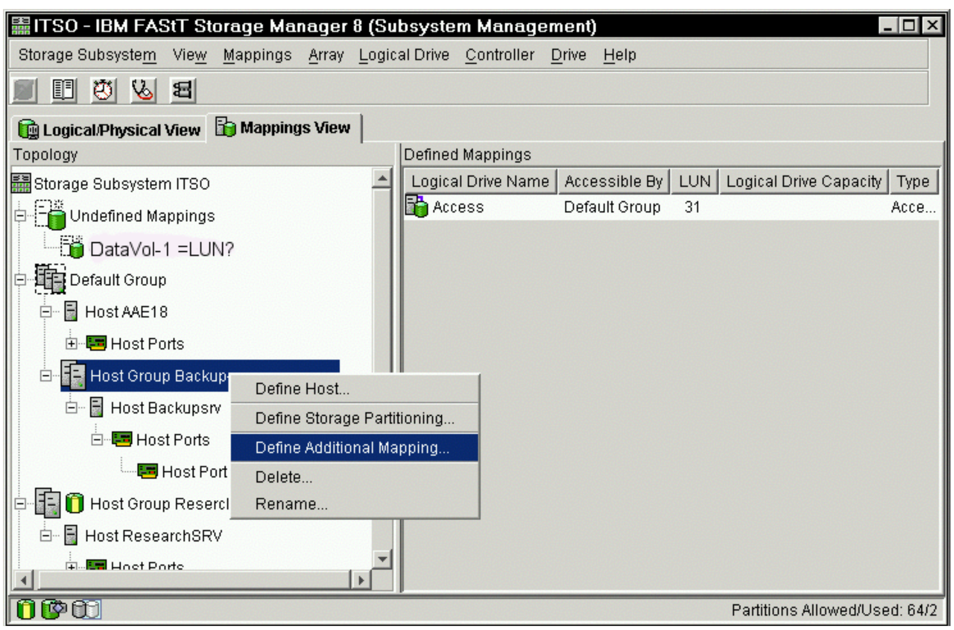

Figure 5-20 Define Additional Mapping

4. The Define Additional Mapping window (Figure 5-21) opens. Follow these steps:
   a. Select the FlashCopy drive.
   b. Select host or host group.
   c. Set the LUN number.
   d. Click **Add**.

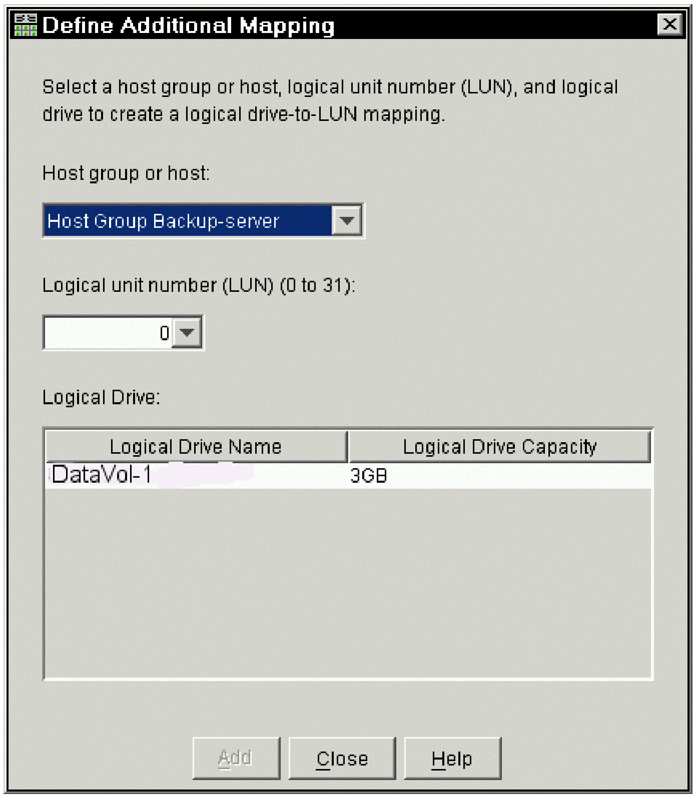

Figure 5-21 Define Additional Mapping window

In the example shown in Figure 5-20 and Figure 5-21, the FlashCopy logical disk “DataVol-1” is made available to the backup server host group. This enables the backup server to access and mount the disk as part of its own file system. It also allows a local backup to performed by the backup application software.

It is possible to map the FlashCopy logical disk to the same server that owns the base logical disk. However, note that the two logical disks, immediately after creating the FlashCopy, appear exactly the same (a block by block copy). Many operating systems do not tolerate seeing an exact duplicate volume. You may need to complete other steps before you can access it. The mapping is shown in Figure 5-22.

Important: If you use this FlashCopy on a regular basis (for example, for backup purposes), use the Disable FlashCopy and Re-create FlashCopy options to reuse the FlashCopy. Using these options preserves the existing mappings to the FlashCopy logical drive. For more information, see 5.1.2, “Disabling and recreating a FlashCopy logical drive” on page 141.

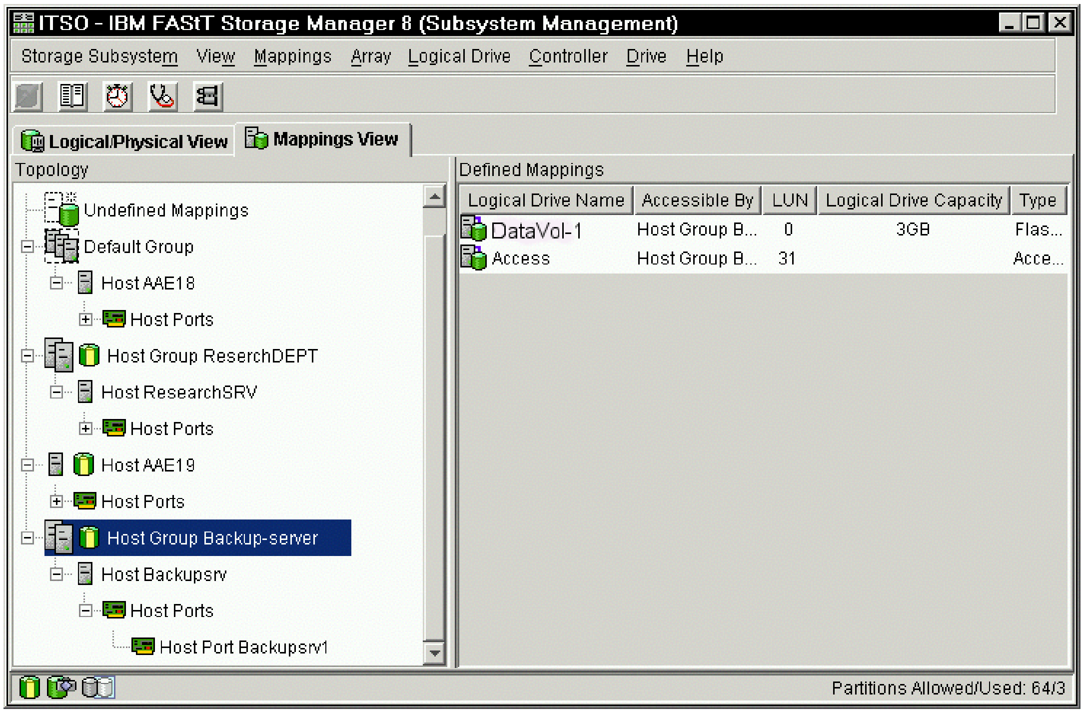

Figure 5-22 Mapped FlashCopy logical disk

Finally use specific operating system and host utilities to mount and use the mapped FlashCopy drive. The basic procedure is:

1. Run the host-based hot_add utility to register the FlashCopy logical drive with the host operating system.
2. Run the host-based SMdevices utility to associate the mapping between the physical device name and the logical drive name.
3. Mount the logical drive to the host.

For information on specific host operating system procedures, see Appendix A, "Additional host-specific instructions for FlashCopy logical drives" on page 269.

5.2.5 Viewing the FlashCopy drive status

The status of the FlashCopy logical drive can be determined by viewing the icons that will change depending of the state of the drive. The logical drive component property display is also useful in determining the state of the logical drives.

Use the FlashCopy Repository Logical Drive - Properties dialog to view the FlashCopy repository logical drive base and capacity properties. You may also use this dialog to specify the capacity percentage full and the action to be taken if the FlashCopy repository logical drive becomes full.

The progress of modification operations is displayed at the bottom of the dialog.

FlashCopy icon states

To view the FlashCopy icon, open the Storage Management Device Manager GUI Physical/Logical view. The icon states are described in Figure 5-23.

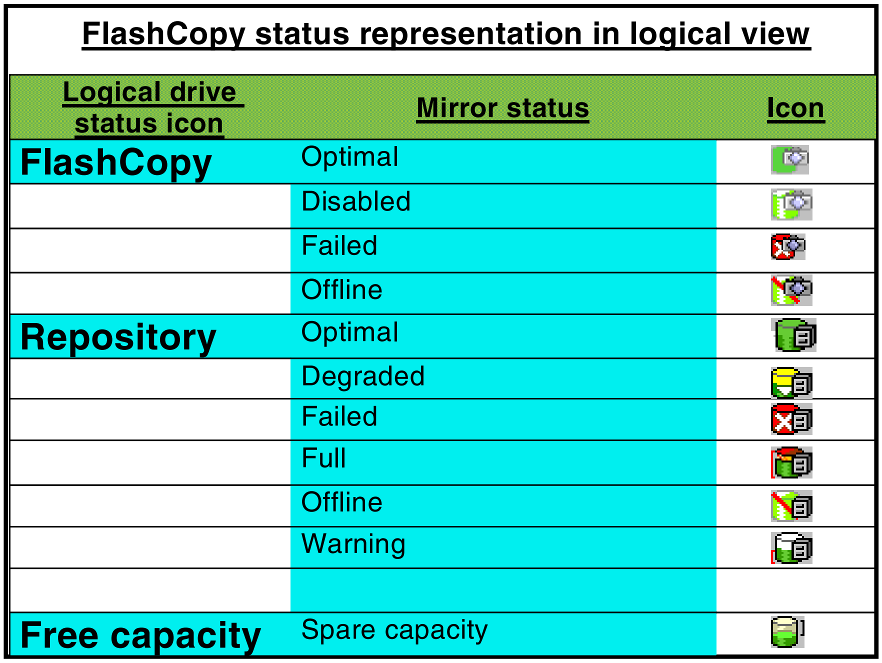

Figure 5-23 Status symbols

FlashCopy repository properties

To view the FlashCopy repository properties, complete the following procedure:

1. Select a FlashCopy repository logical drive in the Logical View of the Subsystem Management window.
2. Click **Logical Drive-> Properties**. Or you can right-click and select **Properties** as shown in Figure 5-24.

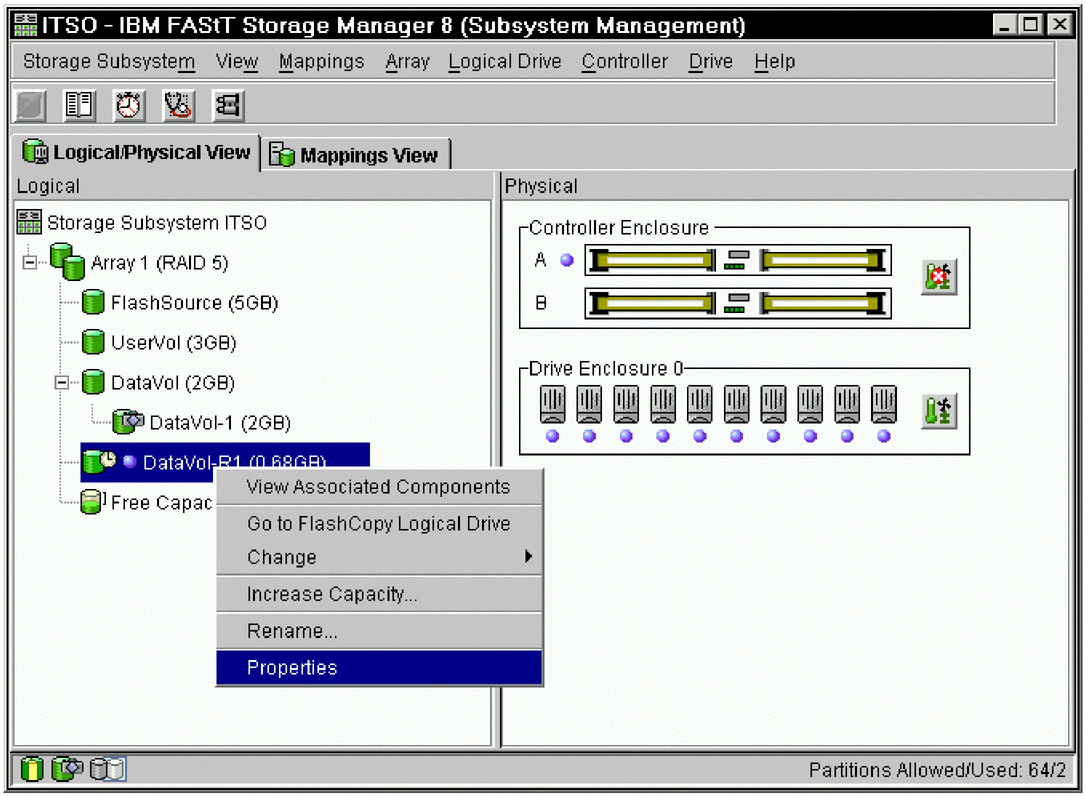

Figure 5-24 Repository logical drive properties

The FlashCopy Repository Logical Drive - Properties window opens as shown in Figure 5-25.

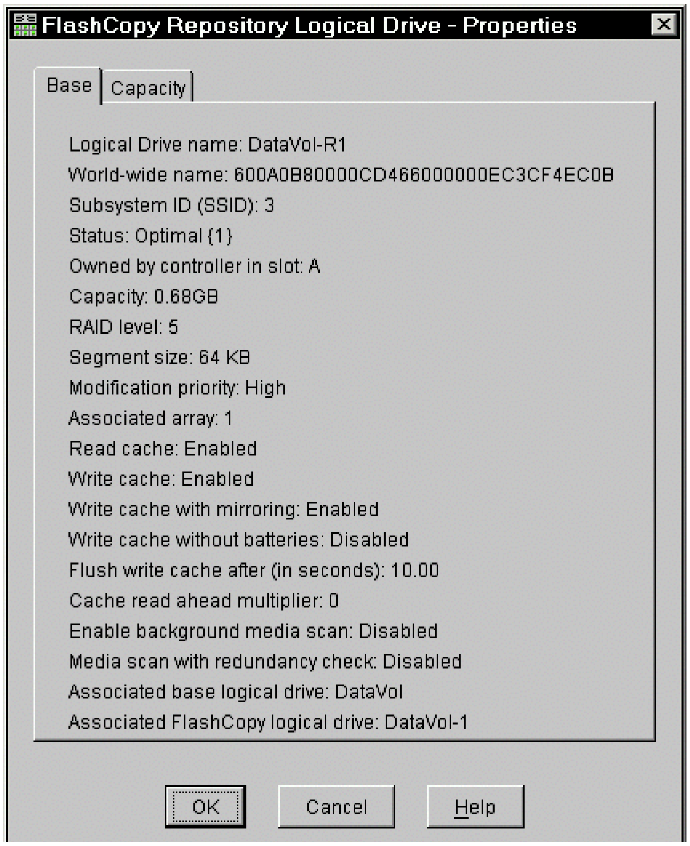

Figure 5-25 Base Repository Logical Drive Properties

The Base tab of the FlashCopy Repository Logical Drive - Properties dialog displays the following information for the selected FlashCopy repository logical drive:

– Logical Drive name
– World-wide name
– Status
– Controller ownership
– Capacity
– RAID level
– Modification priority
– Associated base logical drive
– Associated FlashCopy logical drive

3. Click the **Capacity** tab (Figure 5-26) to view or set the following FlashCopy repository logical drive capacity properties:

   – **FlashCopy repository logical drive percentage full**

   Allows a threshold level to be set for the FlashCopy repository logical drive capacity. Once the defined percentage is reached, a warning is issued indicating that the repository is nearing its capacity. The default percentage setting is 50% of the FlashCopy repository logical drives maximum capacity.

   Use the percent (%) full box to define the percentage at which a warning is issued.

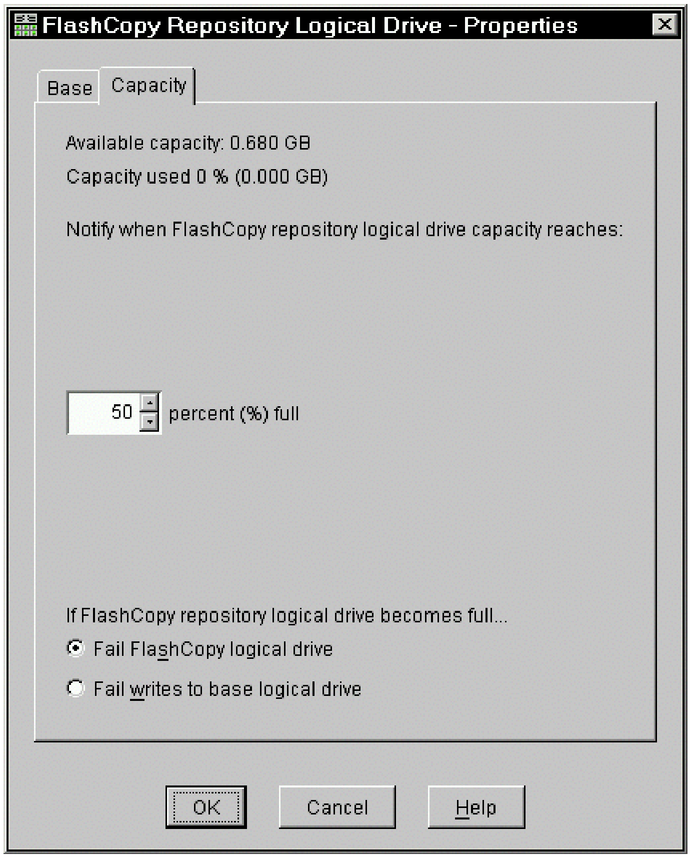

Figure 5-26 Repository drive warning threshold

– **Repository full policy**

Once a FlashCopy repository logical drive reaches its capacity and becomes full, one of the following actions occurs:

- Fail FlashCopy logical drive
- Fail writes to base logical drive

If the FlashCopy repository logical drive is set to *Fail FlashCopy logical drive* (Figure 5-26) when it becomes full, its data will not be recoverable and the FlashCopy cannot be accessed. The only available option (if this situation occurs) is to delete the FlashCopy logical drive or re-create the FlashCopy logical drive to create a new point-in-time image.

If the FlashCopy repository logical drive is set to *Fail writes to the base logical drive* (Figure 5-26), the data is recoverable. However, the FlashCopy repository logical drive capacity must be increased so writes to the base logical drive are not rejected. For more information, see 5.1.6, “Increasing the capacity of a FlashCopy repository logical drive” on page 147.

Viewing the progress of a modification operation

The progress bar at the bottom of the FlashCopy Repository Logical Drive Properties dialog displays the progress of an operation. You can view the progress of the following operations:

- Copyback
- Reconstruction
- Initialization
- Change RAID Level

- Dynamic Logical Drive Expansion
- Add Capacity
- Defragment
- Change Segment Size

Important: The Storage Management software cannot obtain progress information from the storage subsystem controllers if the network management connection to the controllers is down or if the storage subsystem is partially managed. For more information on a partially managed storage subsystem or an unresponsive controller or storage subsystem condition, see the Enterprise Management window online help.

5.2.6 Disabling a FlashCopy logical drive

If you no longer need a FlashCopy logical drive, you may want to disable it because, as long as a FlashCopy logical drive is enabled, your storage subsystem performance may be impacted by the copy-on-write activity to the associated FlashCopy repository logical drive. When you disable a FlashCopy logical drive, the copy-on-write activity stops.

If you disable the FlashCopy logical drive instead of deleting it, you can retain it and its associated repository. Then, when you need to create a different FlashCopy of the same base logical drive, you can use the re-create option to reuse a disabled FlashCopy. This takes less time than creating a new one.

When you disable a FlashCopy logical drive, note that:

- You cannot use that FlashCopy logical drive again until you use the re-create option on that logical drive.
- Only that FlashCopy logical drive is disabled. All other FlashCopy logical drives remain functional.

If you do not intend to re-create a FlashCopy, you can delete that FlashCopy logical drive instead of disabling it.

To disable a FlashCopy logical drive, follow these steps:

1. Select the FlashCopy logical drive. Right-click and select **Disable**, as shown in Figure 5-27.

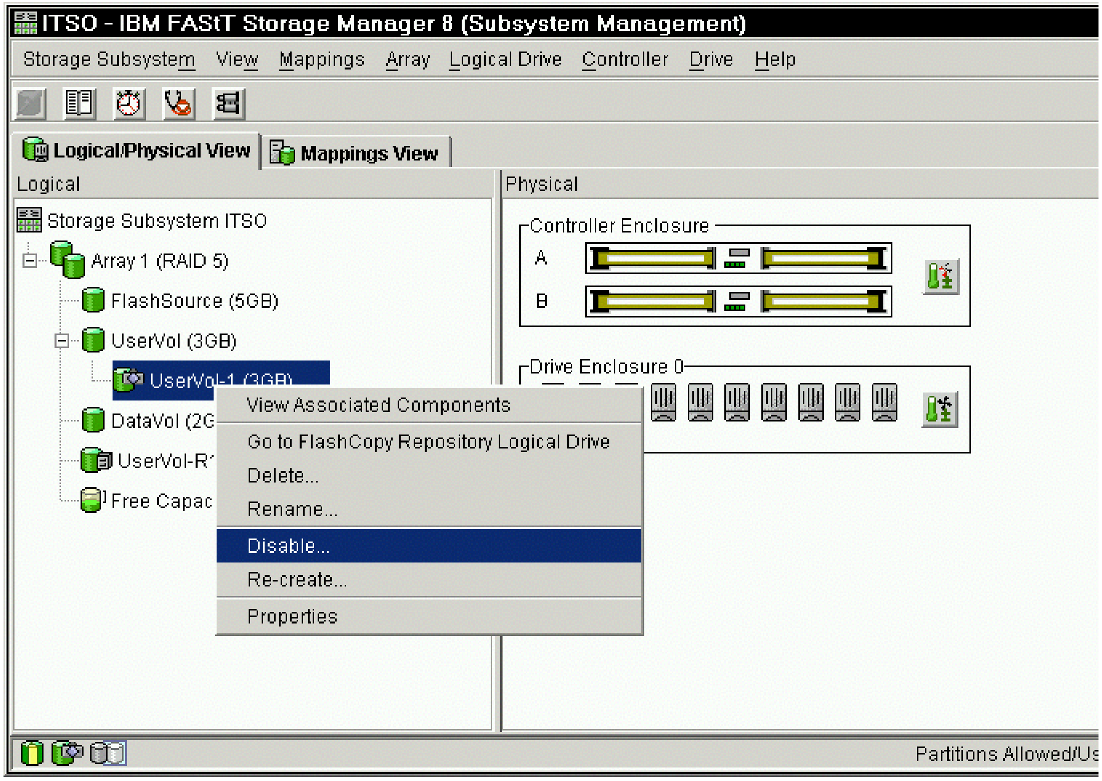

Figure 5-27 Choosing to disable the FlashCopy drive

2. The Disable FlashCopy Logical Drive confirmation window (Figure 5-28) opens. On this window, type `yes` and click **OK** to begin the disable operation.

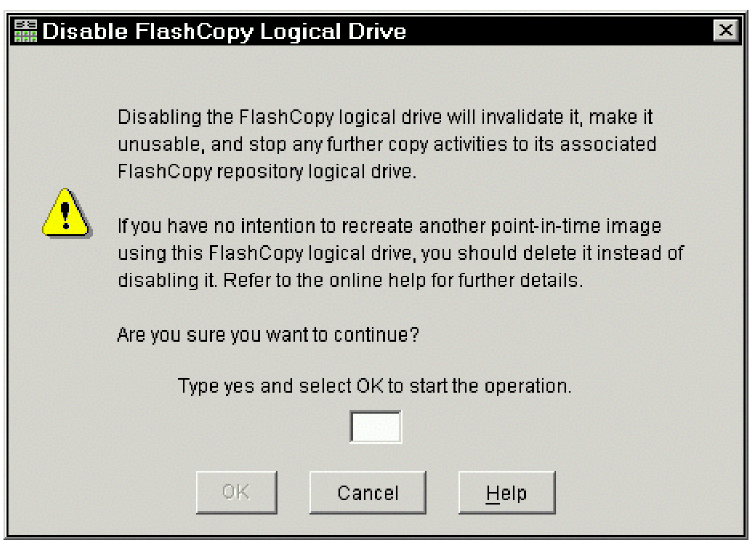

Figure 5-28 Disable FlashCopy confirmation window

The FlashCopy icon in the Physical/Logical View now appears as disabled, as shown in Figure 5-29.

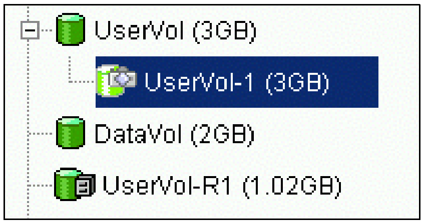

Figure 5-29 Icon showing the disabled FlashCopy logical drive

5.2.7 Re-creating a FlashCopy logical drive

Re-creating a FlashCopy logical drive takes less time than creating a new one. If you have a FlashCopy logical drive that you no longer need, instead of deleting it, you can reuse it (and its associated FlashCopy repository logical drive) to create a different FlashCopy logical drive of the same base logical drive.

When you re-create a FlashCopy logical drive, note the following points:

- The FlashCopy logical drive must be either in an optimal or a disabled state.
- All copy-on-write data on the FlashCopy repository logical drive is deleted.
- FlashCopy and FlashCopy repository logical drive parameters remain the same as the previously disabled FlashCopy logical drive and its associated FlashCopy repository logical drive. After the FlashCopy logical drive is re-created, you can change parameters on the FlashCopy repository logical drive through the appropriate menu options.
- The original names for the FlashCopy and FlashCopy repository logical drives are retained. You can change these names after the re-create option completes.
- When using this option, the previously configured FlashCopy name, parameters, and FlashCopy repository logical drive are used.

To recreate a FlashCopy drive, follow these steps:

1. Select the FlashCopy logical drive. Right-click and select **Re-Create**, as shown in Figure 5-30.
2. The Re-create FlashCopy Logical Drive dialog opens. Type `Yes` and click **OK**.

   The FlashCopy logical drive is disabled and re-created (if it had not previously been disabled) and displays in the Logical View in an Optimal state. The creation timestamp shown on the FlashCopy Logical Drive Properties dialog is updated to reflect the new point-in-time image. Copy-on-write activity resumes to the associated FlashCopy repository logical drive

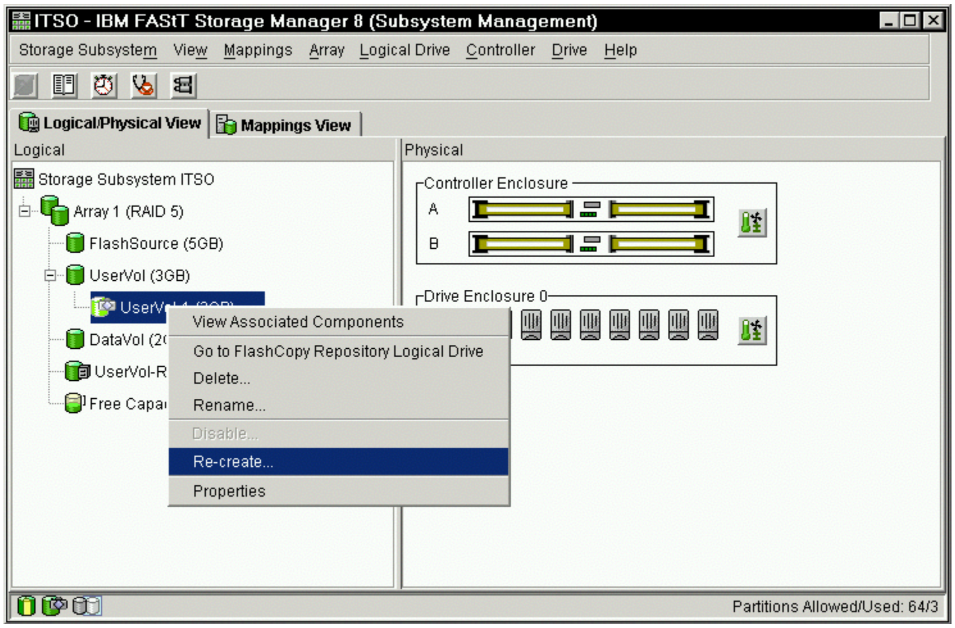

Figure 5-30 Recreating a FlashCopy

> **Important:**
>
> - To use the Re-create option, the FlashCopy logical drive must be in either an optimal state or a disabled state.
> - If the FlashCopy logical drive is in an *optimal* state, it is first disabled and then re-created. This invalidates the current FlashCopy.

5.2.8 Resizing a FlashCopy repository drive

Use this option to increase the storage capacity of an existing FlashCopy repository logical drive. Typically, this option is used when a warning is received that the FlashCopy repository logical drive is in danger of becoming full.

You can achieve an increase in storage capacity by:

- Using free capacity available on the array of the FlashCopy repository logical drive.
- Adding unconfigured capacity (in the form of unused drives) to the array of the FlashCopy repository logical drive. Use this option when no free capacity exists on the array.

> **Important:** A maximum of two drives may be added at one time to increase FlashCopy repository logical drive capacity.

A FlashCopy repository logical drive's storage capacity cannot be increased if:

- One or more hot spare drives are in use in the logical drive.
- The logical drive has a *Non-Optimal* status.
- Any logical drive in the array is in any state of modification.
- The controller that owns this logical drive is in the process of adding capacity to another logical drive (each controller can add capacity to only one logical drive at a time).

- No free capacity exists in the array.
- No unconfigured capacity (in the form of drives) is available to add to the array.

To resize a FlashCopy repository drive, follow these steps:

1. Highlight a FlashCopy repository logical drive in the Logical View of the Subsystem Management window.
2. Click **Logical Drive-> Increase Capacity**. Or right-click and select **Increase Capacity**.

   > **Important:** If no free capacity or unconfigured capacity is available, the Increase Capacity option is not available.

   The Increase Repository Capacity dialog opens. You can see the FlashCopy repository logical drive name, the associated FlashCopy logical drive name, the associated base logical drive name, current capacity, and amount of free capacity available for the selected repository. If free capacity is available, the maximum free space is shown in the Increase capacity by box.

   If there is no free capacity on the array, the free space that is shown in the Increase capacity by box is 0. Drives must be added to create free capacity on the array of the standard logical drive.

3. Use one of the following two methods to increase capacity:

   – Increase FlashCopy repository logical drive capacity using free capacity on the array of the FlashCopy repository logical drive:

     i. Accept the final capacity increase or use the Increase capacity by box to adjust the capacity. Click **OK**.

     ii. A confirmation dialog is displayed. Type `Yes` and click **OK** to continue.

     The Logical View is updated. The FlashCopy repository logical drive with its capacity increased shows a status of *Operation in Progress*, together with its original capacity and the total capacity being added.

     iii. In addition, the Free Capacity node involved shows a reduction in capacity. If all of the free capacity is used to increase the logical drive size, then the Free Capacity node involved is removed from the Logical View.

   – Increase FlashCopy repository logical drive capacity by adding unconfigured capacity (drives) to the array of the FlashCopy repository logical drive:

     i. If no unassigned drives are available and empty slots in the drive enclosures are available, insert new drives.

     If no unassigned drives are available, and there are no empty slots available in the drive enclosures, install another drive enclosure and additional drives.

     ii. Select **Add Drives**.

     The Increase Repository Capacity - Add Free Capacity dialog opens. Enclosure, slot, and usable capacity details for the available free drives are displayed.

     > **Note:** The drives that are displayed have a capacity that is either the same size, or larger than, the capacity of the drives already being used by the array.

     iii. Select a single drive, or two drives, to be added:

     - Press Ctrl and click to select the nonadjacent drives.
     - Press Shift and click to select the adjacent drives.

iv. Select **Add**.

The Add Free Capacity dialog is closed. Check the **Drives to add [enclosure, slot]** area to ensure the correct drives are added.

v. Accept the final capacity or use the Increase capacity by box to adjust the capacity.

vi. Select **OK**.

vii. A confirmation dialog is displayed. Type `Yes` to confirm the operation, and click **OK** to continue.

The Logical View is updated. The FlashCopy repository logical drive having its capacity increased shows a status of *Operation in Progress*, together with its original capacity and the total capacity being added. In addition, the Free Capacity node involved in the increase shows a reduction in capacity.

If all of the free capacity is used to increase the logical drives size, then the Free Capacity node involved is removed from the Logical View.

If a Free Capacity node did not exist prior to the addition of capacity and not all of the capacity that is added will be used to increase the FlashCopy repository logical drives capacity, a new Free Capacity node is created and displayed in the Logical View.

Unassigned drives (unconfigured capacity) added to increase the FlashCopy repository logical drives capacity changes in the Physical View to assigned drives, and becomes associated to the array of the FlashCopy repository logical drive.

4. View the progress of the increase the capacity process. Highlight the FlashCopy repository logical drive. Click **Logical Drive-> Properties**, or right-click and select **Properties**.

   The FlashCopy Repository Logical Drive - Properties dialog opens. A progress bar at the bottom of the dialog indicates the status of the capacity increase.

5.2.9 Deleting a FlashCopy drive

Use this option to delete a FlashCopy logical drive that is no longer needed for backup or application testing purposes. This option results in an increase of free capacity in the array or additional unconfigured capacity.

Attention:

- Deleting a logical drive causes loss of all data on the logical drive. Back up the data and stop all I/O before performing this operation, if necessary.
- If a file system is mounted on the logical drive, unmount it before you attempt this operation.
- Deleting a base logical drive automatically deletes the associated FlashCopy logical drive(s) and FlashCopy repository logical drive(s).
- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy repository logical drive.
- You cannot delete a FlashCopy repository logical drive using the Delete Logical Drive option, but you can using the **`delete`** command in the Script Editor or CLI. See the Enterprise Management window online help system for more information on using the **`delete`** command.

To delete a FlashCopy drive, follow these steps:

1. Select the FlashCopy logical drive in the Logical View.

2. Click **Logical Drive-> Delete**, or right-click and select **Delete**, as shown in Figure 5-31.

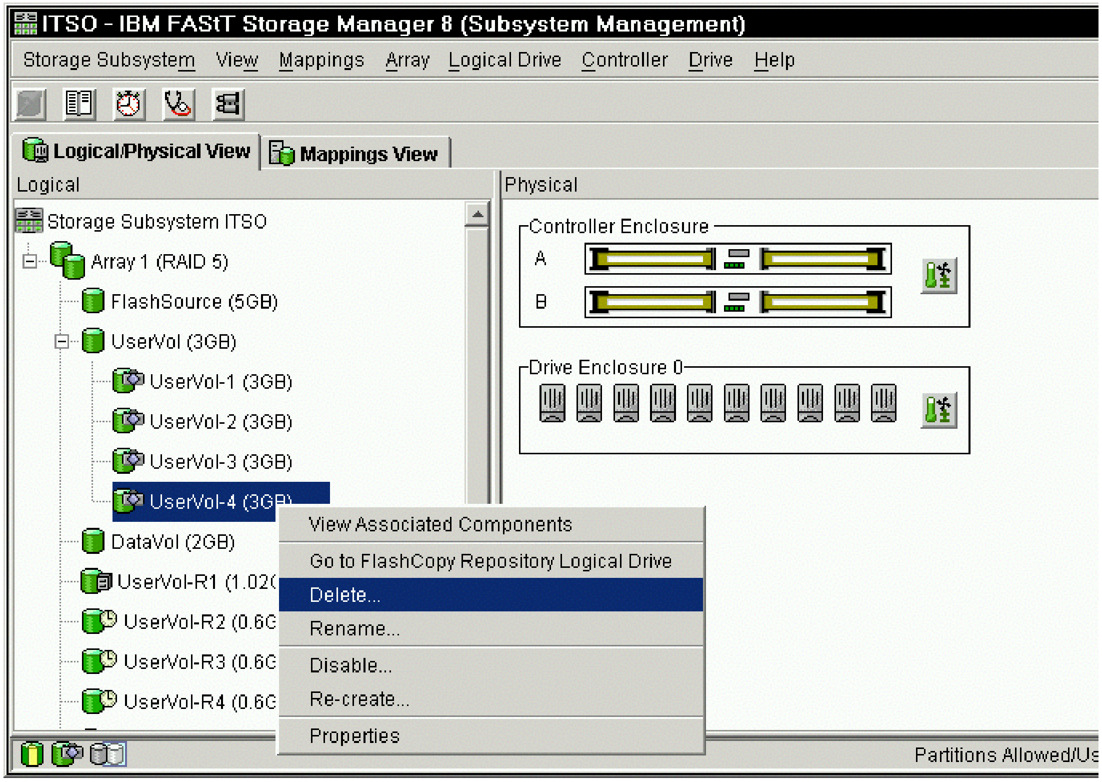

Figure 5-31 Deleting the FlashCopy logical drive

3. The Delete FlashCopy Logical Drive dialog opens as shown in Figure 5-32. Type `Yes` and click **OK**.

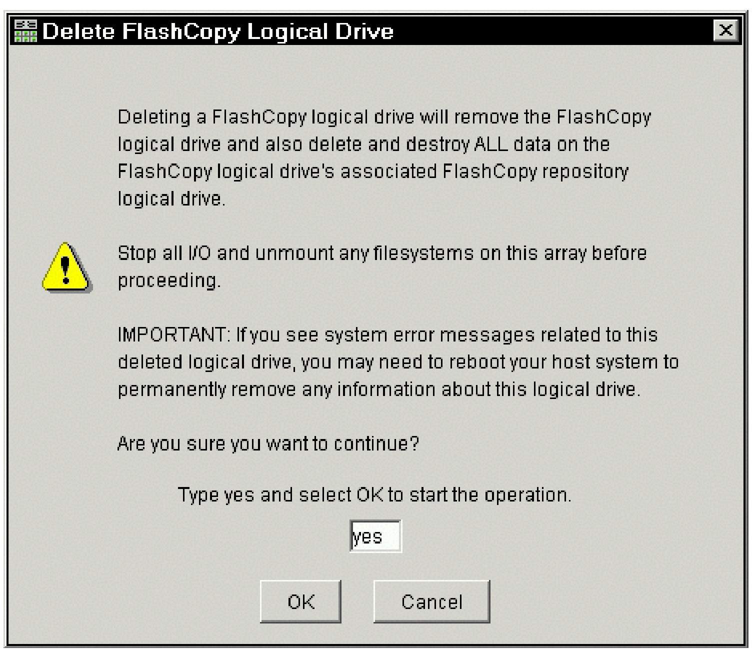

Figure 5-32 Delete FlashCopy Logical Drive dialog

The FlashCopy logical drive and FlashCopy repository logical drive are deleted. Any data on them is destroyed.

5.3 Scenarios using scripts

This section presents three scenarios where we implement the FlashCopy feature controlled only by scripts. These scripts can then be integrated, for example, in your backup scenario.

We describe the three scenarios here. You can use them as a guideline to develop scripts that fit your needs exactly:

- Windows 2000 Basic Disks, FlashCopy logical drive mapped to a second host
- Windows 2000 Dynamic Disks, FlashCopy logical drive mapped to a second host
- Linux, FlashCopy logical drive mapped to the same host as the source logical drive

In “Command line interface” on page 60, we present SMcli, which we use in the scripts to control the FAStT Storage Server.

For each scenario, we provide two scripts:

- One to recreate an already defined FlashCopy
- One to disable the FlashCopy

The scripts ensure that the attached operating systems are handling correctly the FlashCopied volumes. For the two Windows scenarios, we provide an initial script that establishes the first FlashCopy.

We tried to use environment variables where possible to make it easy for you to adapt the scripts to your needs.

A possible implementation for a database server is outlined in these steps:

1. Pause the productive database on server A.
2. Make a FlashCopy of the database volume and map it to server B (this is done with the recreate script).
3. Resume database on server A.
4. Back up the database off server B.
5. Disable the FlashCopy again (done with the disable script).

Now the two scripts can be integrated to run just before and after the backup process itself.

For the two Windows scenarios, we use several command line tools that are provided by Microsoft in the Resource Kit. The Linux operating system already contains the needed tools:

- **Diskpart**: A command line version of the Disk Administrator.
- **Sleep**: A tool that waits a given time.
- **Remote Command Service**: A service installed on a Windows 2000 host that allows remote execution of scripts or commands (similar to `rsh` on UNIX systems).
- **rcmd**: The counterpart for the Remote Command Service
- **netsvc**: A command line tool that allows the control of services on a remote system.

You can find more information on the Microsoft Windows 2000 Resource Kit at:

http://www.microsoft.com/windows2000/techinfo/reskit/tools

There are two known issues while handling with SCSI devices in Windows 2000. They are documented in the Microsoft Knowledge Base:

- Restart Required Message When You Scan a Device in Device Manager (Q306185):

  `http://support.microsoft.com/default.aspx?scid=kb;en-us;Q306185`

- Problems When You Remove SCSI Devices Without Notifying Windows (Q318327):

  `http://support.microsoft.com/default.aspx?scid=kb;en-us;Q318327`

We didn't discover any problems during our tests, but they may occur in your setup.

There is another known issue with a FlashCopy functionality that is described in IBM Retain Tip H175907. In Windows, if an open file handle to the FlashCopy logical drive exists during a recreation, Windows still displays the old data content. Because of this behavior in our scripts, the drives are not only disabled, but the mapping is also removed, so the disk device itself disappears for Windows. After the re-creation, the mapping is re-established and the disk can be brought online again.

5.3.1 Windows 2000 basic disks

In this scenario, we use the following setup:

- Host AAE18 is running Windows 2000. The logical drive *FlashSource* is mapped from the Storage Subsystem *ITSO*. The disk is used as a basic disk and drive letter E: is assigned to it.
- The host AAE18 has the Remote Command Service installed. It is set to manual startup and is only started if needed by our scripts.
- The FAStT host software is installed in the default location on host AAE18.
- Host AAE19 is also running Windows 2000. The FlashCopy logical drive *FlashTarget*, which is a point-in-time copy of the logical drive *FlashSource*, will be used by this host as a basic disk with drive letter F:
- The host AAE19 has the FAStT host software installed in the default location.
- The host AAE19 has installed *rcmd*, *diskpart*, and *netsvc* from the Microsoft Windows 2000 Resource Kit.
- Both host systems are either in the same domain or have the same account. This allows the remote execution of commands using the Remote Command Service.
- The scripts are executed in AAE19.

We discuss the scripts in detail here. You can find the complete scripts in Appendix D, “FlashCopy example scripts” on page 327.

You can do the initial FlashCopy manually or with the initial script provided in Appendix D, “FlashCopy example scripts” on page 327. Because we only run this script once, we do not explain it here in detail. The main difference is the command for the FAStT Storage Server to create a FlashCopy with new parameters compared to recreating an already defined FlashCopy. The commands are explained in Appendix C, “Command line reference” on page 317.

To mount and unmount the disk properly in Windows 2000, we use **`mountvol`**, which is included in the operating system. The **`mountvol`** command allows us to assign and unassign a drive letter to a specific disk or volume in Windows 2000.

Initial script

You can perform all the steps in this script manually. Since this process is used only once, we only include this script to show the scripting facilities available for creating FlashCopy logical drives. We do not explain other parts of the script, apart from the parts for creating the actual FlashCopy logical drive and its repository.

```
rem Create FlashCopy Script with Repository on existing array.
set CREATE_FlashCopy_SCRIPT="on error stop; create FlashCopyLogicalDrive
baseLogicalDrive=%SOURCE_DRIVE_NAME% repositoryArray=%REPOSITORY_ARRAY%
userLabel=%TARGET_DRIVE_NAME% repositoryUserLabel=%REPOSITORY_NAME%
repositoryRaidLevel=%REPOSITORY_RAID_LEVEL%;"
```

The CREATE_FlashCopy_SCRIPT defines a first snapshot of the base logical drive BASE_LOGICAL_DRIVE chosen with the parameter baseLogicalDrive=name. The only other required parameter is the information for storing the repository. In our example, it is stored on the array REPOSITORY_ARRAY, which is already defined.

It is also possible to use the parameter repositoryDriveCount=*n*, where *n* is the number of disk drives used for a new array. The storage software chooses the drives automatically. And it is also possible to use the parameter repositoryDrives=*n1*,*m1*,n2,m2 and define a list of drives to be used for the new array, where *n1* is the enclosure number and *m1* the slot number of the first disk drive to use. The setting repositoryRaidLevel=n allows you to define the raid level of the array that will hold the repository logical drive.

All other parameters are optional. If left empty, they will use the default values.

The userLabel=name is the name of the FlashCopy logical drive that will be created. Because we use this name later in our scripts, we define it here. Also, to distinguish later between several different FlashCopy definitions, we label the repository with the parameter repositoryUserLabel=name.

The settings you can choose in the Advanced tab of the manual method can also be defined within the command line. We explain them shortly in the following list:

- **repositoryPercentofBase=n**: Defines the size of the repository, in percent, of the base logical drive. The default is 20%.
- **repositoryFullPolicy=failBaseWrites** or **failFlashCopy**: Defines in the policy what happens if the repository is full. You can either fail the FlashCopy logical drive, which is the default, or fail the base logical drive.
- **repositoryTresholdPercent=n**: The level, in percent, to which the repository must be filled before a warning is generated.

The command line created here is then passed to the storage manager command line tool and sent to the FAStT Storage Server for execution. The remaining parts of the script handle the Windows specific tasks to bring the FlashCopy logical drive online as a normal disk in the operating system.

Recreate script

The FlashCopy is defined on the FAStT Storage Server, but is momentarily disabled because it is not used. To activate the FlashCopy again, map it and bring it online on host AAE18. We provide the following script.

In the first part of the script, we define all environment variables that are needed:

```
rem Path to the SMFlashCopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMFlashCopyassist"
```

```
rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path tho the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem Host name of the machine FlashCopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the FlashCopy source drive.
set SOURCE_FlashCopy_DRIVE_LETTER="E:"

rem The drive letter of the FlashCopy target drive.
set TARGET_FlashCopy_DRIVE_LETTER=F:

rem Name of the target logical drive to hold the FlashCopy of on FAStT.
set TARGET_DRIVE_NAME=\"FlashTarget\"

rem Volume ID of the FlashCopy target drive.
set TARGET_FlashCopy_DRIVE_VOLUME_ID=\\?\Volume{062d708e-732b-11d6-ab27-0003477b412d}\

rem LUN number to map the target drive to the host.
set LUN_NUMBER=1
```

The only non-trivial parameter in this list is TARGET_FlashCopy_DRIVE_VOLUME_ID. This value varies for each logical drive on every host. It is also called the *disk signature*. To extract this information, the logical drive FlashTarget must be online and assigned a drive letter, which is F: in our example. Then use the following command:

```
mountvol F: /l
```

This outputs the volume ID, which must then be put in the script.

The parameter LUN_NUMBER determines under which logical unit number the logical drive FlashTarget is mapped to the host system AAE19. This must be a unique value that is not used by another drive mapping from the same FAStT to this host.

The following lines generate the command line input we use later in the script from the variables defined earlier:

```
rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FlashCopy_DRIVE_LETTER%"

rem Recreate FlashCopy script.
set RECREATE_FlashCopy_SCRIPT="on error stop; recreateFlashCopy logicalDrive
[%TARGET_DRIVE_NAME%];"

rem Create mapping script to map target drive to host %HOST_NAME%.
```

```
set MAP_FlashCopy_SCRIPT="on error stop; create mapping logicalDrive=%TARGET_DRIVE_NAME%
logicalUnitNumber=%LUN_NUMBER% host=%HOST_NAME%;"
```

The REMOTE_SCRIPT contains the command line we execute on the remote system AAE18.

The RECREATE_FlashCopy_SCRIPT is passed to the FAST and executed there. The command recreates the previously defined FlashCopy logical drive TARGET_DRIVE_NAME. The **on error stop** command assures that the script exits with an error if the FlashCopy cannot be recreated. The FlashCopy logical drive must be disabled before you can recreate it. It is not possible to issue the recreateFlashCopy command to an active FlashCopy logical drive. The are no other options available.

The MAP_FlashCopy_SCRIPT finally maps the newly recreated FlashCopy logical drive back with the logical unit number LUN_NUMBER to the host HOST_NAME. If the logical drive would map to a host group rather than a host, the line would be:

```
set MAP_FlashCopy_SCRIPT="on error stop; create mapping logicalDrive=%TARGET_DRIVE_NAME%
logicalUnitNumber=%LUN_NUMBER% hostgroup=%HOST_GROUP_NAME%;"
```

There are no other options available for the **create mapping** command.

Now we start with the execution of the first commands:

```
echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"
```

This starts the Remote Command Service on the remote host AAE18 and allows the usage of the **rcmd** tool now. To start the service, you must have the appropriate rights on the remote system, for example, an account with domain administrator rights.

At this point, all I/O going to the source logical drive should be stopped. Otherwise, the FlashCopy may not be consistent and even unusable.

```
echo Flushing drive %SOURCE_FlashCopy_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%
```

The command we generated in REMOTE_SCRIPT is executed on the remote host. Using the SMFlashCopyassist tool provided with the FAStT host software, we ensure that the operating system writes all data buffered for the drive E: to the storage subsystem.

```
echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"
```

Finally we stop the Remote Command Service to avoid any misusage.

```
echo Recreating FlashCopy logical drive %TARGET_FlashCopy_DRIVE%.
call %SM_CLI% -n %FAStT_NAME% -c %RECREATE_FlashCopy_SCRIPT%
if errorlevel 1 goto recreate_FlashCopy_fail_exit

echo Mapping flashcopy logical drive %TARGET_FlashCopy_DRIVE%.
call %SM_CLI% -n %FAStT_NAME% -c %MAP_FlashCopy_SCRIPT%
if errorlevel 1 goto mapping_fail_exit
```

The two lines we defined are now used with the storage manager command line tool and sent to the storage subsystem FAStT_NAME. Since we are using in-band management, we use the `-n` switch. If you plan to use out-of-band management via a network, you don't need any switch. If FAStT_NAME is properly defined in your DNS, the command would be:

```
call %SM_CLI% %FAStT_NAME% -c %RECREATE_FlashCopy_SCRIPT%
```

The commands recreate the FlashCopy logical drive and map it back to the host HOST_NAME.

At this time, the I/O to the source logical drive can be resumed.

We now have to bring the disk back online in Windows 2000:

```
echo Bus rescan using hot_add.
%HOTADD%

echo Wait 5 seconds.
sleep 5

echo Mount the volume under old drive letter %TARGET_FlashCopy_DRIVE_LETTER%.
mountvol %TARGET_FlashCopy_DRIVE_LETTER% %TARGET_FlashCopy_DRIVE_VOLUME_ID%
```

The hot_add utility, which is part of the FAStT Utility package, rescans all SCSI busses for new devices and finds the remapped FlashCopy logical drive. The time delay of five seconds ensures that Windows 2000 can include the device properly in the device manager.

Finally we use again the **mountvol** command to bring recreated FlashCopy logical drive online with the same drive letter TARGET_DRIVE_LETTER as before. This is the step where we again use the volume ID, also known as the disk signature, we extracted before.

The rest of the script only deals with error handling of any of the commands that exited abnormally and the deletion of all environment variables we used throughout the script.

Disable script

The FlashCopy is valid. The FlashCopy logical drive is mapped and is online on host *AAE18*. Because of performance reasons, the FlashCopy logical drive *FlashTarget* should be disabled when not in use.

The scripts start again with all the environment variables needed in the script. Since they are the same as in the recreate script, they are not discussed here again.

```
rem Path to the SMFlashCopyassist utility.
set SNAPASSIST="c:\Program Files\IBM_FAStT\Util\SMFlashCopyassist"

rem Path to the SMFlashCopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMFlashCopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem Host name of the machine FlashCopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem FAStT name of the machine executing this script.
```

```
set FAStT_NAME="ITSO"

rem The drive letter of the FlashCopy source drive.
set SOURCE_FlashCopy_DRIVE_LETTER="E:"

rem The drive letter of the FlashCopy target drive.
set TARGET_FlashCopy_DRIVE_LETTER="F:"

rem Name of the target logical drive to hold the FlashCopy of on FAStT.
set TARGET_DRIVE_NAME=\"FlashTarget\"
```

The main difference is that we don’t need the volume ID of the FlashCopy logical drive initial script.

Again we prepare the command lines that will be used with the Remote Command Service and the FAStT command line tool:

```
rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FlashCopy_DRIVE_LETTER%"

rem Disable FlashCopy script.
set DISABLE_FlashCopy_SCRIPT="on error stop; disableFlashCopy logicalDrive
[%TARGET_DRIVE_NAME%];"

rem Unmapping script to map target drive to host %HOST_NAME%.
set MAP_FlashCopy_SCRIPT="on error stop; delete mapping logicalDrive=%TARGET_DRIVE_NAME%
host=%HOST_NAME% ;“
```

The REMOTE_SCRIPT is the command line we execute on the remote host AAE19.

The DISABLE_FlashCopy_SCRIPT is passed to the FAST and executed there. The command disables the active FlashCopy logical drive TARGET_DRIVE_NAME. The **on error stop** command assures that the script exits with an error if the FlashCopy cannot be disabled. The FlashCopy logical drive must be active before it can be disabled. No other options are available.

The MAP_FlashCopy_SCRIPT finally deletes the mapping of the FlashCopy logical drive. The reason for not only disabling, but also unmapping the logical drive, is that there is no mechanism to properly unmount or take offline a device in Windows 2000 through a command line. If only the disk is unmounted (drive letter unassigned) and the FlashCopy on the FAStT disabled, Windows still tries to access the physical device and generate a large number of event log entries.

As in the recreate script, we flush all file system buffer of the operating system on the remote system using the Remote Command Service:

```
echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% “Remote Command Service"

echo Flushing drive %SOURCE_FlashCopy_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% “Remote Command Service"
```

Because the FlashCopy logical drive is online in the operating system, we also must flush the buffers of this disk. Otherwise, Windows may try to save buffers for this disk while the FlashCopy logical drive is already disabled.

```
echo Flushing target drive %TARGET_FlashCopy_DRIVE_LETTER%.
%SNAPASSIST% -f %TARGET_FlashCopy_DRIVE_LETTER%
```

```
if errorlevel 1 goto flush_target_fail_exit
```

After flushing the buffers for this disk, we can unmount the disk using the **mountvol** tool and removing the drive letter *F:*

```
echo Unmount target drive %TARGET_FlashCopy_DRIVE_LETTER%.
mountvol %TARGET_FlashCopy_DRIVE_LETTER% /d
if errorlevel 1 goto unmount_fail_exit
```

The next two commands execute our prepared scripts on the FAStT Storage Server and delete the mapping for the FlashCopy logical drive. Now that Windows 2000 cannot access this device anymore, the FlashCopy logical drive can safely be disabled.

```
echo Unmapping FlashCopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -c %MAP_FlashCopy_SCRIPT%
if errorlevel 1 goto mapping_fail_exit

echo Disabling FlashCopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -c %DISABLE_FlashCopy_SCRIPT%
if errorlevel 1 goto disable_FlashCopy_fail_exit
```

To clean up the device manager entries in Windows 2000 we invoke the hot_add tool:

```
echo Bus rescan to remove stale drives.
%HOTADD%
```

Unfortunately Windows displays a message stating that a device was removed unsafely, which is not logged in the event log. There is no possibility to disable this warning message. As stated earlier, this is caused by the lack of a command line tool that allows the safe removal of devices.

The remaining parts of the script handle error handling and post cleanup.

5.3.2 Windows 2000 dynamic disks

In this scenario, we use the following setup:

- Host AAE18 is running Windows 2000. The logical drives *FlashSource-1, FlashSource-2* and *FlashSource-3* are mapped from the Storage Subsystem *ITSO*. The disks are used as dynamic disks. A spanned volume is defined on these three disks, and drive letter *E:* is assigned to it.
- The host AAE18 has the Remote Command Service installed, but it is set to manual startup and will only be started if needed by our scripts.
- The FAStT host software is installed in the default location on host AAE18.
- Host AAE19 is also running Windows 2000. The FlashCopy logical drive *FlashTarget-1*, *FlashTarget-2* and *FlashTarget-3*, which are point-in-time copies of the logical drives *FlashSource-1*, *FlashSource-2*, and *FlashSource-3*, will be used by this host again as a spanned volume on dynamic disks with drive letter *F:*.
- The host AAE19 has the FAStT host software installed in the default location.
- The host AAE19 has installed *rcmd*, *sleep*, *diskpart* and *netsvc* from the Microsoft Windows 2000 Resource Kit.

Both host system are either in the same domain or have the same account that allows the remote execution of commands using the Remote Command Service.

If you import the dynamic disk configuration from a foreign Windows 2000 system, you cannot re-import a changed configuration from this host again. This is a known problem and is documented in the Microsoft Knowledge Base, under “Internal Error - Disk Group Exists and Is Imported” Error Message While Importing Foreign Disk (Q260113).

We discuss the scripts in detail here. You can find the entire script in Appendix D, “FlashCopy example scripts” on page 327.

The main difference between the scripts dealing with dynamic disks and the scripts explained in 5.3.1, “Windows 2000 basic disks” on page 177, are the handling of the disks in Windows and the usage of the storage manager command line tool. Since we have to handle three separate FlashCopies in this scenario, the commands for the FAStT Storage Server are stored in a separate file.

Initial script

Again we provide an initial script that creates the first FlashCopies of our three base logical drives and imports this disk configuration to host *AAE18*. Since we deal with dynamic disks now, it is not sufficient to only rescan the disk configuration using the diskpart tool. To import a foreign volume consisting of several dynamic disks, one of those must be first imported. Than we can set the volume residing on this disk online and assign a drive letter. Diskpart can only be used interactively or by a file holding the commands.

The command lines we use to script the FAStT Storage Server are stored in a separate file called *initial.cli*, which is used by the storage manager command line tool.

Recreate script

The FlashCopies are defined on the FAStT Storage Server but momentarily disabled because they are not used. To activate the FlashCopies again, map them and bring them online on host *AAE18.* We explain the script here.

As with the scripts for the basic disks, we define all needed environment variables first:

```
rem Path to the SMFlashCopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMFlashCopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\scripts\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Path to the diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Pathto the script for diskpart.
set DISKPART_SCRIPT="c:\scripts\recreate.dp"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the FlashCopy source drive.
```

```
set SOURCE_FlashCopy_DRIVE_LETTER="E:"
```

The variable DISKPART_SCRIPT points to the file that we are using with the diskpart tool. Diskpart can only be used with a separate script file or interactively, which does not suit our needs.

The command line generation part of this script is relatively short. The first command is again the command line we execute on the remote system. The second parameter points to another file *recreate.cli*. This file stores all the commands we pass to the FAStT Storage Server for execution. Because we have to make of FlashCopy of three logical drives, we decided to use a command line file for the storage manager command line utility instead of generating each command in the script itself as we did for the scenario with the basic disks:

```
rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f%SOURCE_FlashCopy_DRIVE_LETTER%"

rem Recreate FlashCopy script.
set RECREATE_FlashCopy_SCRIPT="c:\temp\scripts\recreate.cli"
```

At this point, the I/O going to the base logical drive should be stopped. We then flush the operating system file buffers for the base logical drive:

```
echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% “Remote Command Service"

echo Flushing drive %SOURCE_FlashCopy_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% “Remote Command Service"
```

Now that we have assured that the base logical drive is in a consistent state, we can FlashCopy the three logical drives.

```
echo Recreating FlashCopy logical drive and remapping.
call %SM_CLI% -n %FAStT_NAME% -f %RECREATE_FlashCopy_SCRIPT%
if errorlevel 1 goto recreate_FlashCopy_fail_exit
```

As you can see, one option changed. In the example, we used the -c switch, which expects a FAStT command directly. The -f switch points to a file holding the commands.

The commands in the RECREATE_FlashCopy_SCRIPT named recreate.cli are:

```
// Recreate and map the three FlashCopy logical Drives
on error stop;
// Create the FlashCopy
recreateFlashCopy logicalDrive [“FlashTarget-1"] ;
recreateFlashCopy logicalDrive [“FlashTarget-2"] ;
recreateFlashCopy logicalDrive [“FlashTarget-3"] ;
// Map drives
create mapping logicalDrive="FlashTarget-1" logicalUnitNumber=0 host="AAE18" ;
create mapping logicalDrive="FlashTarget-2" logicalUnitNumber=1 host="AAE18" ;
create mapping logicalDrive="FlashTarget-3" logicalUnitNumber=2 host="AAE18" ;
```

The same commands are used as with the basic disks to recreate the FlashCopy logical drives. We recreate all three FlashCopy logical drives and map them back to host *AAE18*.

Since the logical drives were not only disabled but unmapped we have to rediscover the disk devices itself first in Windows 2000 before we can actually bring the disks online.

```
echo Bus rescan.
%HOTADD%
```

```
echo Wait 5 seconds.
sleep 5

echo Bringing FlashCopy Target drive online with drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit
```

The **diskpart** commands we provide in the file DISKPART_SCRIPT called *recreate.dp* are:

```
rescan
select disk 2
online
select disk 3
online
select disk 4
online
select volume 0
online
assign letter=F
```

We first rescan the disk configuration of Windows 2000 and then set all disk drives that are part of the spanned volume online.

You must adapt the script here, as we had, for example, only two local drives, a CD-ROM, and a SCSI disk. You can start diskpart and use **list disk** to view all disk that are attached to your system at the moment. Using this information extracts the numbers of the disks and puts them in the script recreate.dp.

The same step is necessary for the spanned volume we want to bring online. After the disks are imported and set online, using diskpart again with the **list volume** command, you can extract the correct number of this volume.

The last command assigns a drive letter to the spanned volume and is now online and ready to use. The FlashCopy of the whole spanned volume is recreated and usable now.

Disable script

When the FlashCopy is not used anymore, it should be disabled. We again provide a script for this purpose.

As in the other scripts, all variables are defined at the top of the script.

```
rem Path to the SMFlashCopyassist utility.
set SNAPASSIST="c:\Program Files\IBM_FAStT\Util\SMFlashCopyassist"

rem Path to the SMFlashCopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMFlashCopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool
set RCMD_TOOL="C:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"
```

```
rem Path to the diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Path to the script for diskpart.
set DISKPART_SCRIPT="c:\scripts\disable.dp"

rem Path to the resacn script for diskpart.
set DISKPART_RESCAN_SCRIPT="c:\temp\scripts\rescan.dp"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the FlashCopy source drive.
set SOURCE_FlashCopy_DRIVE_LETTER="E:"

rem The drive letter of the FlashCopy target drive.
set TARGET_FlashCopy_DRIVE_LETTER="F:"
```

As you see later, we need two scripts for diskpart this time. The first one is DISKPART_SCRIPT for unmounting the spanned volume. The second one is DISKPART_RESCAN_SCRIPT, which is used after unmapping the FlashCopy logical drives to remove stale volume information from the disk administrator.

Now the command lines are generated to be executed in the next step:

```
rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FlashCopy_DRIVE_LETTER%"

rem Disable FlashCopy script.
set DISABLE_FlashCopy_SCRIPT="c:\temp\scripts\disable.cli"
```

The first step that is executed is the buffer flush of the base logical drive on the remote system:

```
echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FlashCopy_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"
```

Then the spanned volume residing on the three FlashCopy logical drives is also flushed so we can safely unmount it without the risk of having unwritten data still cached in the operating system:

```
echo Flushing target drive %TARGET_FlashCopy_DRIVE_LETTER%.
%SNAPASSIST% -f %TARGET_FlashCopy_DRIVE_LETTER%
if errorlevel 1 goto flush_target_fail_exit
```

Now we can safely remove the drive letter from the spanned volume using the first diskpart script DISKPART_SCRIPT:

```
echo Removing FlashCopy Target drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit
```

In the script DISKPART_SCRIPT called disable.dp, we remove the drive letter from the spanned volume we select in the first command. The volume number is the same as the in the diskpart script recreate.dp used during the recreation of the FlashCopy logical drives. Please refer to the explanation there to learn how you can extract the correct volume number.

```
select volume 0
remove
```

Now the spanned volume is safely removed from the host *AAE18*. We can unmap and disable the three FlashCopy logical drives building our spanned volume:

```
echo Disabling and unmapping FlashCopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -f %DISABLE_FlashCopy_SCRIPT%
if errorlevel 1 goto disable_FlashCopy_fail_exit
```

DISABLE_FlashCopy_SCRIPT is again only a collection of the commands we need to pass to the FAStT Storage Server. They are outlined in 5.3.1, “Windows 2000 basic disks” on page 177.

```
// Unmap and disable the three FlashCopy logical Drives
on error stop;
// Unmap drives
delete mapping logicalDrive="FlashTarget-1" host="AAE18" ;
delete mapping logicalDrive="FlashTarget-2" host="AAE18" ;
delete mapping logicalDrive="FlashTarget-3" host="AAE18" ;
// Disable the FlashCopy
disableFlashCopy logicalDrive ["FlashTarget-1"] ;
disableFlashCopy logicalDrive ["FlashTarget-2"] ;
disableFlashCopy logicalDrive ["FlashTarget-3"] ;
```

Now we removed not only the spanned volume, but also the disk devices from Windows 2000. To ensure a proper removal, we rescan our devices using the hot_add utility. We also rescan the disk configuration using diskpart to remove any obsolete drive information. This assures we can properly import the spanned volume again with the recreate script.

```
echo Bus rescan to remove stale devices.
%HOTADD%

echo Rescanning disks to remove stale drives and volumes.
%DISKPART_TOOL% -s %DISKPART_RESCAN_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit
```

The DISKPART_RESCAN_SCRIPT called rescan.dp consists only of the command to rescan the disk configuration.

```
rescan
```

Now our point-in-time copy of the spanned volume is properly removed from host *AAE18* and disabled on the FAStT Storage Server.

5.3.3 Linux disks

In this scenario, we use the following setup:

- Host *AAE18* is running Linux Redhat 7.2. The logical drive *FlashSource* is mapped from the Storage Subsystem *ITSO*. The disk is using either an ext2 or ReiserFS file system and is mounted under */mnt/source.*
- The FAStT host software is installed in the default location on host AAE18.
- The FlashCopy logical drive FlashTarget is also mapped to host AAE18 and mounted under */mnt/target*.

- Both logical drives are already configured using the FAStT MSJ.

We don't provide an initial script for Linux since the reconfiguration of the FAStT MSJ requires manual steps we cannot cover in a script.

The basic steps in the Linux scripts are the same as for the Windows scrips, but Linux provides all the needed utilities itself. And there is no need to take care of volume IDs, disk signatures, or disk devices that are present but cannot accessed as a disabled FlashCopy logical drive.

The scripts are written for the bash shell, which is a very common standard in Linux environments. It may be necessary to adapt the syntax of the script if you plan to use another shell.

The script can be used with either. Only the parameter TARGET_FILE_SYSTEM in the script must be changed.

Recreate script

In the first part, we define all variables we need during the script:

```
SMCLI_TOOL="/opt/IBM_FAStT/client/SMcli"

HOST_NAME="100.100.100.19"
FAStT_NAME=ITSO

TARGET_FlashCopy_DRIVE="FlashTarget"
TARGET_PARTITION="/dev/sdc1"
TARGET_MOUNT_POINT="/mnt/target/"
TARGET_FILE_SYSTEM="ext2"
```

With this information, we can create the command line we will send to the FAStT Storage Server for recreating the FlashCopy logical drive:

```
SMCLI_SCRIPT="recreateFlashCopy logicalDrive [$TARGET_FlashCopy_DRIVE];"
```

To flush the buffers of all disks used by the operating system, we use the **sync** command, to ensure the action, we repeat it after three seconds:

```
sync
sleep 3
sync
```

Now we can recreate our FlashCopy logical drive on the FAStT Storage Server. The commands sent to the FAStT Storage Server are the same as for the Windows scenarios and are explained there.

```
echo Recreating FlashCopy Target on FAStT
$SMCLI_TOOL -n "``$FAStT_NAME" -e -c "``$SMCLI_SCRIPT"
```

Now we can mount the FlashCopy logical drive back to the operating system and use it:

```
echo Mounting target drive
mount -t $TARGET_FILE_SYSTEM $TARGET_PARTITION $TARGET_MOUNT_POINT
```

Disable script

The disable script has the same layout as the recreate script. We define all the variables:

```
SMCLI_TOOL="/opt/IBM_FAStT/client/SMcli"

HOST_NAME="100.100.100.19"
FAStT_NAME=ITSO
```

```
TARGET_FlashCopy_DRIVE="FlashTarget"
TARGET_PARTITION="/dev/sdc1"
```

We can now create the command line for the storage manager command line tool:

```
SMCLI_SCRIPT="disableFlashCopy logicalDrive [$TARGET_FlashCopy_DRIVE];"
```

The next step is to flush all the file system buffers in the operating system:

```
sync
sleep 3
sync
```

We now can safely unmount the FlashCopy logical drive so the operating system will not try to access it anymore:

```
echo Unmounting target drive
umount $TARGET_PARTITION
```

The last step is to send the command line generated to the FAStT Storage Server to disable the FlashCopy logical drive:

```
echo Disabling FlashCopy Target on FAStT
$SMCLI_TOOL -n "``$FAStT_NAME" -e -c "``$SMCLI_SCRIPT"
```

6

Remote Volume Mirroring

This chapter introduces and discusses the concepts of the Remote Volume Mirror (RVM) feature. It describes what a Remote Volume Mirror is, how information is replicated between storage subsystems, and the connection of the storage subsystems and fabric configuration. The chapter also explains and gives step-by-step procedures for using, monitoring, and recovering a Remote Volume Mirror installation.

Remote Volume Mirroring provides protection from catastrophic events, including natural disasters, such as earthquakes and flood, as well as events such as fire, arson, and terrorism. It achieves this by maintaining two synchronized copies of a data set in two remote locations.

6.1 Introduction to the Remote Volume Mirror option

The Remote Volume Mirror option is a premium feature that comes with the IBM FAStT Storage Manager 8.2 software and is enabled by purchasing a premium feature key. The Remote Volume Mirror option is used for online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Volume Mirror option enables you to promote a second storage subsystem to take over responsibility for normal input/output (I/O) operations (Figure 6-1).

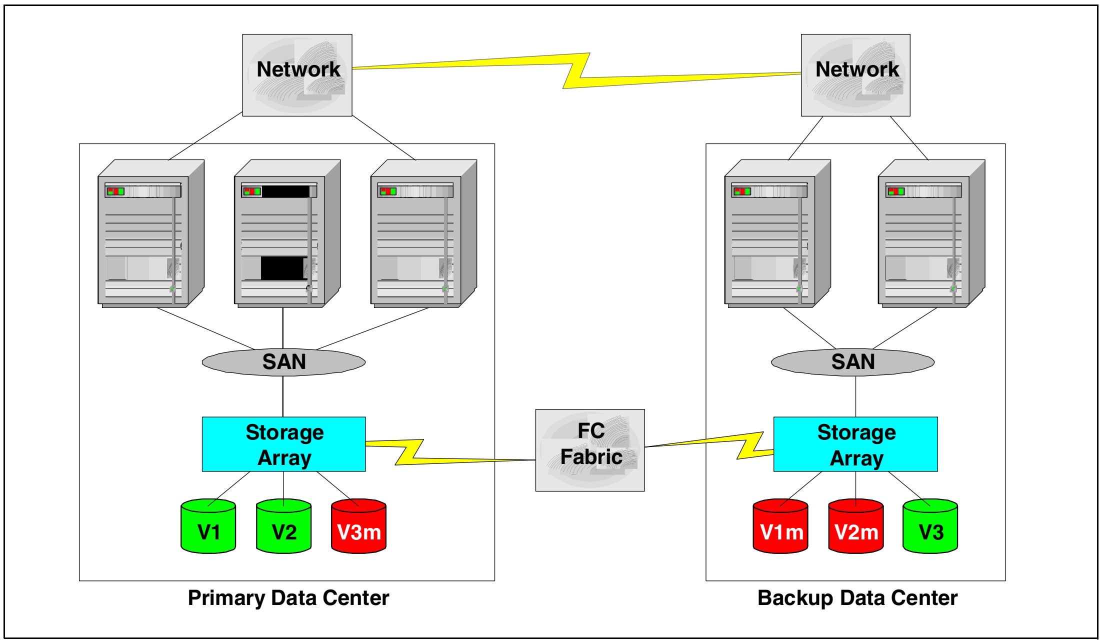

Figure 6-1 Remote Volume Mirroring

A mirroring relationship is on a volume basis:

- Associates two volumes (primary and secondary) using Storage Manager software.
- Data is copied to secondary volume in the background.

The mirroring is synchronous. The write must be completed to both volumes before the I/O is complete.

A minimum of two arrays is required. One storage system can have primary volumes being mirrored to other arrays and hold secondary volumes from multiple arrays.

6.2 Primary and secondary logical drives

When you create a Remote Volume Mirror, a mirrored logical drive pair is defined and consists of a primary logical drive at the primary storage subsystem and a secondary logical drive at a secondary storage subsystem. A standard logical drive may only be defined in one mirrored logical drive pair. The maximum number of supported mirrored logical drive pairs is determined by the storage subsystem model.

6.3 Mirror repository

A mirror repository logical drive is a *special logical drive* in the storage subsystem created as a resource for the *controller owner* of the primary logical drives in a remote logical drive mirror relationship. The controller stores status information only on this mirrored repository logical drive, including information about remote writes that are not yet written to the secondary logical drive.

Notes:

- No host data is written to the repository logical drive. It is only used for status and control data in relation to the Remote Volume Mirrored relationships.
- Only one repository drive is created for each controller in the storage system. That is one for Controller A and one for Controller B in the primary system and one for Controller A and one for Controller B in the secondary system (two per FAStT500 or FAStT700). The number of mirror relationships that a system supports does not effect the number of repository logical drives.

When you activate the Remote Volume Mirror option on the storage subsystem, two mirror repository logical drives are created by the system – one for each controller in the storage subsystem. An individual mirror repository logical drive is not needed for each mirror logical drive pair.

The example Figure 6-2 shows how the mirror repository drives are depicted in the logical view of the Storage Management software. These drives cannot be used or mapped for any other purpose.

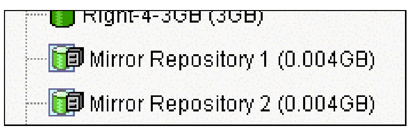

Figure 6-2 Mirror repository view

6.4 Mirror relationships

The primary and secondary role in a Remote Volume Mirror relationship is implemented at the logical drive level not at the storage subsystem level. A logical drive can only be in one mirror relationship at time. The logical drives in a storage subsystem can be in either a primary or secondary role. This means that some of the logical drives in *System A (primary)* can be mirrored to logical drives in *System B (secondary)*. While at the same time, other logical drives in *System B (primary)* can be mirrored to logical drives in *System A (secondary)*. This means that the storage subsystem can also have a combination of logical drives in a primary role and logical drives in a secondary role.

Whether the logical drive is in a primary or secondary logical drive, it counts toward the maximum number of mirror logical drive pairs that can be defined in one FAStT storage subsystem.

The mix of primary and secondary drives between the two storage subsystems is not important. However, the total Remote Volume Mirror relationships for a given storage system cannot exceed the maximum of 16 for a FAStT700 and eight for a FAStT500. For example, a storage system with a maximum of 16 Remote Volume Mirror relationship pairs may have logical drives in both primary and secondary roles.

The primary drive accepts the computer I/O and stores the data. When a mirror relationship is first established, the entire data content of the primary drive is copied to the secondary drive. Once this is completed, the mirror is considered to be synchronized. During the synchronization process (copying data), the primary drive remains available to the host for write and read I/O operations.

In a normal operation, when a write request is made to the primary logical drive, the controller owner of the primary logical drive also initiates remote write request to the secondary drive. The write request does not complete until both write requests are performed. This additional write request keeps the data on the two logical drives in the Remote Volume Mirror relationship synchronized. Whenever the data on the primary drive and the secondary drive becomes unsynchronized, the controller owner of the primary drive initiates a full synchronization.

The secondary logical drive is used to store data copied from its associated primary logical drive. The controller owner of the secondary logical drive receives remote writes from the controller owner of the primary logical drive and will not accept host computer read or write requests. The secondary logical drive is unavailable to host computer programs while in a mirror relationship as a secondary. In the event of a disaster or unrecoverable error of the primary storage subsystem, a role reversal is performed to promote the secondary logical drive to the primary logical. Host computers can then access the newly promoted logical drive and normal operations can continue.

Figure 6-3 shows the relationship between primary and secondary logical drive on the primary storage system as it would appear in the physical view of the Storage Management GUI for the local storage system.

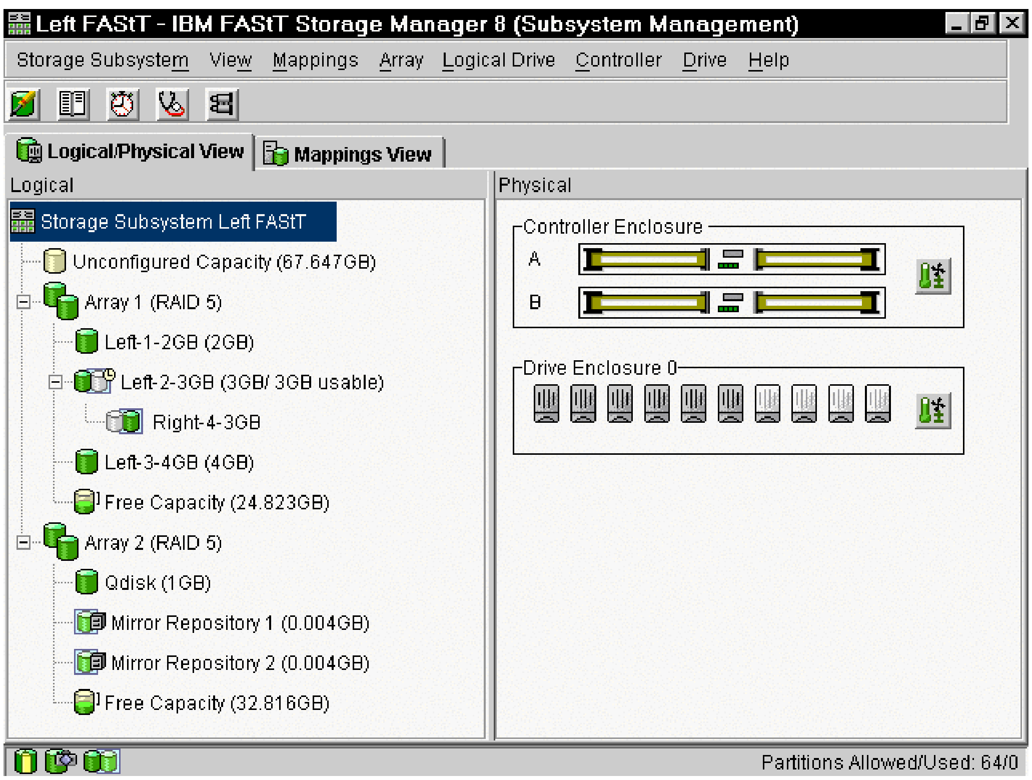

Figure 6-3 Primary storage system

Figure 6-4 shows the secondary drive “Right-4-3GB” as it would appear in the physical view of the Storage Management GUI for the remote storage system.

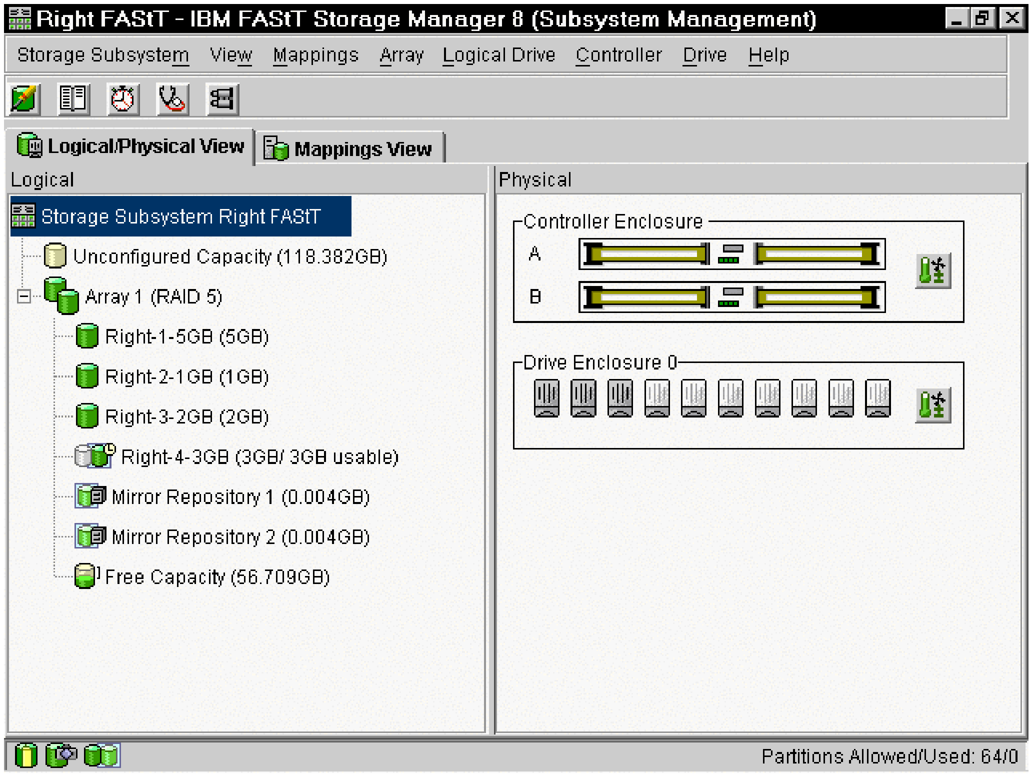

Figure 6-4 Remote storage subsystem

6.5 Logical drive limits

Enabling the Remote Volume Mirror option limits the number of logical drives that can be created in a single storage subsystem. When RVM option is enabled, the total number of logical drives that are supported for each storage subsystem is reduced from the number of logical drives you would have without the Remote Volume Mirror option enabled. The reduced maximum number of logical drives that a FAStT storage subsystem can support with firmware version 05.20.xx.xx depends on the model of the FAStT storage subsystem. Refer to Table 6-1 for the current limitations.

With future releases of firmware, the maximum number logical drives many vary. Refer to the current release notes and software documentation for the latest updates.

Table 6-1 Logical drive limits

Model	Maximum drives with RVM option enabled	Maximum drives without RVM option enabled	Comments
FAStT200	N/A	512	RVM not supported on FAStT200
FAStT500	128	512	Maximum number of logical drives includes the FlashCopy repository logical drives and the Remote Volume Mirror repository logical drives
FAStT700	256	2048	

6.6 Fabric configuration

The Remote Volume Mirror option requires two dedicated host ports for connection between primary and secondary storage systems.

When the Remote Volume Mirror option is activated, one Fibre Channel host-side I/O port on each controller is dedicated only to mirroring operations. For example, in the primary storage subsystem, controller host ports A2 and B2 are dedicated to mirroring operations. In the secondary storage subsystem, controller host ports A2 and B2 are also dedicated to mirroring operations. See Figure 6-5 for the port locations.

One mini-hub connects to a host port in a controller. Each controller (A and B) has two host ports. For this reason, two mini-hubs are required to be dedicated to Remote Volume Mirroring on each of the storage system that requires the feature activated.

The level of redundancy within the fabric depends on the fabric design and Fibre Channel switch configuration. This book does not specifically address Storage Area Network (SAN) design issues. However, you can find three sample configurations in 6.17, “Remote Volume Mirroring solution design” on page 226.

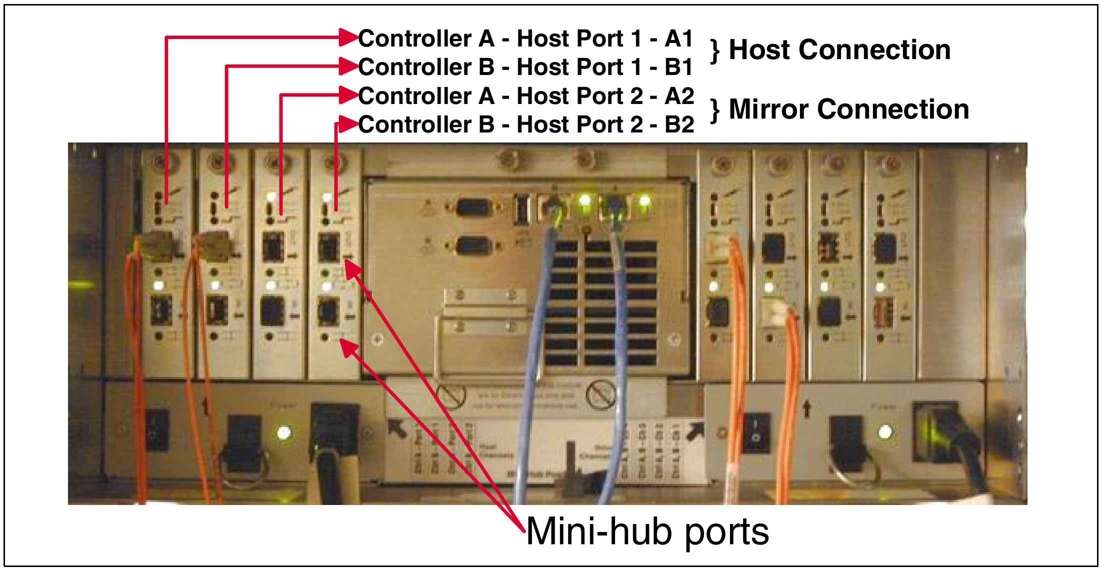

Figure 6-5 Port locations

6.6.1 Distance limits

The distance between primary and secondary storage subsystems is limited by the distance limits of Fibre Channel inter-switch links (ISL). See Table 6-2.

Table 6-2 Fibre Channel distance limits

Fiber cable type	Laser type	Distance limit in kilometers	Distance limit in miles
Single mode 9 micron	Long wave	10 km.	6.25 miles
Multi mode 50 micron	Short wave	0.5 km.	0.32 mile

Note: Link distance greater than 10 km. can be achieved by using extended fabric features and Fibre Channel extension via protocol conversion such as FC-ATM-FC. These options have not been tested or certified with the FAStT Storage Server.

6.6.2 Switch zoning

It is a mandatory requirement that you keep FAStT links, used as Remote Volume Mirroring links, in a separate zone of there own. We recommend that you create four separate zones within the fabric:

- The first for a host connection to controller “A”
- The second for host connection to controller “B”
- The third for Controller “A” Remote Volume Mirror links between storage subsystem
- The fourth for Controller “B” Remote Volume Mirror links between storage subsystem

Figure 6-6 shows one example of zoning a fabric for FAStT and Remote Volume Mirroring.

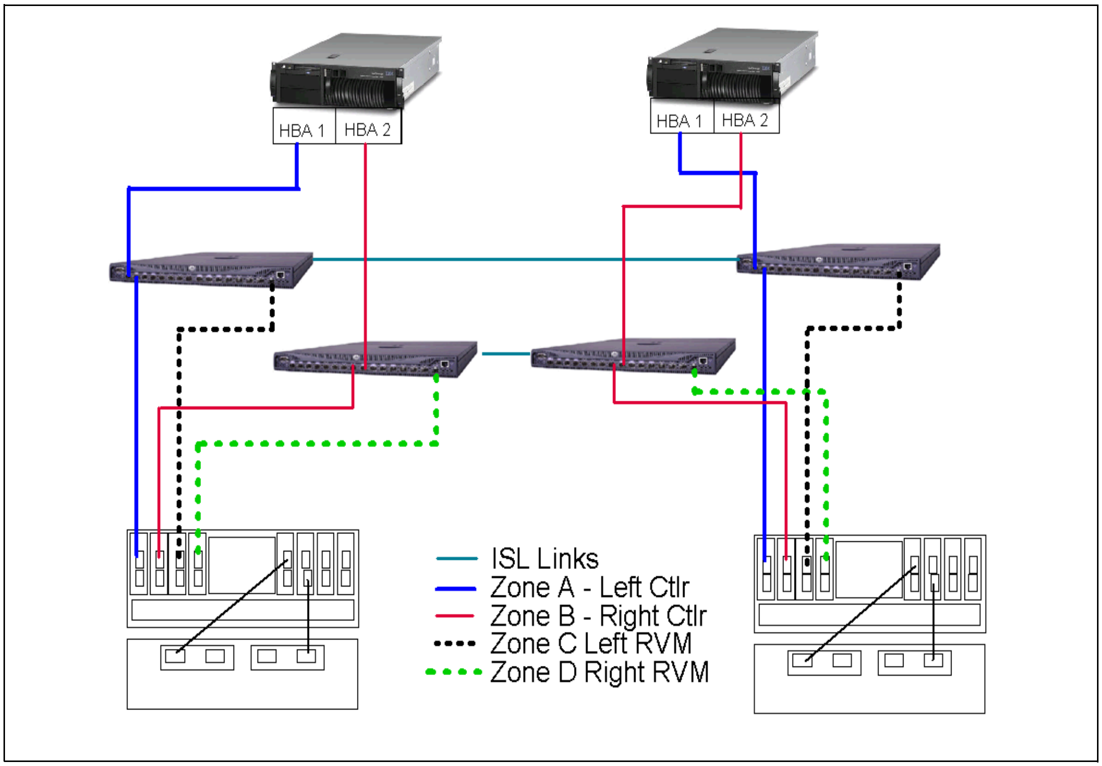

Figure 6-6 RVM zoning example

6.7 Data replication

Data replication between the primary logical drive and the secondary logical drive is managed at the controller level. It is transparent to the attached host computers and applications.

When the controller owner of the primary logical drive receives a write request from a host computer, the controller logs information about the write to its mirror repository first. Then it writes the data to the primary logical drive. Next, the controller initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the secondary storage subsystem.

Note: The owning primary controller only writes status and control information to the repository logical drive. It is not used store host data.

After the host computer write request is written to the primary logical drive and the data is successfully copied to the secondary logical drive, the controller removes the log record on the mirror repository logical drive. Then it sends an I/O completion indication back to the host computer. See Figure 6-7 for the data replication process.

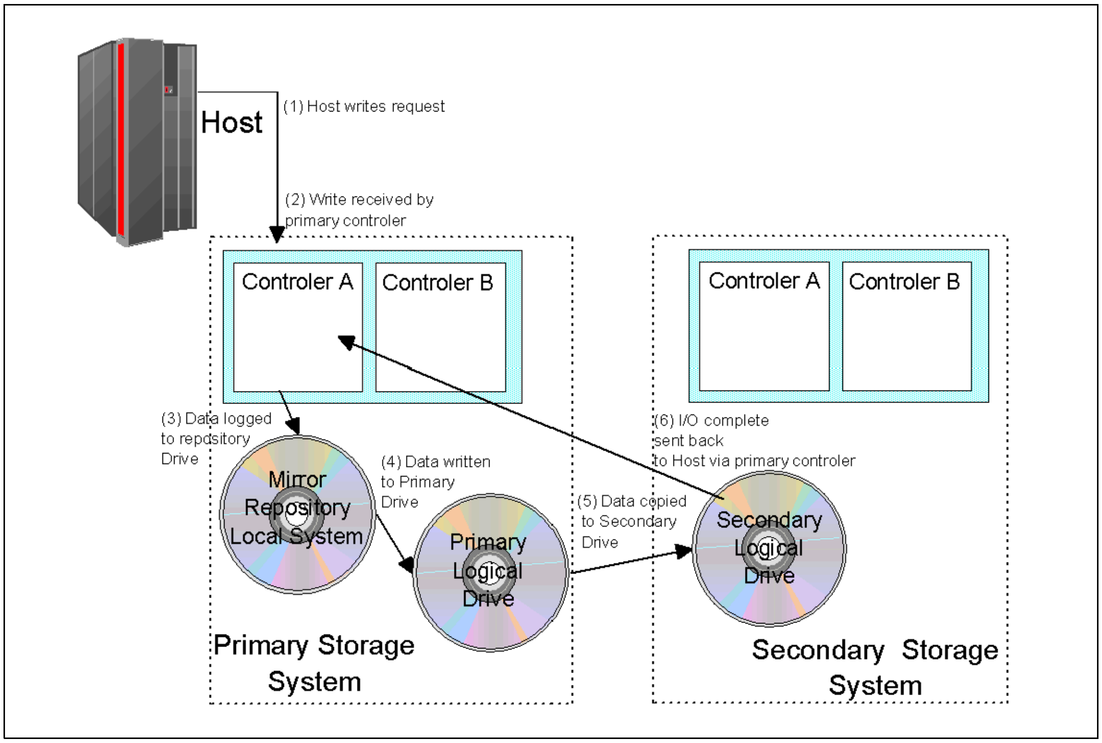

Figure 6-7 Remote Volume Mirror logical data flow diagram

The host does not receive the I/O completion until both the primary and remote write have completed.

On a read I/O request only, the primary controller and primary logical drive is accessed. There is no communication with the secondary subsystem on read.

6.8 Performance considerations

This section lists the factors that can effect the performance of remote volume mirroring:

- The controller owner of a primary logical drive performs a full synchronization in the background while processing local I/O writes to the primary logical drive and associated remote writes to the secondary logical drive. Because a full synchronization diverts controller processing resources from I/O activity, this can impact performance of the host computer application.

  To reduce the performance impact, you can set the synchronization priority to establish how the controller owner prioritizes a full synchronization relative to other I/O activity. The time required for a full synchronization varies depending on the data capacity and I/O activity that is to be synchronized. However, if the slowest synchronization priority setting is selected, the time taken for a full synchronization is approximately eight times longer than the shortest synchronization setting.

- After the Remote Volume Mirror is synchronized, on a host write I/O, the primary owning controller writes status information only to the repository drive and writes host data to the primary drive and then to the remote secondary drive. The last step is to return I/O complete to the host. On a host read I/O, the data is accessed from the primary drive only in the normal way. The RVM process may slightly affect the write performance. Although during our testing, we did not detect loss in performance due to RVM being activated. A read I/O has no impact on performance.

- When the mirror logical drive pair is in the Synchronization-in-Progress status, all host computer write data is copied to the remote system in processing the I/O. Controller I/O bandwidth and I/O latency affect host computer write performance. Host read performance is not affected by the mirror relationship.
- The overall system performance is impacted when data is copied from the primary logical drive to the secondary logical drive.

6.9 Loss of connection to secondary drive

During processing write requests, a primary controller may be able to write to the primary logical drive, while a Fibre Channel link failure between the primary and remote storage systems may prevent communication with the remote secondary controller. Then the remote write cannot complete to the secondary logical drive, and the primary and secondary logical drives are no longer correctly mirrored. The primary controller transitions the mirrored pair into a Link Failure status and sends an I/O completion to the primary host computer, as shown in Figure 6-8. The primary host computer can continue to write to the primary logical drive but remote writes do not take place.

When connectivity is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, the status changes to Unsynchronized, and a full synchronization automatically takes place. The mirrored pair transitions from an Unsynchronized status to a Synchronization-in-Progress status.

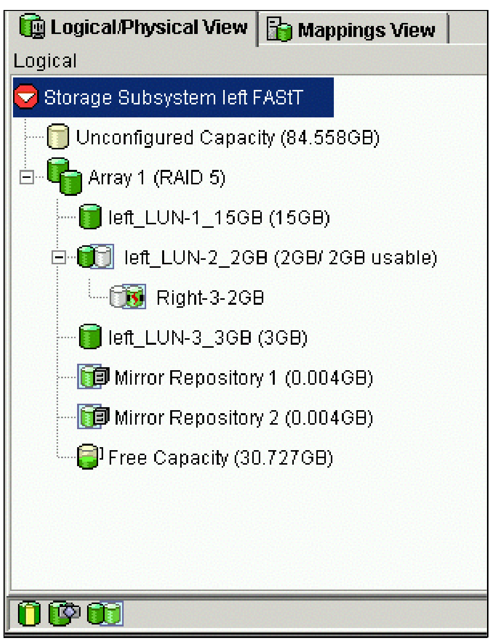

Figure 6-8 Loss of mirror communication

The status of the mirror status can be determined by its symbols as shown in Figure 6-9.

Remote mirror status representation in logical view

Logical Drive Status	Remote Mirror Status	Primary	Secondary
Optimal	Synchronized		
	Synchronization in progress		
	Unsynchronized		
Failed	Synchronized		
	Synchronization in progress		
	Unsynchronized		
Degraded	Synchronized		
	Synchronization in progress		
	Unsynchronized		
Offline	Synchronized		
	Synchronization in progress		
	Unsynchronized		
Missing	Synchronized		
	Synchronization in progress		
	Unsynchronized		
Unresponsive	Synchronized		
	Synchronization in progress		
	Unsynchronized		

Figure 6-9 Mirror status symbols

Note: A loss of communication between the primary and secondary do not result in the controllers attempting to change ownership of drives. The only time ownership changes is on a host path failover. This results in the secondary mirror to change ownership to match the primary on the next write I/O. For more information, see the following section.

6.10 Logical drive ownership

The preferred owner controller of the primary drive only attempts to communicate only with its mirror controller in the secondary storage system. For example, controller “A” in the primary only communicates with controller “A” in the secondary. Controller “B” only communicates with Controller “B”. For this reason, the controller that is the preferred owner of the drive on the primary dictates the preferred owner on the secondary.

If primary Controller A cannot communicate with secondary Controller A, no controller ownership changes take place and the Remote Volume Mirror link is broken for that mirror logical drive pair.

If the host RDAC or AVT initiate a failover to controller “B”, primary controller B communicates with secondary controller B. The next I/O request to the primary logical drive results in a secondary drive ownership change to controller “B”.

6.11 FlashCopy and Remote Volume Mirroring

You cannot create a FlashCopy of Remote Volume Mirror secondary drives. FlashCopy is supported for Remote Volume Mirror primary drives. If you want to FlashCopy the mirror drives on the secondary site, you must break the link between the two sites and then promote the secondary drives to primary drives. You can then carry out the FlashCopy at the secondary site. If you want to resume mirroring between the primary and secondary sites, you must first terminate the FlashCopy. Then a full resynchronization of the primary drives occurs.

6.12 Supported hardware and operating systems

Host computers accessing mirrored logical drives must have one of the following operating systems installed to use the Remote Volume Mirror option. Currently the following operating systems are supported on various FASTt models (see Table 6-3):

- Windows NT 4.0 Server Edition
- Windows 2000
- Linux (RedHat Version 7.x, SuSE 7.3, Turbo 7.0)
- NetWare Versions 5.1, 6.0

Support for the following operating systems is expected later:

- Solaris 2.6, 7 and 8
- HP-UX Version 11.0 or later
- AIX 4.33, 5.1

Note: At the time this redbook was redbook, support for all operating systems in the list had not been announced. Check the Web addresses in the following table for the most current information on hardware and operating system support.

Table 6-3 Supported server matrix

Model	Link to latest supported server matrix
FAStT700	http://www.storage.ibm.com/hardsoft/products/fast700/supserver.htm
FAStT500	http://www.storage.ibm.com/hardsoft/products/fast500/supserver.htm
FAStT200	http://www.storage.ibm.com/hardsoft/products/fast200/supserver.htm

6.13 Supported configurations

A supported configuration uses certified hardware and software. It also has been implemented in accordance with all current IBM release notes and installation documentation. Refer to the Web addresses in Table 6-3 for the current support matrix.

6.14 Using Remote Volume Mirroring

This section presents a guide to enabling, configuring, and implementing Remote Volume Mirror between two FAStT storage subsystem. For the purpose of demonstrating the Remote Volume Mirror function, the solution shown in Figure 6-10 was created.

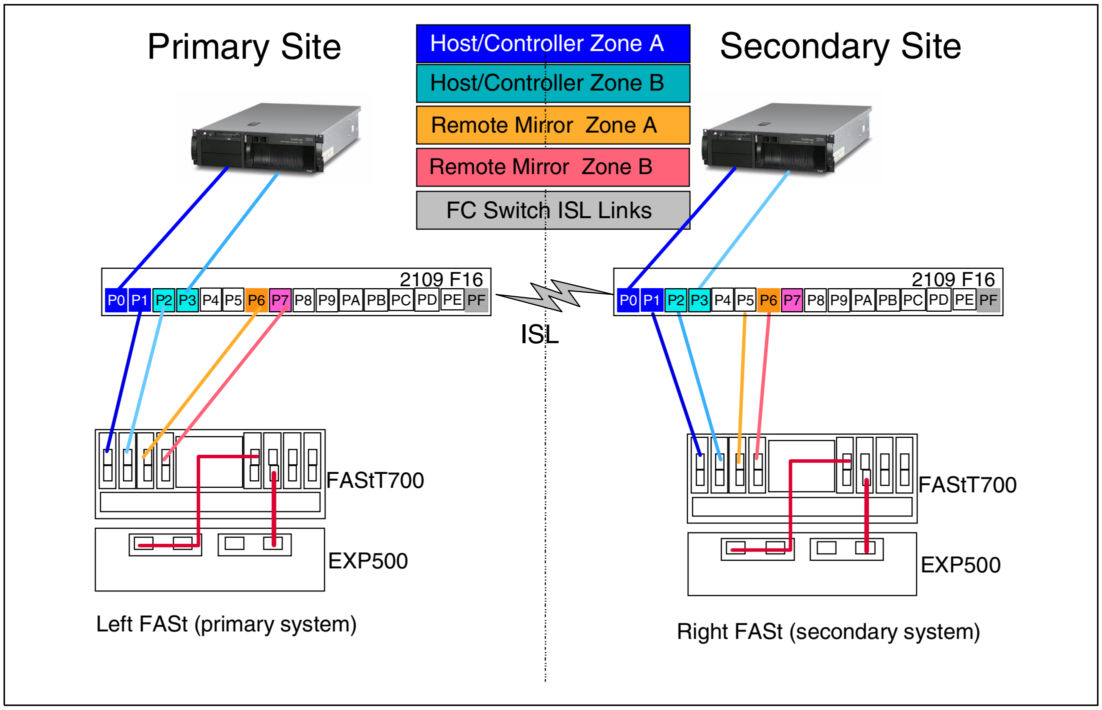

Figure 6-10 Test system

6.14.1 Enabling the Remote Volume Mirror premium feature option

The Remote Volume Mirror option is a premium feature that is included as part of the IBM FAStT Storage Manager 8.2 installation. However, the option is not enabled or activated by default after installation. You must purchase an IBM Storage Manager Remote Volume Mirror option for each storage subsystem in your Remote Volume Mirror configuration.

This option contains a *GenKey file* that enables the Remote Volume Mirror option premium feature. If you want to purchase the Remote Volume Mirror premium feature option, contact your IBM Technical Support representative.

Important: All storage subsystems participating in mirror relationships must have installed firmware level 05.2x.xx.xx or higher. For more information about upgrading firmware and installing the storage-management software, refer to Chapter 7, “Migrating from an earlier version of FAStT Storage Manager” on page 231.

If the current status of the Remote Volume Mirror option is Disabled/Deactivated or Disabled/Activated on any of the storage systems, it must be enabled first using the premium option feature key.

6.14.2 Remote Volume Mirror status

To check whether the Remote Volume Mirror option is activated on a FAStT Storage Server, complete the following steps for each FAStT storage system you want to use within your mirroring solution:

1. Start the Storage Management software.
2. From the Enterprise Management window, click **Tools-> Automatic Discovery**. Click **OK**.
3. Verify that all storage subsystems designated to participate in mirror relationships are displayed in the Device Tree view of the Enterprise Management window.
4. From the Enterprise Management window, select a storage subsystem that will participate in mirror relationships.
5. Click **Tools-> Manage Device** to open the Subsystem Management window for that storage subsystem.
6. From the Subsystem Management window, click **Help-> About**.
7. Verify that the version of the storage-management software is 08.2x.xx.xx or later. If this requirement is not met, refer to Chapter 7, "Migrating from an earlier version of FAStT Storage Manager" on page 231, for storage-management software upgrade procedures.
8. Click **OK**.
9. From the Subsystem Management window, click **View-> Storage subsystem Profile**.
10. Under the **All** tab, verify that the current firmware version is 05.2x.xx.xx or later. If this requirement is not met, refer to Chapter 7, "Migrating from an earlier version of FAStT Storage Manager" on page 231, for firmware upgrade procedures.

> **Tip for Enterprise Management window help**: Use this online help system to learn more about working with the entire management domain. You can access these help systems from within the Storage Manager 8.2 Client. From either a Enterprise Management or Subsystem Management window, click **Help** or press F1.

11. Click **Close**.
12. Observe the Remote Volume Mirroring icon in the status area as shown in Figure 6-11. The icon represents the Disabled and Deactivated status of the Remote Volume Mirror option.

There are four possible states of the Remote Volume Mirror option:

- Disabled/Deactivated
- Disabled/Activated
- Enabled/Deactivated
- Enabled/Activated

> **Important:** The Remote Volume Mirror option must be in the Enabled/Activated state on the secondary storage subsystem as well as on the primary storage subsystem to create and maintain mirror logical drive pairs.

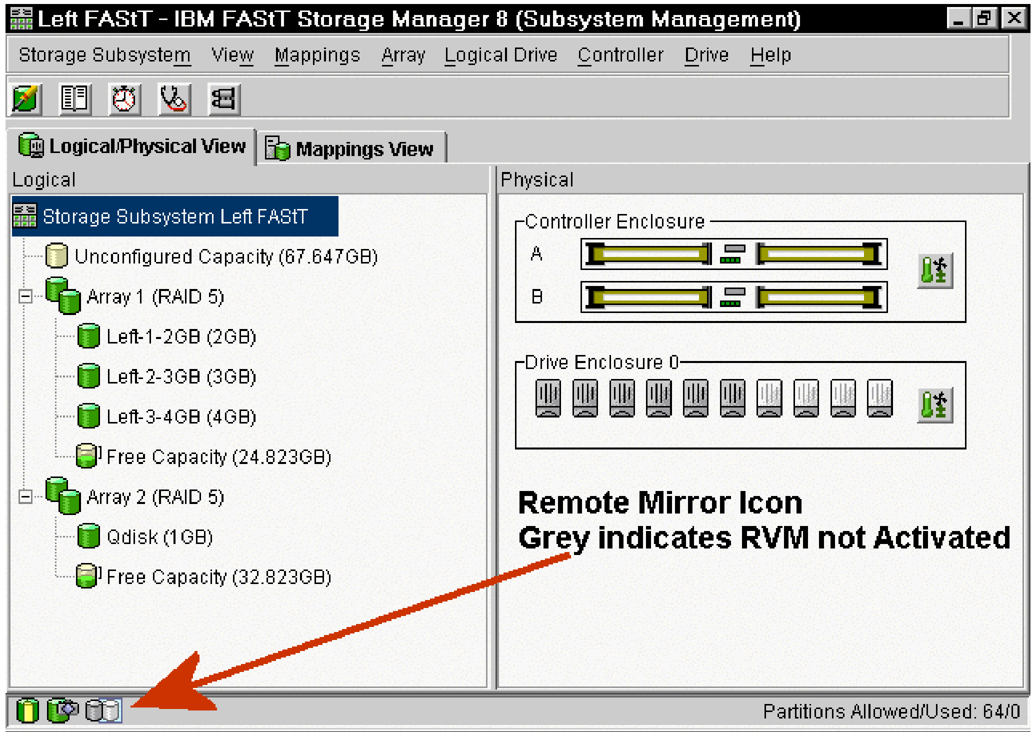

Figure 6-11 RVM not activated

Tip: The status can be determined by hovering the mouse pointer for a few seconds over the Remote Volume Mirror icon.

6.14.3 Enabling and activating the Remote Volume Mirror option

This section explains how to enable and activate the Remote Volume Mirroring option.

Enabling the premium option feature

To enable the Remote Volume Mirror option, follow these steps:

1. From the Subsystem Management window, click **Storage subsystem-> Features-> Enable**.
2. The Select Feature Key File window opens. Select the folder in which you placed the generated key file.

   Note: The Select Feature Key File program filters files with the .key extension. Refer to the instructions on the *IBM FAStT Storage Manager Version 8.2 Remote Mirror Option Premium Feature Key* CD-ROM.

3. Select the appropriate key file, and then click **OK**.
4. The Enable Feature window opens. Click **Yes**.
5. The Remote Volume Mirror option is now enabled. The icon in the Premium Feature status area no longer displays a red slash. To further verify the status of the option, click **Storage subsystem-> Features-> List**.

Activating the Remote Volume Mirror option

Activating the Remote Volume Mirror option prepares the storage subsystem to create and configure mirror relationships. When the option is activated, ports A2 and B2 of the storage controller become reserved and dedicated to Remote Volume Mirror option usage. In addition, a mirror repository logical drive is created for each controller in the storage subsystem.

> **Note:** Before activating Remote Volume Mirror function, the status must be Enabled/Deactivated.

If the current status of the Remote Volume Mirror option is Enabled/Deactivated, complete the following procedure to activate the option:

1. From the Subsystem Management window, click **Storage subsystem-> Remote Mirror-> Activate** (see Figure 6-12).

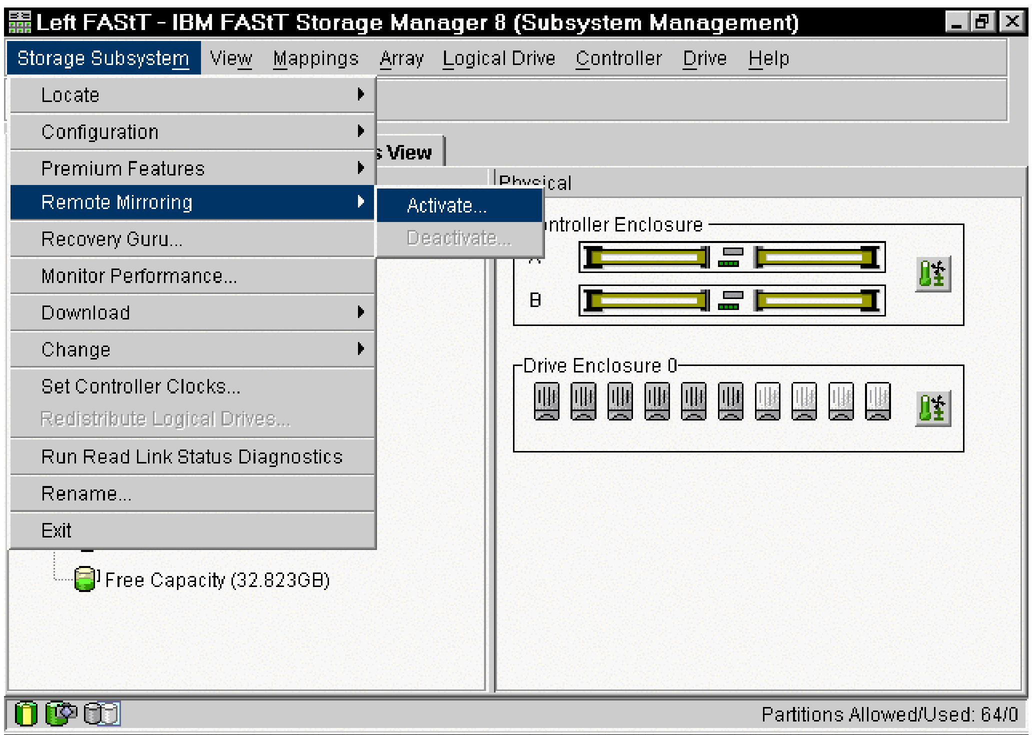

Figure 6-12 Activating RVM

2. The Activate Remote Logical Drive Mirroring - Introduction window opens (Figure 6-13). Decide where the mirror repository logical drives will reside. Select one of the following options:

   – *Free capacity on existing arrays*: If this option is selected, select a corresponding logical drive.

      i. Select a logical drive.
      ii. Select **Finish** or **Next**.

   – *Unconfigured capacity (create new array)*: If this option is selected, more information must be collected about the new logical drive.

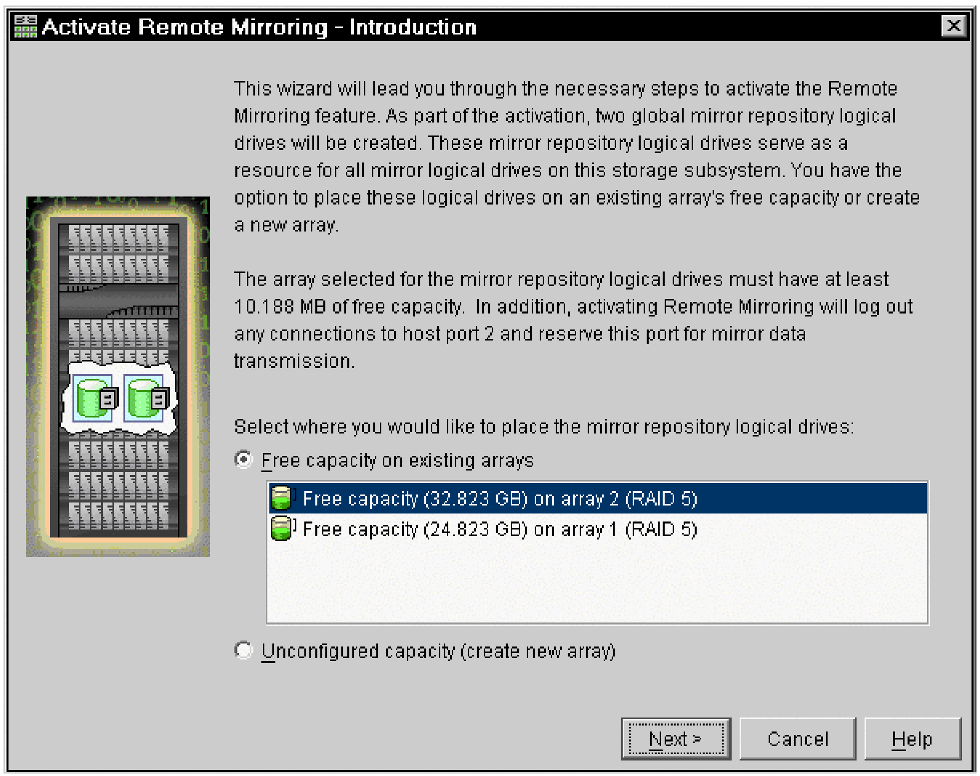

Figure 6-13 Activate Remote Logical Drive Mirroring - Introduction window

i. In the Create New Array window (Figure 6-14), select the RAID level for the new logical drive.

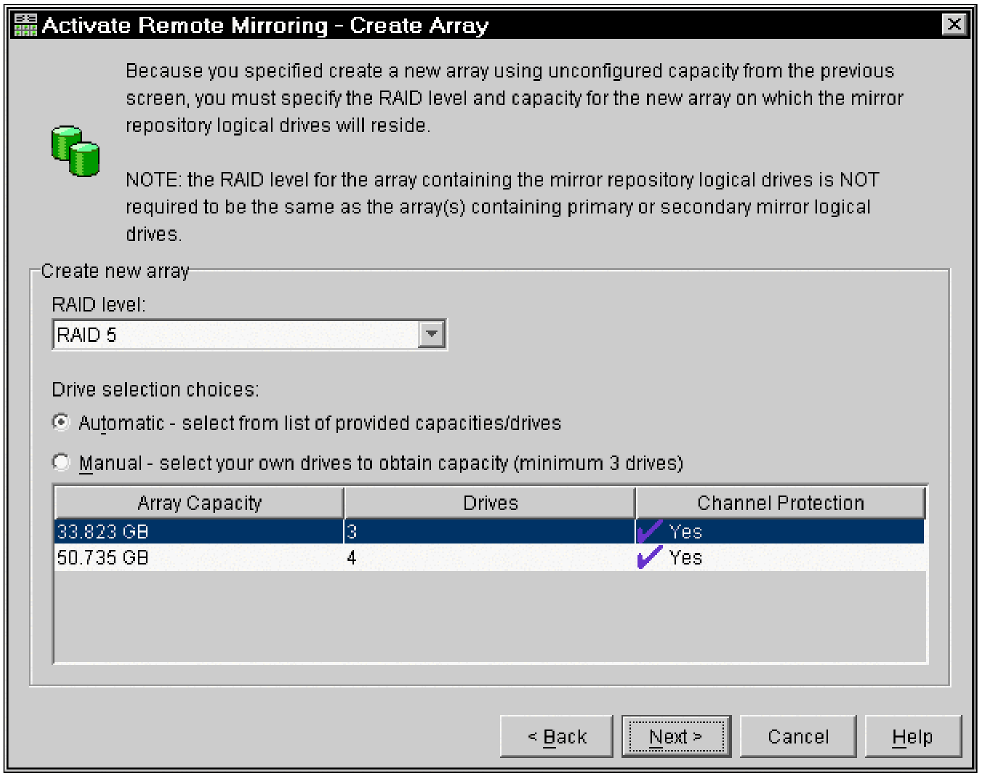

Figure 6-14 Create New Array window

ii. Under the drive selection choices, select one of the following options:

- **Automatic**: The drives are chosen automatically according to available capacity.

- **Manual**: The user can specify which drives contain the mirror repository logical drive. For this option, the press and hold the Ctrl key and click to select the drives you want to include in your array. Then click **Next**.

iii. An information preview window opens (Figure 6-15). Review the information and click **Finish**.

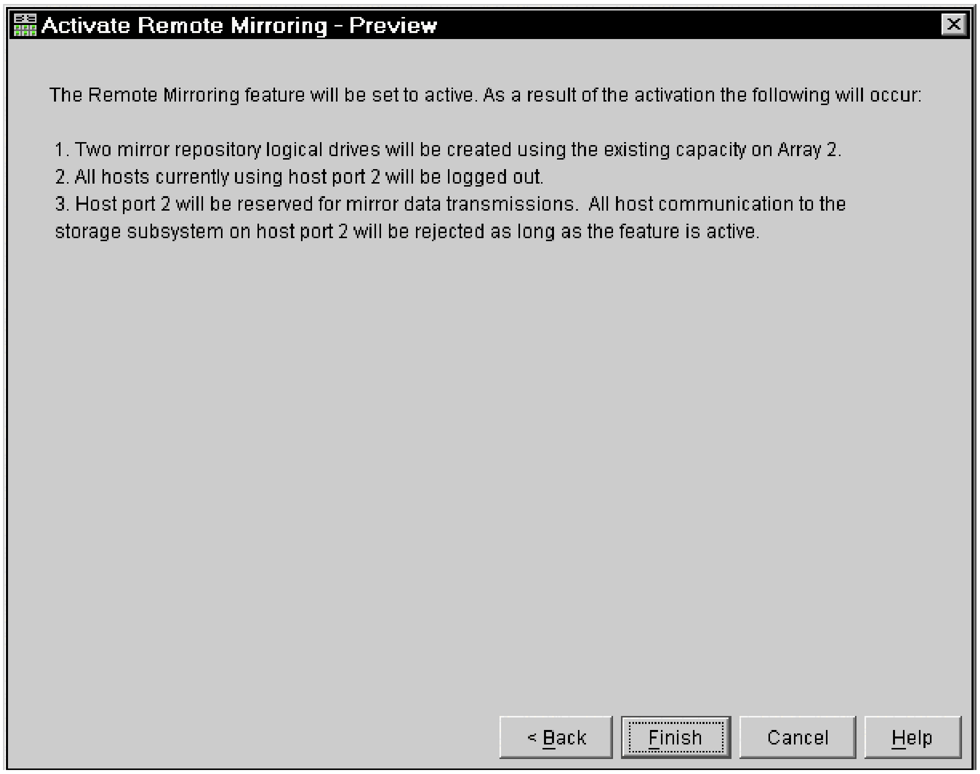

Figure 6-15 RVM Preview window

3. A completion window appears (Figure 6-16), advising that the Remote Volume Mirroring feature is active. Review the information and click **OK**.

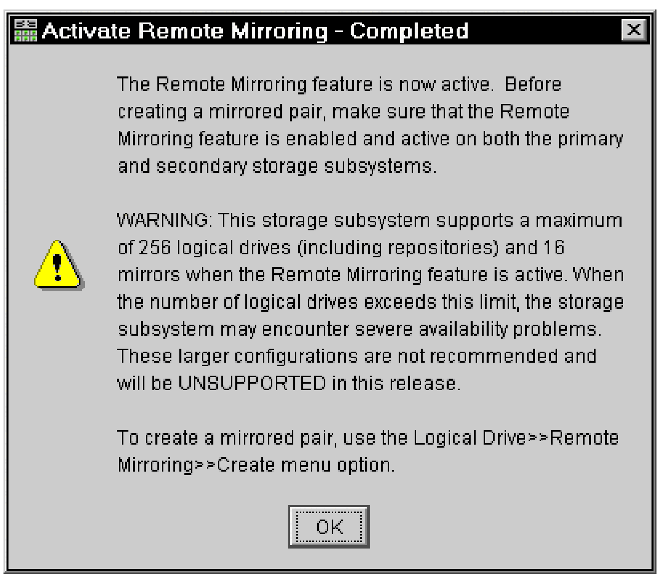

Figure 6-16 Activate Remote Mirroring - Completed window

6.14.4 Creating Remote Volume Mirror relationships

Before you create mirror relationships, the logical drive must exist at both the primary and secondary storage subsystems. The secondary logical drive on the remote storage subsystem must be at least the same capacity as the primary drive. Otherwise, it is not possible to select the drive to create a mirror relationship. The storage subsystem that the primary drive resides in is called the *primary storage subsystem*. Similarly, the logical drive residing in the secondary storage subsystem is the secondary logical drive.

Note: The secondary logical drive must be of equal to, or greater than, in size the primary logical drive.

You can access the Subsystem Management window online help for more information about creating logical drives.

Once logical drives exist at both sites, you can create mirror relationships by using the Logical/Physical View of the Storage Management window for the primary storage system and by following these steps:

1. Select the logical drive. Then right-click and select **Create Remote Mirror** (Figure 6-17).

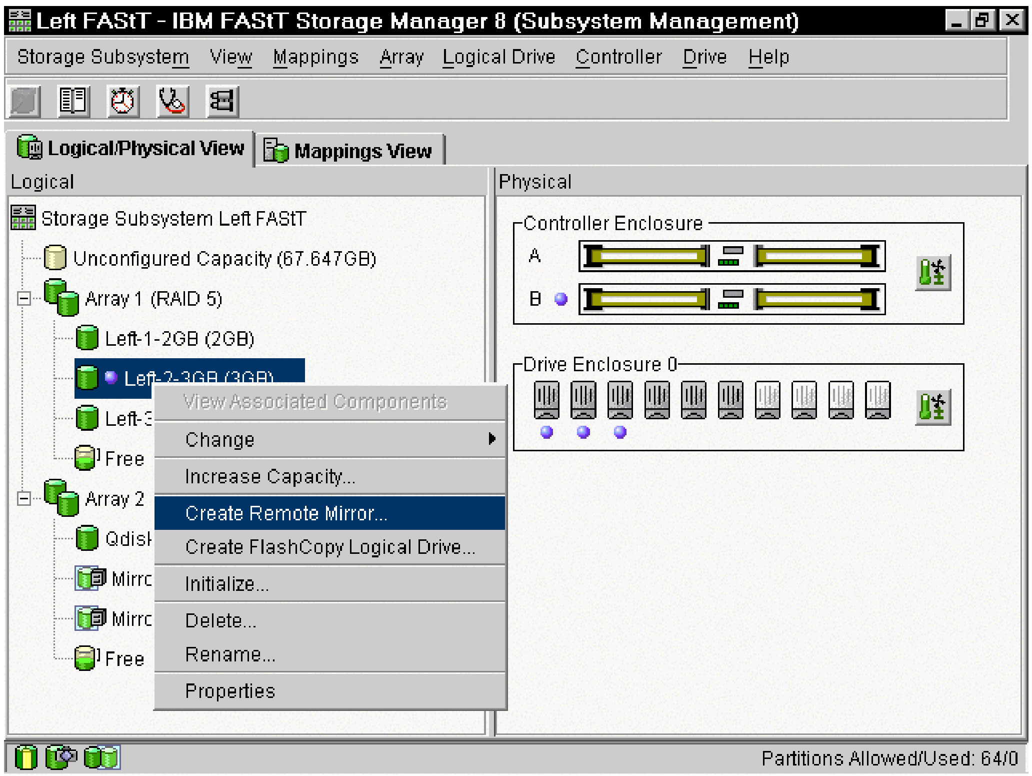

Figure 6-17 Selecting the Create Remote Mirror wizard

2. The introduction window (Figure 6-18) of the Create Remote Mirror wizard appears. Review the information and click **Next** to continue.

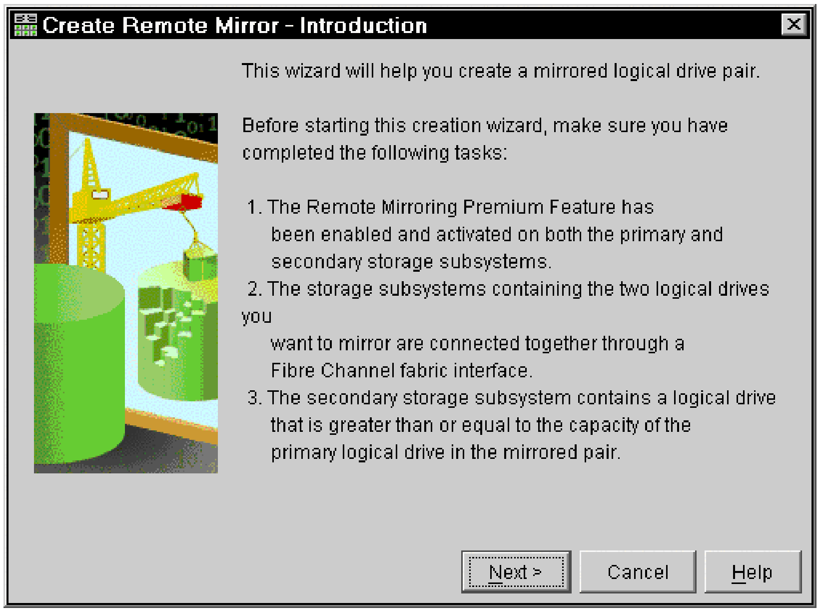

Figure 6-18 Create Remote Mirror - Introduction window

3. When the select storage subsystem window appears (Figure 6-19), highlight the FAStT storage subsystem that you want to be the secondary. Click **Next** to create the remote storage subsystem.

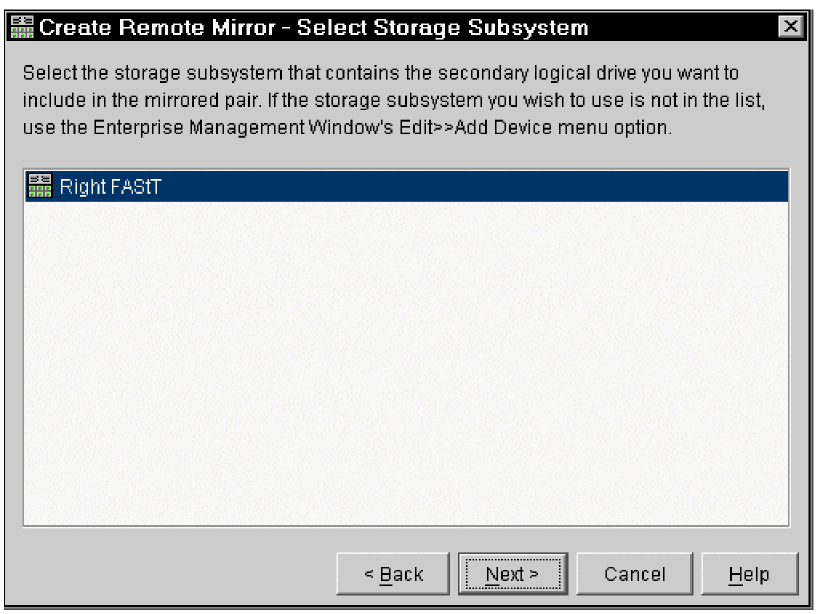

Figure 6-19 Select Storage Subsystem window

The Remote Volume Mirror option supports a maximum of two storage subsystems per configuration. The wizard displays all available storage subsystems within the management domain that has the Remote Volume Mirror option enabled and activated. If you have an existing Remote Volume Mirror relationship with a given storage subsystem, do not select another storage subsystem from the list.

4. In the next window (Figure 6-20), select the logical drive that will become the secondary logical drive in the mirror relationship. If no logical drives are listed, the secondary storage

subsystem that you selected does not have logical drives with enough capacity available to mirror the selected logical drive. Click **Next**.

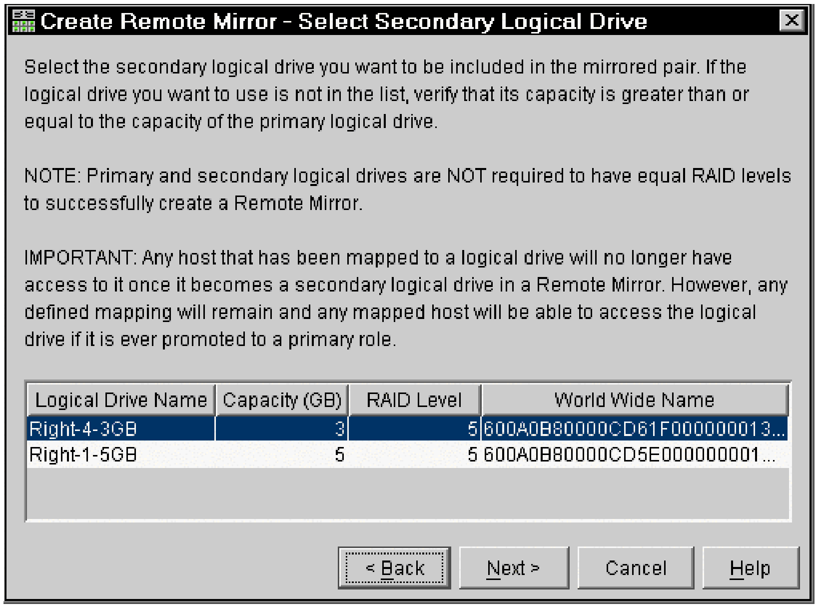

Figure 6-20 select secondary logical drive window

5. In the Set Synchronization priority window (Figure 6-21), select the Synchronization Priority Level. Click **Finish**.

> **Synchronization Priority Level**: This level of a mirror relationship defines the amount of system resources used to synchronize the data between the primary and secondary logical drives of a mirror relationship. If the highest priority level is selected for a mirror relationship, the data synchronization uses a high amount of system resources to increase mirror performance, which may decrease performance for all other functions, including other mirror relationships. If the lowest synchronization level is selected, there is less impact on complete system performance, but the mirror relationship synchronization might be slower.

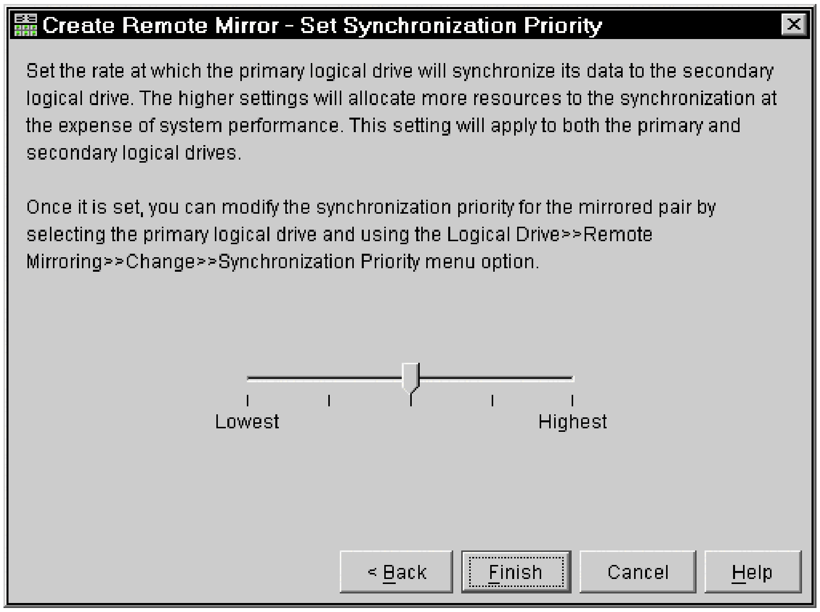

Figure 6-21 Set Synchronization Priority window

6. If you do not exceed the maximum number of mirrored logical drive pairs that can be created for the storage subsystem, you are prompted to create another mirror logical drive pair by the Create Remote Logical Drive mirror wizard. If you do not want to create another mirror logical drive pair, click **No**.

6.14.5 Viewing mirror relationships

The mirror status can be viewed by the Storage Management device logical view. On the primary storage system, the logical drive icon changes to show the icon of a mirror device and the icon of the Remote Volume Mirror attached beneath it. While synchronization is in progress, a clock symbol appears on the primary mirror (see Figure 6-22).

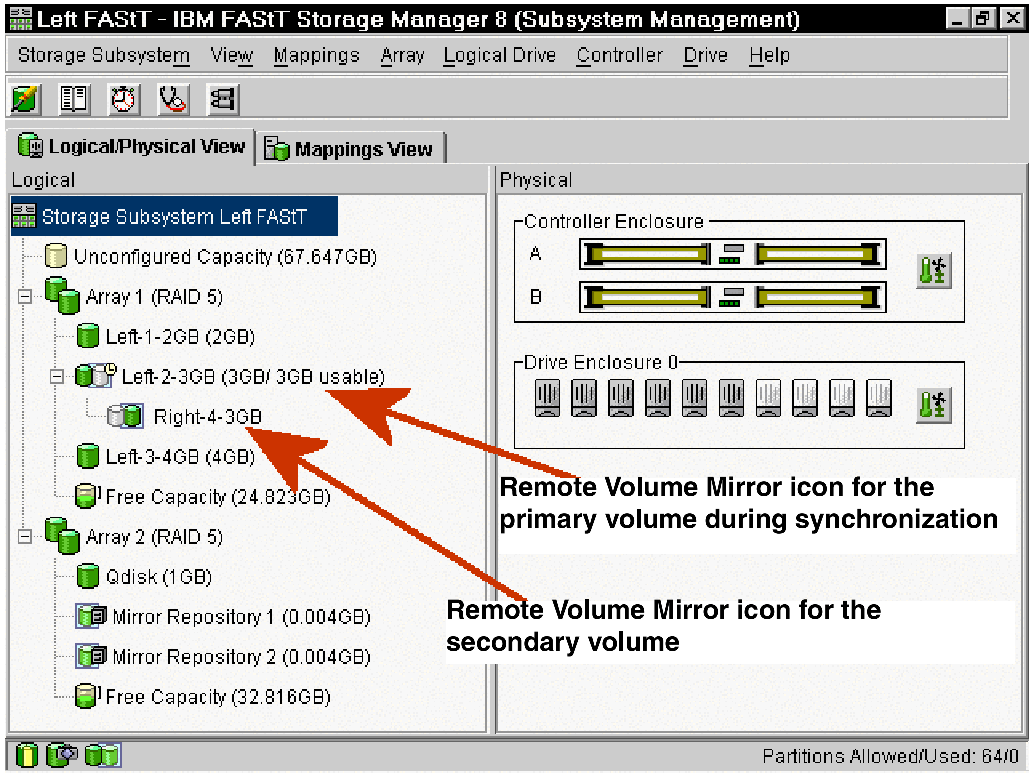

Figure 6-22 Primary storage system

From the remote storage system, the secondary drive icon also changes to a mirror icon with the clock displayed while synchronization is in progress (Figure 6-23). For a detailed list of Remote Volume Mirror drive icon status, see Figure 6-9 on page 200.

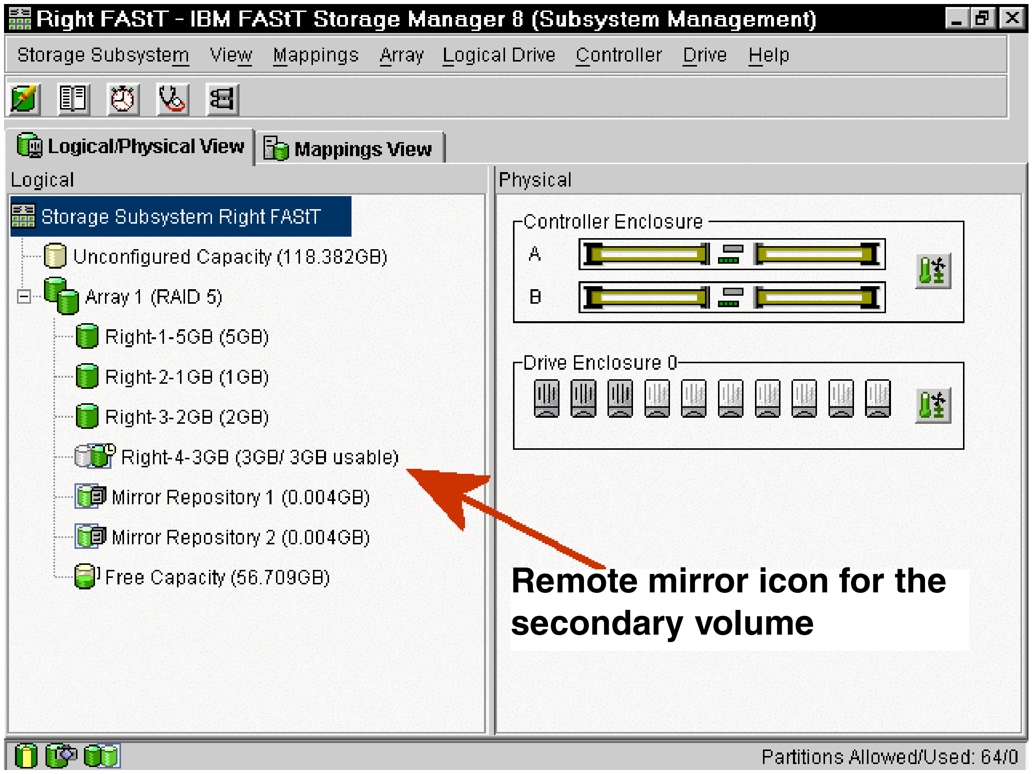

Figure 6-23 Secondary storage system

The mirror relationship that exists between the primary logical drive and the secondary logical drive can be examined by various methods. Such methods include using the storage subsystem profile, the Mirroring Properties window, and the View Associated Components window. Each of these methods is explained in the following sections.

Storage subsystem profile

The storage subsystem profile is one way to view information about any or all components of the storage subsystem. You can easily view details for all logical drives, such as the primary, secondary, and mirror repository logical drives, through the storage subsystem profile. The storage subsystem profile also contains specific information for components associated with mirror relationships.

Mirror properties

The Mirroring Properties window displays all the physical characteristics of a single logical drive in the mirror relationship. You can view the properties by highlighting the logical drive for the mirror, right-clicking, and selecting **Properties** as shown in Figure 6-24.

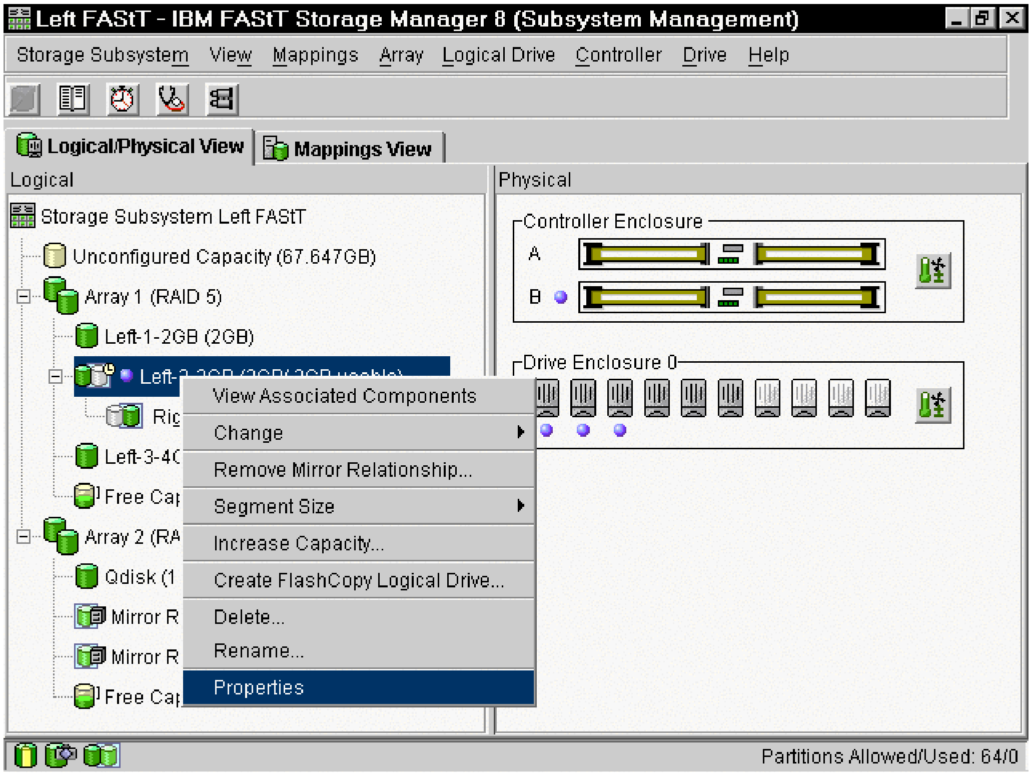

Figure 6-24 Selecting the drive properties

A window showing the base logical drive properties appears (Figure 6-25). The window displays the same information as the storage subsystem profile for the selected logical drive, but it is specific only to that logical drive.

In addition under the Mirroring tab, the synchronization progress is displayed if the selected logical drive is synchronizing data. This is normally when a newly-defined mirror relationship or in an existing mirror relationship after the broken mirror link is restored.

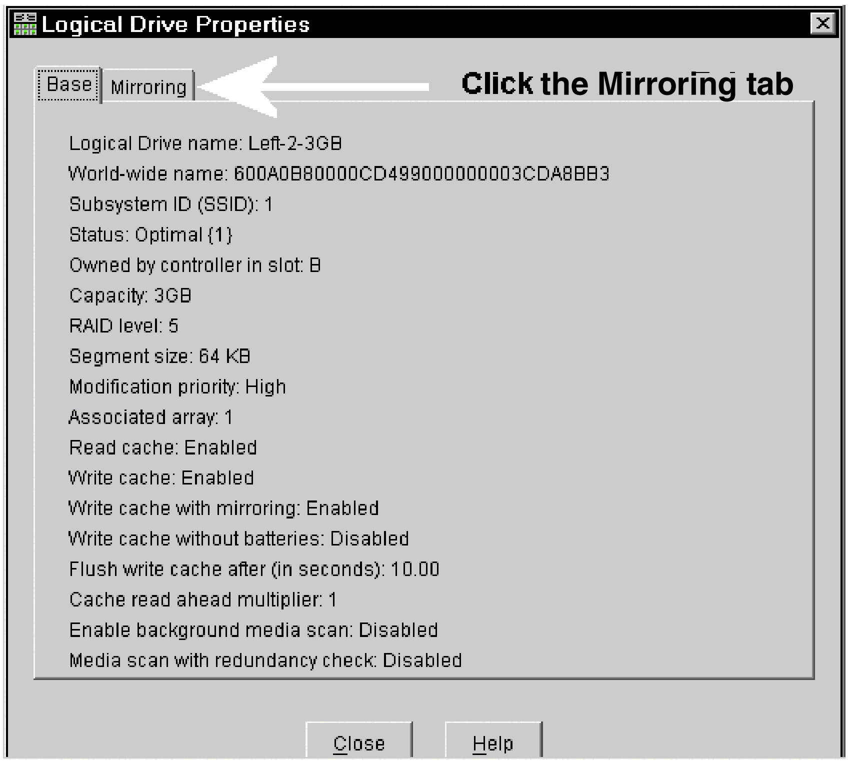

Figure 6-25 Base drive properties

For each logical drive that needs to be examined in the mirror relationship, click the **Mirroring** tab (Figure 6-26). On this tab, you van view the mirroring properties for each logical drive. You can also use the properties window to view the synchronization progress of a recently created mirror relationship.

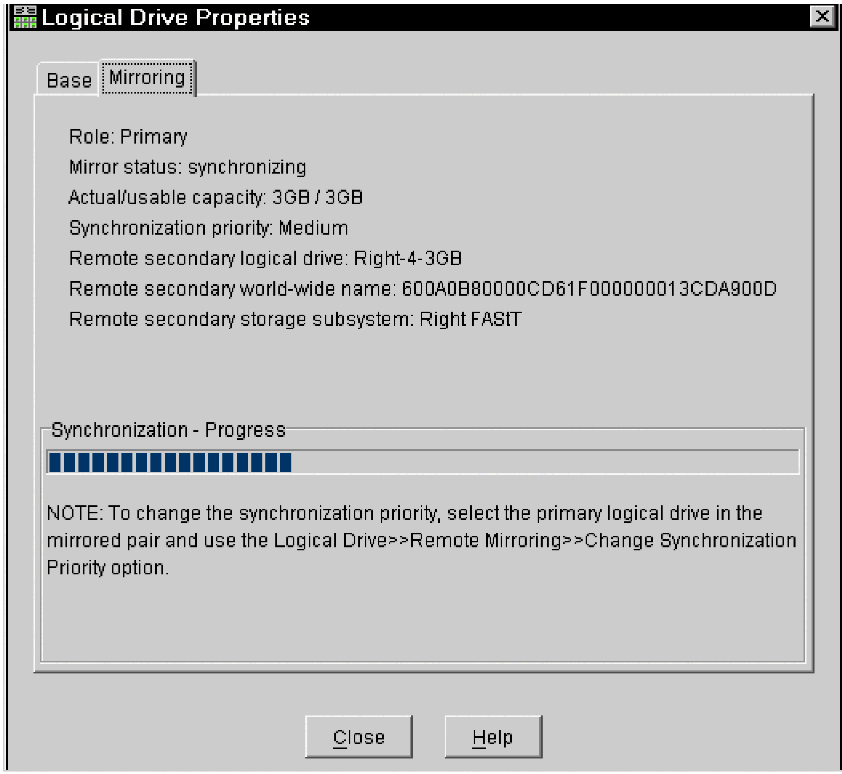

Figure 6-26 Mirror properties

The View Associated Components window

The View Associated Components window (Figure 6-27) provides a graphical representation of the logical drives participating in the mirror relationship. In addition, details are provided for all components, rather than just for the logical drive initially selected. To view all associated components in a mirror relationship, including primary, secondary, and mirror repository logical drives, complete these steps:

1. Select the primary or secondary logical drive in a mirror relationship.
2. Right-click and select **View Associated Components** .

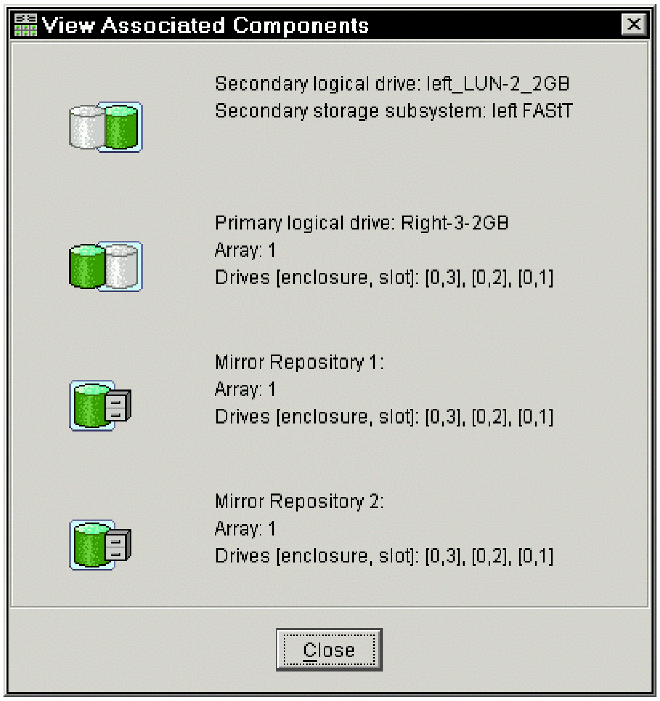

Figure 6-27 Mirror associated components

6.14.6 Changing the mirror synchronization priority

To change the synchronization priority level for a mirror relationship, complete the following procedure:

1. In the Logical/Physical View of the Subsystem Management window, select a primary logical drive of a mirror relationship.
2. Click **Logical Drive-> Remote Mirror-> Change-> Synchronization Priority**.
3. The Change Synchronization Priority window opens (Figure 6-21 on page 212). The primary logical drive chosen in the first step is selected by default in the Select logical drives selection area. Select one or more logical drives.
4. Select the synchronization priority level. The five levels are Lowest, Low, Medium, High, and Highest. All selected logical drives change to the same synchronization priority level. Click **OK**.
5. On Confirmation window, click **Yes**.
6. On the Completed window, click **OK**.

6.14.7 Removing mirror relationships

Removing a mirror relationship between a primary and secondary logical drive does not affect any of the existing data on either logical drive. The link between the logical drives is removed, but the primary logical drive continues normal I/O operation. You can perform this action for backup routines, particularly from the secondary logical drive. A mirror relationship between the two logical drives can be re-created unless one of the logical drives is deleted.

To remove a mirror relationship between two logical drives, complete the following steps:

1. From the Subsystem Management window, select a local primary logical drive or local secondary logical drive of a mirror relationship.

2. Click **Logical Drive-> Remote Mirror-> Remove Mirror Relationship**.
3. The Remove Mirror Relationship window displays all mirror relationships associated with this storage subsystem. Select one or more mirror relationships to be removed.
4. Carefully review the information presented in the Confirmation window. Click **Yes**.

6.14.8 Mapping secondary mirror drives

A secondary mirror drive can be LUN-mapped to a host or host group like any standard virtual drive. The difference is that the mirrored drive is invisible to any host system until it is promoted to a primary drive.

This option is useful in that it enables the administrator to pre-configure the mapping of the secondary mirrored drives prior to the changing roles. This makes the transition in a failover or backup situation easier by having the drive already mapped to the host.

6.15 Disaster recovery

As modern business pressures increasingly require 24-hour data access, system administrators are required to ensure that critical data is safeguarded against potential disasters. In preparing for a potential disaster, system administrators must develop a disaster recovery plan that details the procedures to prepare for and prevent disasters, as well as the actions required to respond to and recover from disasters if they occur.

Remote Volume Mirroring of data is one tool to add in recovery and business continuity. It is not a replacement for routine data backup to removable media. A mirror is only one copy of the data that is linked to the original. If the data on the original becomes corrupted due to programing or human error, the mirror that is a copy of the original may also be corrupted. If this occurs, the only option may be to recover from other data sources such as flash copies or tape backups.

6.15.1 Disaster plan

Before implementation of the Remote Volume Mirroring solution, you must take steps in preparation for a possible disaster or unrecoverable error. Correct planning and preparation at this stage ensures the effortless implementation of documented procedures that must be taken in for a to recovery from various disaster and error scenarios.

A disaster recovery plan should include the creation of a disaster or recovery toolkit. A *disaster recovery toolkit* contains essentially all the required information, resources, hardware, software, and documentation required to enable a component system administrator to recover the data and return the system to production with the minimum possible downtime. This toolkit must be in an easily accessible location.

The toolkit should be updated each time any system changes occur that may alter the steps required for the system recovery.

6.15.2 Key file backups and availability

Backing up critical data regularly is vital to ensure against disasters or an unrecoverable error. You must perform backups regardless of whether the Remote Volume Mirror option is in use.

Be sure to backup all critical data on both the primary storage subsystems participating in remote logical drive mirroring prior to creating a remote logical drive mirror. Perform periodic backups of the primary storage subsystem when remote logical drive mirrors are in use.

You can make backups to tape or disk while I/O continues using FlashCopy, which captures a point-in-time image of a logical drive. For more information on using FlashCopy logical drives, refer to Chapter 5, “FlashCopy and the FAStT Storage Server” on page 139, or the Subsystem Management window online help.

6.15.3 Recovery from a switch failure

Normally a high availability SAN fabric is designed to allow for a switch failure not being a single point of failure. If a switch fails, at least another switch is configured to be an alternate path for all I/O.

For example, a SAN fabric may be implemented with a single switch at each site that is responsible for the communication between the two sites for the Remote Volume Mirror option and all host-to-storage subsystem I/O. If a switch fails in this configuration, the system administrator must decide whether to perform a complete site failover. This decision must be based on the immediate availability of a backup switch or the time it will take to receive a replacement.

If a replacement switch is not readily available, one of the following temporary scenarios must occur until the switch is replaced:

- **Direct connection between host computers and storage subsystems**: This bypasses the switch and allows normal I/O to continue. All mirror relationships are suspended until a normal configuration is resumed. Since the storage subsystems have a finite number of host connections, all host computers might not be able to access the storage subsystem.
- **Entire site failover**: The switch at the secondary site is still functional. This enables all host computers to access the switch and storage subsystems. Mirror relationships are suspended until the primary storage subsystem is recovered. If a complete site failover is needed to continue normal operations, the secondary drives must be promoted to primary drives. See 6.16, “Reversing the roles of the primary and secondary drives” on page 221, for information on performing this task.

6.15.4 Recovery from a storage subsystem failure

To recover storage systems when one of the following situations occurs, you must perform a complete site failover from the primary site to the secondary site so that normal operations can continue:

- The primary storage subsystem is damaged or destroyed.
- The entire primary site is damaged or destroyed.

For a secondary site failure, no site failover is necessary. If the primary logical drive is not damaged, then a full synchronization occurs when the site is recovered.

6.15.5 Entire site failover

Entire site failover is necessary when an event occurs that inhibits normal operation to the primary site for an extended period of time. Follow these steps:

1. Manually reverse the role of the secondary logical drive of the mirror relationship to the primary role. For more information, see 6.16, “Reversing the roles of the primary and secondary drives” on page 221.

2. If storage partitioning is pre-configured so hosts can access the secondary logical drives, go to step 4. Otherwise, continue with step 3.
3. Configure the storage partitioning so that host to logical drive access is identical to the primary site configuration. When configured, continue with step 4. For more information on setting up storage partitioning, refer to 4.4, “Configuring storage partitioning” on page 105. Or consult the Subsystem Management window online help.
4. If the host computers at the secondary site are properly configured for host failover, go to step 6. Otherwise, continue with step 5.
5. Re-install host software and reconfigure settings as needed.
6. Run the hot_add utility on hosts that have only mapped logical drives from the secondary storage subsystem. Restart all hosts that have logical drive mappings from the primary and secondary storage subsystems. Then, resume normal operation with the secondary site acting as a new, fully-functioning primary site.

   The entire site failover is complete.

   > **Note:** To revert back to the primary site, you must reconfigure and recreate the Remote Volume Mirror environment using the old primary site as the secondary. The next step depends on the resulting status of the storage subsystem from the disaster.

7. Once the primary storage subsystem and all of the contained data are recovered and fully operational, recreate the mirror relationships.

6.15.6 Recreating a mirror relationship

When the damaged site is back online and properly configured, mirror relationships can be resumed. Re-create a mirror relationship by completing the following steps:

1. From the active secondary site, define a mirror relationship using the logical drive on the recovered primary site as the secondary logical drive. For more information, see 6.14, “Using Remote Volume Mirroring” on page 202.
2. Ensure storage partitioning is properly defined on the recovered primary site so that it can take over normal operation from the secondary site.
3. Ensure the host software is properly configured so that the host computers at the recovered primary site can take over I/O from the secondary site host computers.
4. After the full synchronization has completed, perform a manual role reversal so that the recovered primary site now possesses the active primary logical drive, and the secondary logical drive now exists on the secondary site. For more information, see 6.16, “Reversing the roles of the primary and secondary drives” on page 221.

The Remote Volume Mirror configuration is now optimal.

6.16 Reversing the roles of the primary and secondary drives

A role reversal is the process of promoting the secondary logical drive to be the primary logical drive within the mirrored logical drive pair. It also involves demoting the primary logical drive to be the secondary logical drive.

A role reversal is performed using one of the following methods:

- **Changing a secondary mirrored logical drive to a primary logical drive**. This option promotes a selected secondary logical drive to become the primary logical drive of the mirrored pair. It is used when an unrecoverable error has occurred. For step-by-step instructions, refer to the following section.

- **Changing a primary mirrored logical drive to a secondary logical drive**: This option demotes a selected primary logical drive to become the secondary logical drive of the mirrored pair. It is used during normal operating conditions. See 6.16.2, “Changing the primary to the secondary logical drive” on page 226.

Role reversals can also be performed using the Set Command in the Storage Manager Script window or the CLI. For more information, refer to the Enterprise Management window online help or see “Command line interface” on page 60.

6.16.1 Changing the secondary to the primary logical drive

A secondary logical drive is usually promoted to a primary logical drive role when an unrecoverable error occurs on the storage subsystem that contains the primary logical drive. Then the secondary logical drive needs to be promoted so that host computers can access data and normal operations can continue. This option is not available unless the Remote Volume Mirror option is activated.

> **Note:** When the secondary logical drive becomes a primary logical drive, any host computers that are accessing the logical drive through a logical drive-to-LUN mapping are now able to access the new drive on the secondary storage systems. This can be achieved via logical drive-to-LUN mapping of the secondary storage subsystem.

To promote a secondary drive to become a primary, complete the following steps:

1. Open the logical view of the Storage Management software for the secondary storage subsystem.
2. Select the secondary logical drive.
3. Right-click and click **Change-> Role to Primary** as shown in Figure 6-28.

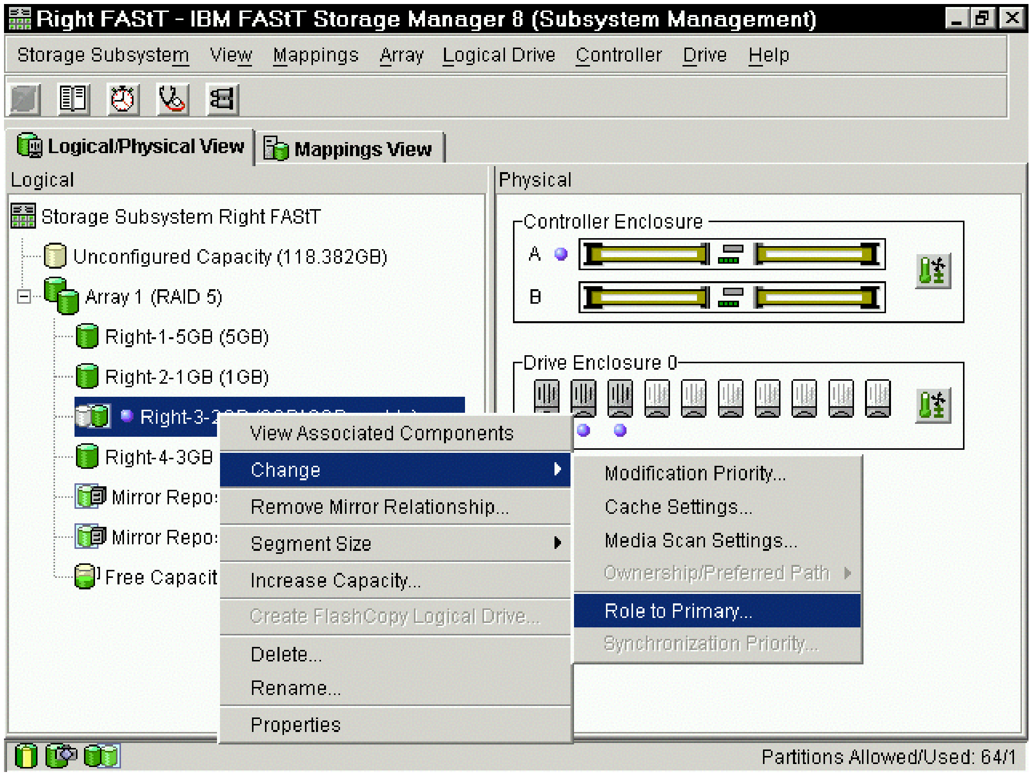

Figure 6-28 Promoting secondary to primary

If a communication problem between the secondary site and primary site prevents the demotion of the primary logical drive, an error message displays (Figure 6-29). However, you are given the opportunity to proceed with the promotion of the secondary logical drive, even though this leads to a Dual Primary Remote Mirror status condition.

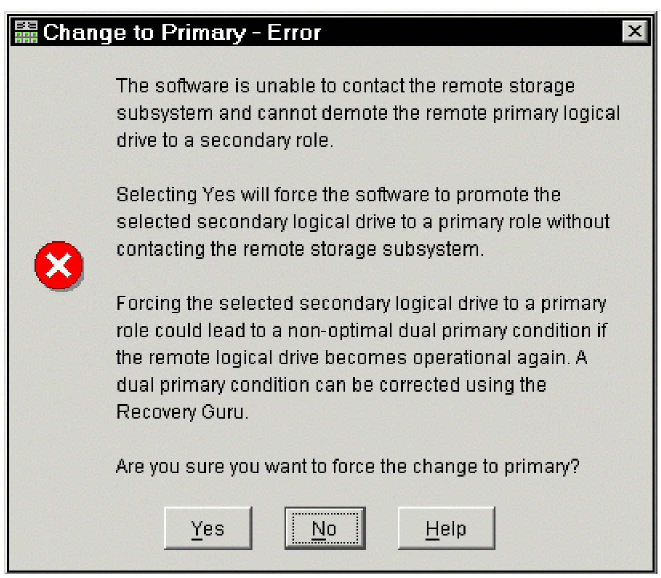

Figure 6-29 Change to Primary - Error message

Once the primary storage subsystem is recovered and the Fibre Channel links between the two subsystems are restored, the dual primary error occurs. The LUN on the original primary storage system displays a Mirror Unsynchronized status as shown in Figure 6-30.

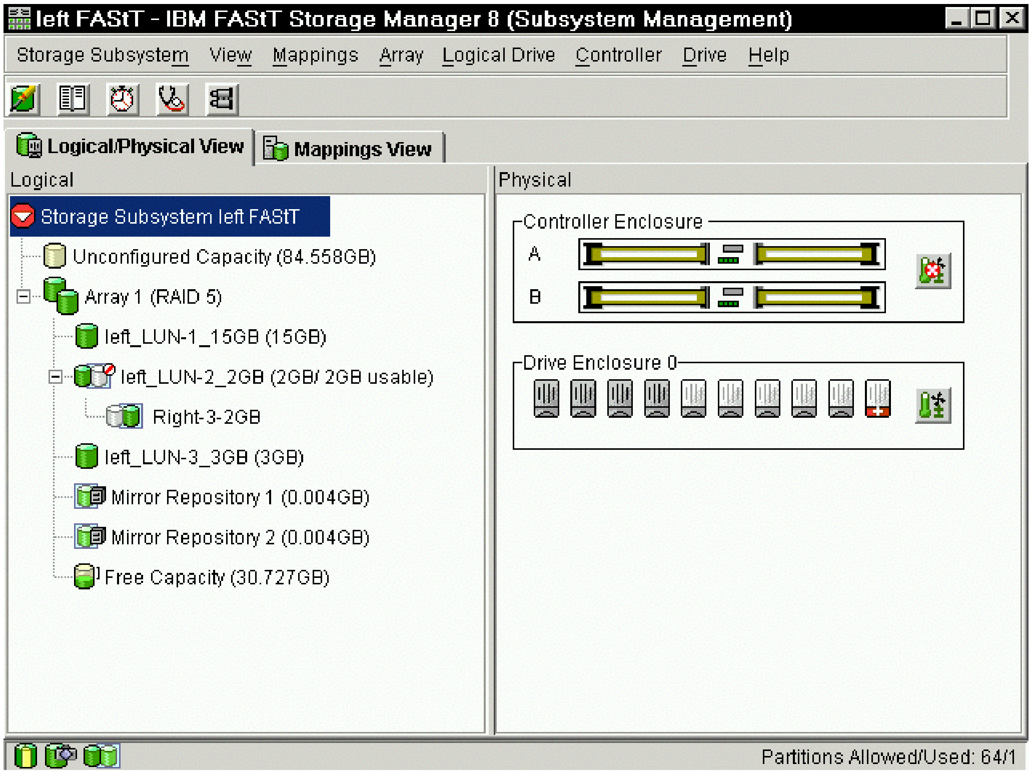

Figure 6-30 Primary drive unsynchronized

The Recovery Guru displays the “Dual Primary Logical Drive Conflict” error message (Figure 6-31), as well as an error description and recovery steps. To launch the Recovery Guru wizard and the window displaying the current error, select the **Dual Primary Logical Drive Conflict** error.

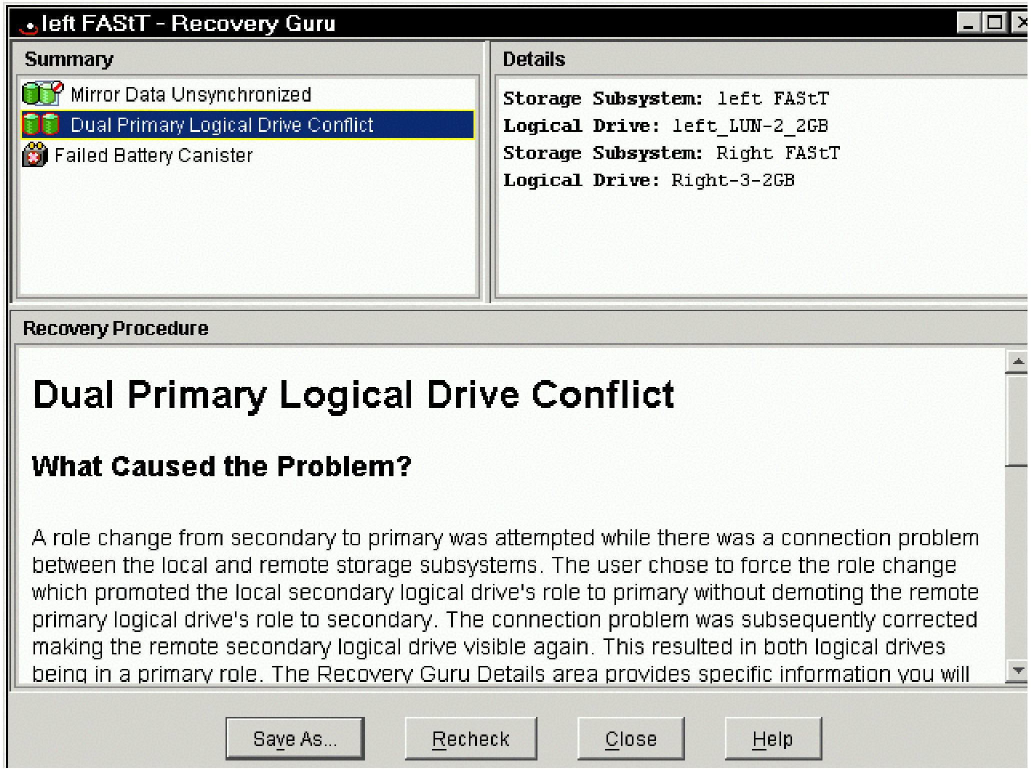

Figure 6-31 Recovery Guru window

The window displays the error and components at fault, gives a description of error, and lists the procedure to recover from the error.

Recovery steps

To recover from the error, follow these steps:

1. Recreate the mirror relationship:
   a. Delete the mirror relationship from either storage subsystem by highlighting either primary logical drive in its respective Subsystem Management window and clicking **Logical Drive-> Remote Mirroring-> Remove Mirror Relationship**.

      Note: The mirror relationship for the remote logical drive is not removed if there is a connection problem between the two storage subsystems. If the mirror relationship remains on the remote logical drive after the above action, highlight that logical drive from its respective Subsystem Management window and click **Logical Drive-> Remote Mirroring-> Remove Mirror Relationship**.

   b. In its respective Subsystem Management window, highlight the logical drive you want to be the primary logical drive. Click **Logical Drive-> Remote Mirroring-> Create**. Follow the instructions in the Create Remote Mirror wizard to re-create the mirror relationship.
2. Select **Recheck** to rerun the Recovery Guru to ensure that the failure is fixed.

6.16.2 Changing the primary to the secondary logical drive

Changing a primary logical drive to the secondary logical drive role is used for role reversals during normal operating conditions. Role reversal does not occur unless the Remote Volume Mirror option is activated.

Important:

- Any host computers that access the primary logical drive through a logical drive-to-LUN mapping can no longer read or write to the logical drive. When the primary logical drive becomes a secondary logical drive, only remote writes initiated by the primary controller are written to the logical drive. If a communication problem between the primary and secondary sites prevents the demotion of the primary logical drive, an error message displays. However, you are given the opportunity to proceed with the demotion of the primary logical drive, even though this leads to a Dual Secondary Remote logical drive mirror status condition. The recovery for the Dual Primary Remote Mirror status condition occurs when Remote Volume Mirror pairs are re-created. If you attempt to change the role of one of the primary drives, an error occurs.
- If the selected primary logical drive has associated FlashCopy logical drives, demoting this logical drive to a secondary role causes the associated FlashCopy logical drives to fail.

To demote a primary logical drive to the secondary logical drive role, follow these steps:

1. Select the primary logical drive in the Logical View. Then, right-click and select **Change-> Role to Secondary**.
2. The Change to Secondary window opens. Click **Yes**. The primary logical drive is demoted to be the secondary logical drive in the remote logical drive mirror.

When the controller owner of the primary logical drive is contacted, the secondary logical drive is automatically promoted to be the primary logical drive in the remote logical drive mirror.

6.17 Remote Volume Mirroring solution design

A SAN design is normally based on storage, flexibility, performance, and redundancy requirements. The Remote Volume Mirror option is one piece of the solution.

This section presents three sample configurations of how the Remote Volume Mirror option can be implemented. The level of redundancy is determined by the type of configuration you choose to use.

The three sample solutions are:

- **Simple departmental with minimum redundancy**: This solution shows a simple departmental solution implemented using two Fibre Channel switches and two FAStT storage systems.
- **Intersite with redundant fabric**: Similar to the previous solution except for the storage systems being located in two physically separate sites.
- **Intersite with FlashCopy drives and tape backup**: A solution that provides for the highest level of redundancy and instant time zero drive backups to disk and tape.

When designing a SAN storage system, it is good practice to complete the following steps:

1. Produce a statement that outlines the solution requirements that can be used to determine the type of configuration you need. It should also be used to cross-check that the solution design delivers the basic requirements. The statement should have easily defined bullet points covering the requirements, for example:
   - Required capacity
   - Required redundancy levels
   - Backup and restore windows
   - Type of data protection needed
   - Network backups
   - LAN free backups
   - Serverless backups
   - FlashCopy
   - Remote volume mirroring
   - Host and operating system types to be connected to SAN
   - Number of host connections required
2. Produce a hardware checklist. It should cover such items that you require you to:
   - Ensure that the minimum hardware requirements are met.
   - Make a complete list of the hardware requirements including the required premium options.
   - Ensure your primary and secondary storage subsystems are properly configured.
   - Ensure that your Fibre Channel switches and cables are properly configured. The remote links must be in a separate zone.
3. Produce a software checklist to cover all the required items that need to be certified and checked. It should include such items that require you to:
   - Ensure that data on the primary and secondary storage subsystems participating in remote logical drive mirroring is backed up.
   - Ensure that the correct version of firmware and storage-management software are installed.
   - Ensure that the Remote Volume Mirror option is enabled on both the primary and secondary storage subsystems.
   - Ensure that the Remote Volume Mirror option is activated and that a mirror repository logical drive is created for each controller on the primary storage subsystem.
   - Ensure that the required primary and secondary logical drives are created on the primary and secondary storage subsystems.

6.17.1 Solution 1: Simple department with minimum redundancy

This is a simple departmental configuration, a low-cost configuration. It consists of two storage subsystems and two Fibre Channel switches connected with a Fibre Channel fabric as shown in Figure 6-32. The primary storage subsystem and secondary storage subsystem may have a maximum connection distance of up to 500 m (0.32 mi.) and can be located in the same building or in the same room.

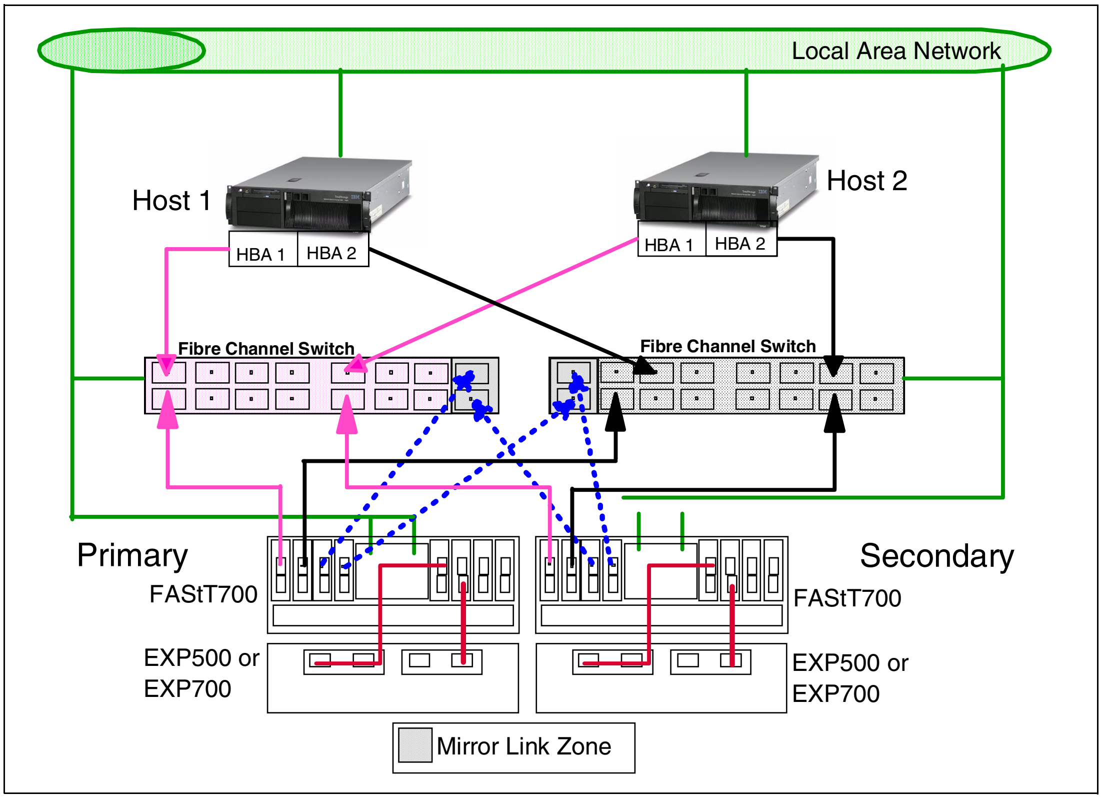

Figure 6-32 Simple departmental solution

The solution provides for fabric switch redundancy, in that a failed switch causes a failover of the logical disk to the opposite controller where the host access is via the second Fibre Channel switch.

Remote Volume Mirroring can be configured in either direction. That is, there may be some with primary logical disks on the primary storage system and secondary disks on the secondary, as well as having primary disks on the secondary system and secondary drives on the primary. Considering that, with this configuration, the two storage subsystems are usually in the same building, spreading the production load over both subsystems may offer the best performance.

If redundancy was *not* a requirement, this solution could be implemented using only one Fibre Channel switch with the correct zoning. You must be aware that a single switch failure can stop all I/O activity.

6.17.2 Solution 2: Intersite with redundant fabric

This configuration is similar to the first solution, except for the storage systems being physically located in different sites. The configuration consists of two storage subsystems and two Fibre Channel switches connected with a Fibre Channel fabric as shown in Figure 6-33. Using standard Fibre Channel long wave Gigabit Interface Converters (GBICs) and single mode cable, the primary storage subsystem and secondary storage subsystem have a maximum connection distance of up to 10 km. (6.25 mi.).

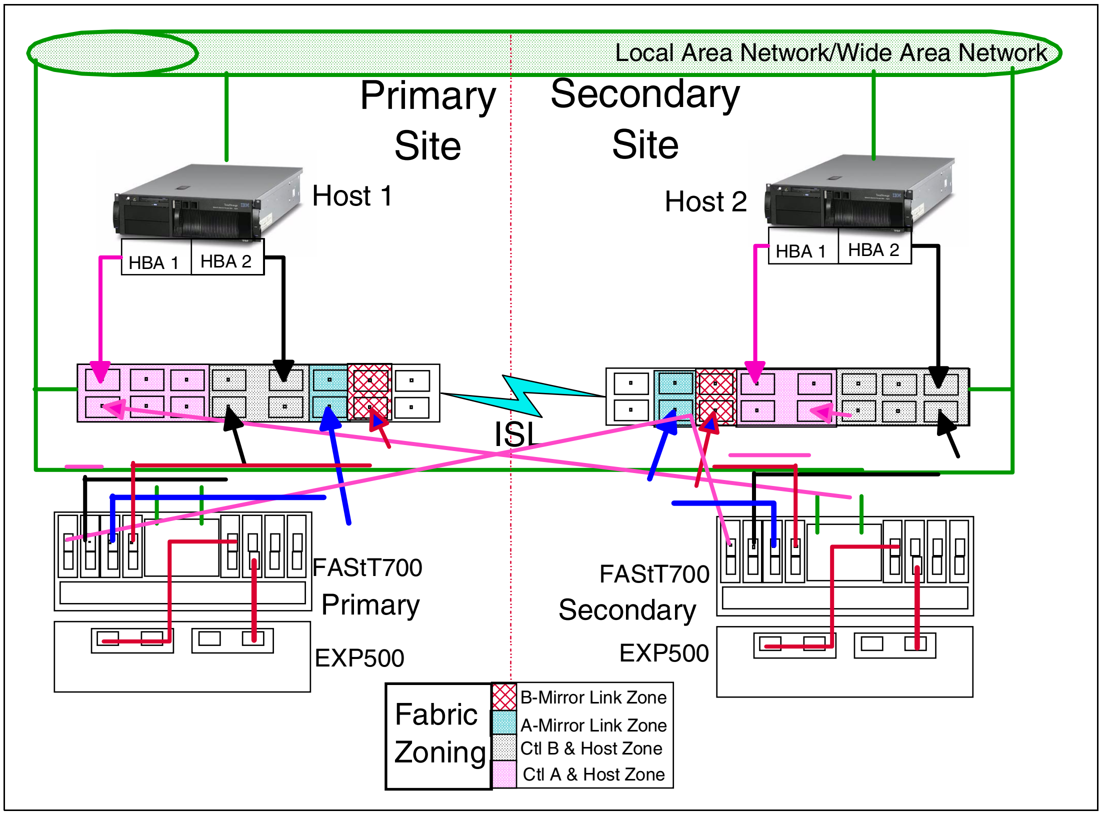

Figure 6-33 Intersite redundant fabric

The configuration provides for redundancy based on a site failover. The fabric has been zoned into four separate zones and each switch share in all four zones.

We recommend the zone in this solution as best practice. The only mandatory zoning is for the Remote Volume Mirror links to be in a zone separate from the FAStT control units and hosts. During testing for this redbook, we found that unpredictable results were obtained when the host scanned devices if the mirrors were in the same zone as the hosts and controllers.

6.17.3 Solution 3: Intersite with FlashCopy drives and tape backup

The highest availability configuration is fully redundant and includes two storage subsystems and four Fibre Channel switches connected with a Fibre Channel fabric, as shown in Figure 6-34. The primary storage subsystem and secondary storage subsystem have a maximum connection distance of up to 10 km. (6.25 mi.).

Note: For performance and high availability configuration, use two 10 km. (6.25 mi.) inter-switch links (ISL).

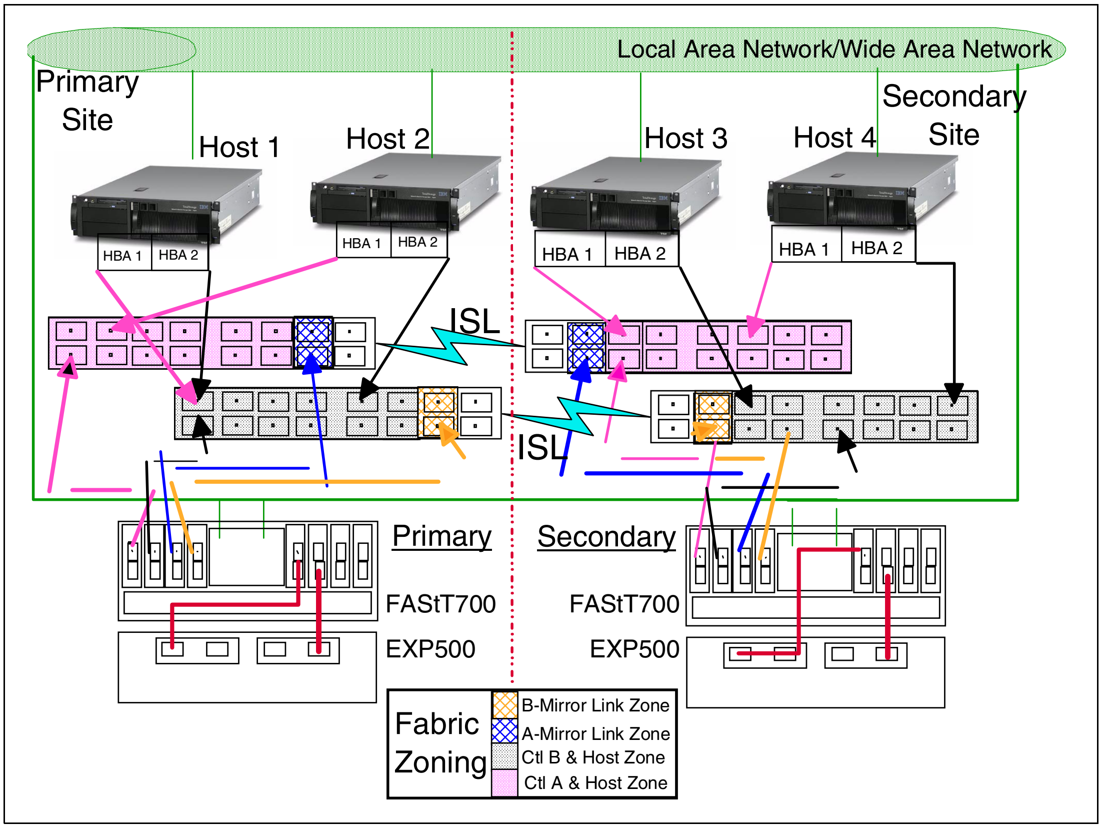

Figure 6-34 intersite high availability solution

Apart from the greater redundancy of dual switch fabric in each site, a greater number of host ports are now available, allowing greater flexibility in the use and connectivity.

With this type of configuration, consider putting the primary drives on the secondary (remote site). This offers several advantages. The first is if the primary site fails, the standby servers in the secondary site can be attached to the original primary disk with a simple procedure.

With the primary drives on the remote site and the secondary drive in the local site give an up-to-date copy at all times, it is still possible through programming or human error to corrupt data and the data corruption to be mirrored to the secondary drives. You now have several different options. You can:

- Make a FlashCopy of the data on the primary drive.
- Make a tape backup of the data from the primary drive.
- Combine both where a FlashCopy is performed and then perform a tape backup of the copied drive.

7

Migrating from an earlier version of FAStT Storage Manager

This chapter explains how to update the Storage Manager software on a host from an earlier release to the latest version 8.2. For this entire procedure, we used a pre-GA version that may exhibit subtle differences from the final release. This chapter also shows you how to update the storage server controller firmware and NVSRAM.

7.1 Prerequisites for migration

Before you start the migration, be sure that your environment meets the following requirements:

- The host software must be updated first. Otherwise, you can lose access to the FAStT Storage Server.
- For FAStT700 (machine type 1742), the controller microcode must be at version 5.00.04.xx or higher.
- For FAStT500 or FAStT200 (machine types 3552 or 3542), the controller microcode must be at version 4.00.02.xx
- If updating FAStT500 or FAStT200 from controller firmware 4.xx.xx.xx (Storage Manager 7), contact IBM or your IBM Business Partner because this is a chargeable option. The publicly available update only updates from version 8.0 to version 8.2.

> **Important:** If you purchased the Storage Manager 8 upgrade for FAStT500 and your package contains version 8.0, download the version 8.2 software and contact IBM level 2 support for the version 8.2 microcode. Do *not* try to update to version 8.0 because there is a known bug that can corrupt your array configuration.

- The host bus adapter firmware and drivers should be at the latest level. Refer to:

  `http://www.pc.ibm.com/qtechinfo/MIGR-42081.html`

- We recommend that you check for updates to your drive microcode and enclosure ESM modules. For further information, see 4.2.4, "Updating the drive and ESM board microcode" on page 94.

7.2 Updating the FAStT host software

This section describes how to update FAStT Storage Manager, both generically and in a Microsoft Windows Cluster environment.

7.2.1 Updating from a previous version of Storage Manager

Be sure that you are familiar with the steps required in a new installation before you continue with the following update procedure.

> **Important:** Do not update your storage subsystem firmware until you update the storage-management software on all host systems and management stations. If you update your storage subsystem controller firmware to version 05.20.xx.xx, you cannot communicate with the controller until you update to the Storage Manager 8.2 software on your management station and host systems. Previously installed versions of Storage Manager do not recognize controllers running version 05.20.xx.xx firmware.

To perform an update from Storage Manager version 7.0, 7.01, 7.02, 7.10, or 8.0 to Storage Manager 8.2, complete the following steps. Not all steps are needed on all operating systems:

1. Uninstall the components from the previous version of the storage-management software in the following order:

   a. Storage Manager Agent
   b. Storage Manager Utility

c. Redundant Disk Array Controller
d. Storage Manager Client

2. Verify that the IBM host bus adapter device driver versions and BIOS are current. If they are not current, upgrade them. For the latest downloadable files matrix, refer to:

   http://www.pc.ibm.com/qtechinfo/MIGR-42081.html

3. If necessary, install the FAStT Runtime Environment (UNIX only). Refer to 4.1, "Driver and host software installation" on page 66.
4. Install the Storage Manager 8.2 Client. Refer to 4.1, "Driver and host software installation" on page 66, for instructions.
5. If required by the operating system, install the Redundant Disk Array Controller. Refer to 4.1, "Driver and host software installation" on page 66, for instructions.
6. If supported, install the Storage Manager 8.2 Agent. Refer to 4.1, "Driver and host software installation" on page 66, for instructions.
7. Install the Storage Manager 8.2 Utility (not applicable to AIX). Refer to 4.1, "Driver and host software installation" on page 66, for instructions.
8. Update the FAStT Storage Server microcode. See 7.3, "Updating the FAStT controller microcode" on page 236.

7.2.2 Updating Storage Manager in a Microsoft Cluster (MSCS) environment

You can update Storage Manager in a Microsoft Cluster environment in one of two ways:

- **Scheduling downtime and taking the cluster offline**: This is the IBM preferred method because it maintains integrity of the cluster and of cluster data. The basic steps are to:
  a. Stop the cluster service on both nodes. Set the cluster service startup type to **manual** on both nodes.
  b. Shut down node B.
  c. Update node A by following the instructions in 7.2.1, "Updating from a previous version of Storage Manager" on page 232.
  d. Shut down node A.
  e. Restart node B.
  f. Update node B by following the instructions in 7.2.1, "Updating from a previous version of Storage Manager" on page 232.
  g. Set the cluster service on node B to **automatic** startup.
  h. Start the cluster service on node B.
  i. Restart node A.
  j. Set the cluster service on node A to **automatic** startup.
  k. Start the cluster service on node A.
- "Rolling upgrade" method:

> **Note:** If you have applications installed in the cluster server environment that do not support a rolling upgrade, you must perform one of the following tasks:
>
> - Move those resources *offline* before the upgrade and then move them back *online* after the upgrade.
> - Perform a scheduled upgrade. This requires scheduling downtime on your cluster server to upgrade the controller firmware and storage-management software.

a. Using the cluster administrator on node A, move all resources to node B.
b. Pause node A.
c. Update node A by following the instructions in 7.2.1, “Updating from a previous version of Storage Manager” on page 232.
d. Resume node A.
e. Using cluster administrator on node B, move all resources to node A.
f. Pause node B.
g. Update node B by following the instructions in 7.2.1, “Updating from a previous version of Storage Manager” on page 232.
h. Resume node B.
i. Move resources back to node B as required.

Regardless of the option you choose, when are are finished, update the FAStT Storage Server microcode as explained in 7.3, “Updating the FAStT controller microcode” on page 236.

7.2.3 Upgrading from Windows NT 4.0 to Windows 2000

If you have an existing storage subsystem running Windows NT 4.0, you can upgrade to Windows 2000 if the following conditions apply:

- Your storage subsystem is not the startup (boot) device.
- Your controllers are running firmware versions 04.00.00.00 or later.

Attention: To avoid possible data loss, you must uninstall any previous versions of the storage-management software before you upgrade your operating system. All storage-management configuration information is deleted during the uninstallation process. A complete re-installation of the storage-management software is required.

To upgrade from Windows NT 4.0 to Windows 2000, complete this process:

1. Verify that the storage subsystem contains no system-dependent files or directories, such as paging files.
2. Uninstall Storage Manager. The Storage Manager components should be uninstalled in the following order:
   a. Storage Manager Agent
   b. Storage Manager Utility
   c. Redundant Disk Array Controller
   d. Storage Manager Client
3. Use the instructions from Microsoft to upgrade to Windows 2000.
4. Re-install Storage Manager. Refer to the instructions in 4.1.1, “Microsoft Windows NT 4.0 and Windows 2000 installation” on page 66, or to the procedure that came with the version of the storage-management software you are currently running.
5. If you are upgrading to Storage Manager v8, update the FAStT Storage Server microcode as explained in 7.3, “Updating the FAStT controller microcode” on page 236.

7.2.4 Upgrading from Windows NT4 to Windows 2000 in a cluster setup

To ensure that the services and resources offered by the cluster are always available, you must perform a *rolling upgrade*. A rolling upgrade requires upgrading the operating system and Storage Management software first in node A and then in node B and finally upgrading the FAStT Storage Server microcode.

Note: If you have applications installed in the cluster server environment that do not support a rolling upgrade, you must perform one of the following tasks:

- Move those resources *offline* before the upgrade and then move them back *online* after the upgrade.
- Perform a scheduled upgrade. This requires scheduling downtime on your cluster server to upgrade the controller firmware and storage-management software.

To upgrade from Windows NT 4.0 Enterprise Edition to Windows 2000 Advanced Server, follow these steps:

1. Using the cluster administrator on node A, move all resources to node B.
2. Pause node A.
3. Uninstall all Storage Manager components. The Storage Manager components should be uninstalled in the following order:
   a. Storage Manager Agent
   b. Storage Manager Utility
   c. Redundant Disk Array Controller
   d. Storage Manager Client
4. Follow the instructions from Microsoft to upgrade node A from Windows NT Enterprise Edition to Windows 2000 Advanced Server.

   Note: During the upgrade, the existing cluster server installation is detected, and clustering for Windows 2000 Advanced Server is installed. When the upgrade is complete, node A rejoins the cluster and remains paused.

5. Install Windows 2000 service pack 2.
6. Verify that host bus adapter drivers are at the latest level.
7. Re-install Storage Manager. Refer to the instructions in 4.1.1, "Microsoft Windows NT 4.0 and Windows 2000 installation" on page 66, or to the procedure that came with the version of the storage-management software you are currently running.
8. Resume node A.
9. Using cluster administrator on node B, move all resources to node A.
10. Pause node B.
11. Uninstall all Storage Manager components. The Storage Manager components should be uninstalled in the following order:
    a. Storage Manager Agent
    b. Storage Manager Utility
    c. Redundant Disk Array Controller
    d. Storage Manager Client
12. Follow the instructions from Microsoft to upgrade node B from Windows NT Enterprise Edition to Windows 2000 Advanced Server.

Note: During the upgrade, the existing cluster server installation is detected, and clustering for Windows 2000 Advanced Server is installed. When the upgrade is complete, node B rejoins the cluster and remains paused.

13. Install Windows 2000 service pack 2.
14. Verify that the host bus adapter drivers are at the latest level.
15. Re-install Storage Manager. Refer to the instructions in 4.1.1, “Microsoft Windows NT 4.0 and Windows 2000 installation” on page 66, or to the procedure that came with the version of the storage-management software you are currently running.
16. Resume node B.
17. Move the resources back to node B as required.
18. If you are upgrading to Storage Manager version 8, upgrade the FAStT Storage Server microcode as explained in 7.3, “Updating the FAStT controller microcode” on page 236.

7.3 Updating the FAStT controller microcode

The microcode of the FAStT Storage Server consists of two packages:

- Actual firmware
- NVSRAM package, including the settings for booting up the FAStT Storage Server

The NVSRAM is similar to the settings in the BIOS of a host system. The firmware and the NVSRAM are not independent. Ensure you install the correct combination of the two packages.

1. The update procedure needs two independent connections to the FAStT Storage Server – one for each controller. It is not possible to perform a microcode update with only one controller connected, so both controllers must be accessible either via Fibre Channel or Ethernet. To update over Fibre Channel, your host system must support in-band management, which Linux and AIX do not support. For AIX and Linux, you are limited to updating over the Ethernet.

   If you plan to update via Fibre Channel, make sure that you have a multipath driver installed on your management host, for example the FAStT RDAC package. This is necessary since the access logical drive moves from one controller to the other during this procedure and the FAStT Storage Server must be manageable during the entire time.

   Attention: Any power or network/SAN interruption during the update process may lead to configuration corruption. Therefore, do not turn off the power to the FAStT Storage Server or the management station during the update. If you are using in-band management and have Fibre Channel hubs or managed hubs, make sure any SAN-connected devices are not powered up during the update, because this causes a loop initialization process and interrupts the process.

2. Make sure you update the host software on all hosts prior to updating the controller firmware. See 7.2.1, “Updating from a previous version of Storage Manager” on page 232.
3. Start the IBM Storage Manager client and select the FAStT Storage Server you want to update from within the Enterprise Management window. See Figure 7-1. Note that the title bar displays as IBM FAStT Storage Manager 8.

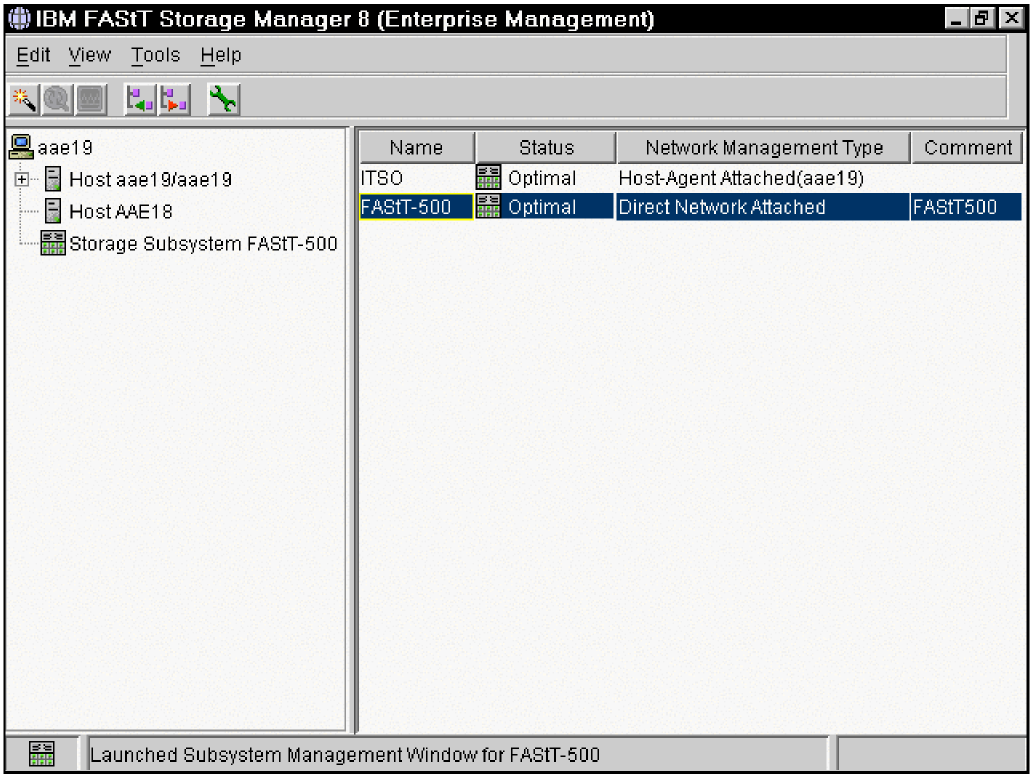

Figure 7-1 The Enterprise Management window

4. The Subsystem Management window opens. Make sure your profile of the storage server is up to date. Click **Storage Subsystem-> Profile** and specify a name and path. Do not store the profile on the FAStT Storage Server itself. Instead, choose a *local* disk. Figure 7-2 shows a sample of the profile output.

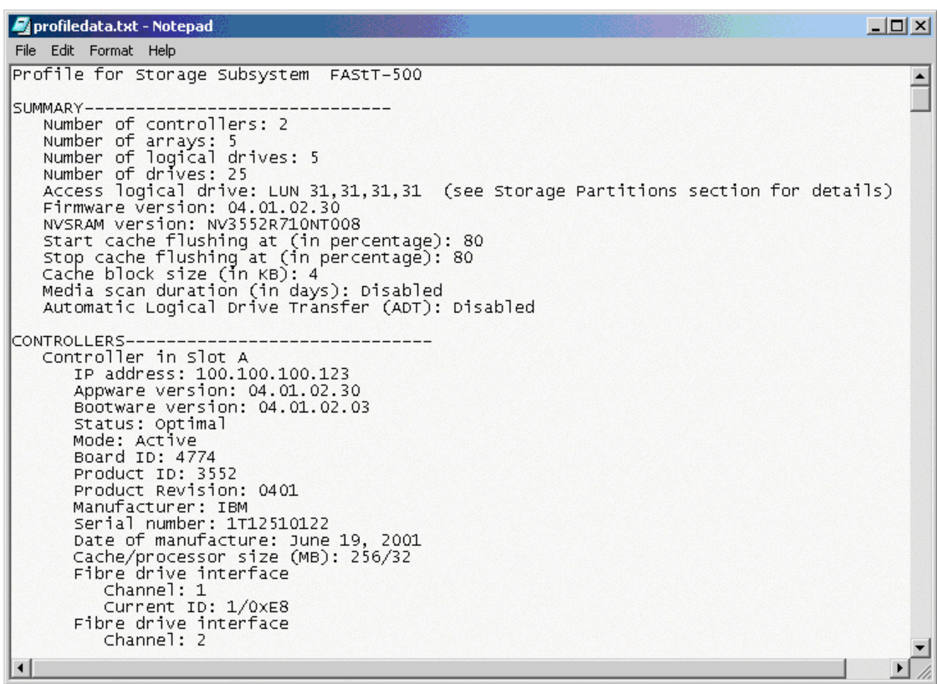

profiledata.txt - Notepad

File Edit Format Help

```
Profile for Storage Subsystem  FASTT-500

SUMMARY------------------------------
   Number of controllers: 2
   Number of arrays: 5
   Number of logical drives: 5
   Number of drives: 25
   Access logical drive: LUN 31,31,31,31  (see Storage Partitions section for details)
   Firmware version: 04.01.02.30
   NVSRAM version: NV3552R710NT008
   Start cache flushing at (in percentage): 80
   Stop cache flushing at (in percentage): 80
   Cache block size (in KB): 4
   Media scan duration (in days): Disabled
   Automatic Logical Drive Transfer (ADT): Disabled

CONTROLLERS------------------------------
   Controller in Slot A
      IP address: 100.100.100.123
      Appware version: 04.01.02.30
      Bootware version: 04.01.02.03
      Status: Optimal
      Mode: Active
      Board ID: 4774
      Product ID: 3552
      Product Revision: 0401
      Manufacturer: IBM
      Serial number: 1T12510122
      Date of manufacture: June 19, 2001
      Cache/processor size (MB): 256/32
      Fibre drive interface
         Channel: 1
         Current ID: 1/0xE8
      Fibre drive interface
         Channel: 2
```

Figure 7-2 Sample of the profile data before update

5. For a visual check of the current controller firmware, etc., right-click each controller in the physical view in turn and select **Properties**. You see a Properties window like the example in Figure 7-3. Note that all this information is saved in the profile.

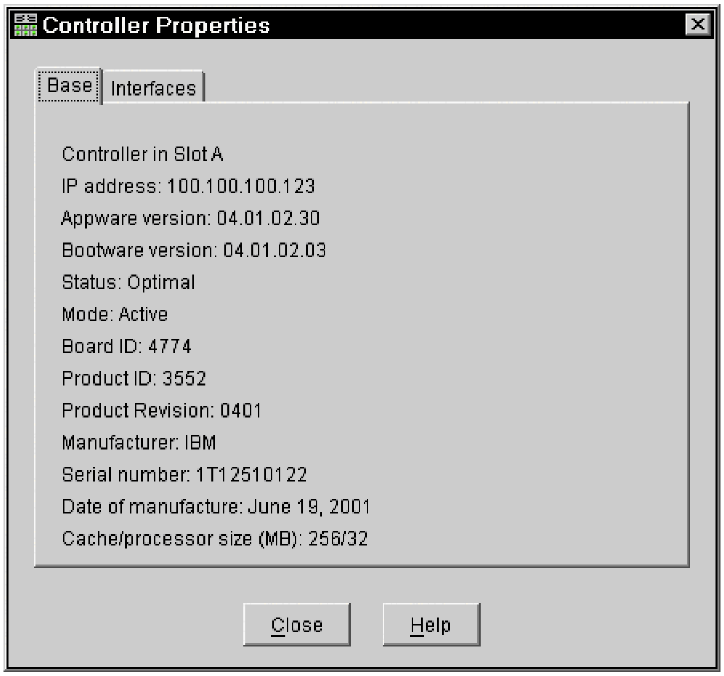

Figure 7-3 Controller Properties before update

6. Make a visual note of the defined mappings to which you can compare after the upgrade. This is shown in Figure 7-4. Again this information is stored in the profile.

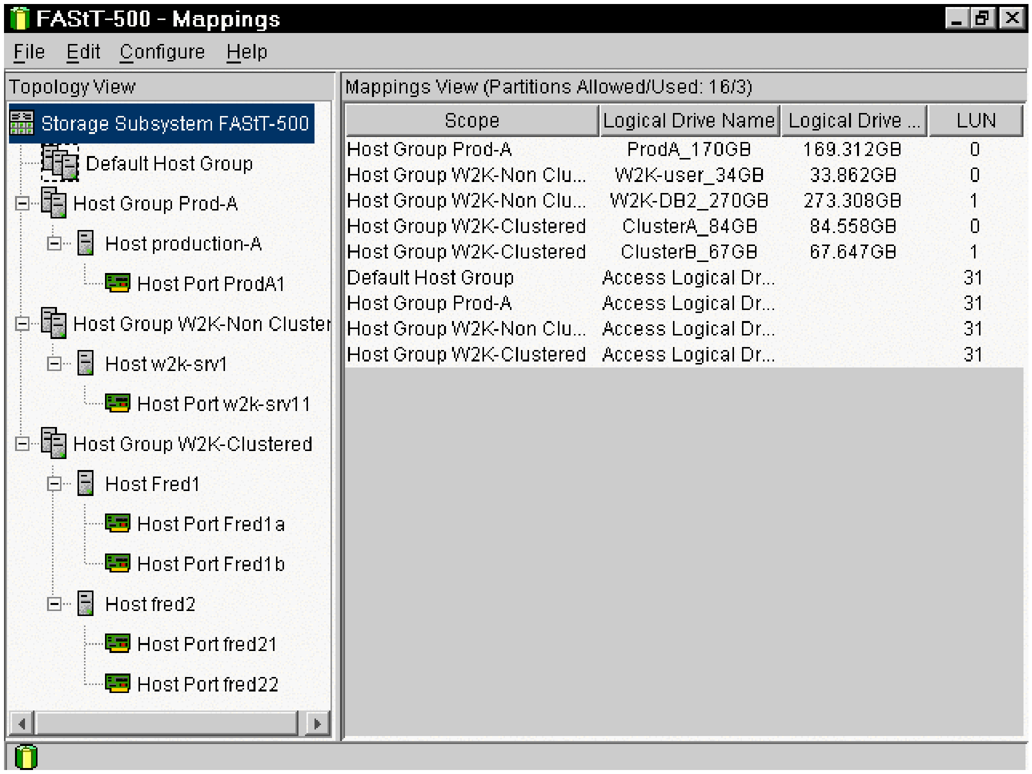

Figure 7-4 Mappings before the update

7. You can verify the premium features key you installed. If your key states that storage partitioning (up to maximum) is enabled, then an update from Storage Manager v7 to Storage Manager v8 will enable all 64 storage partitions. However, if you find that, for example, storage partitioning (up to eight) is enabled, then you have to contact IBM level 2 support to obtain a new key to support 64 storage partitions. The premium features list is shown in Figure 7-5. See also 4.6.10, "Handling premium features" on page 127, for instructions on obtaining a new feature key.

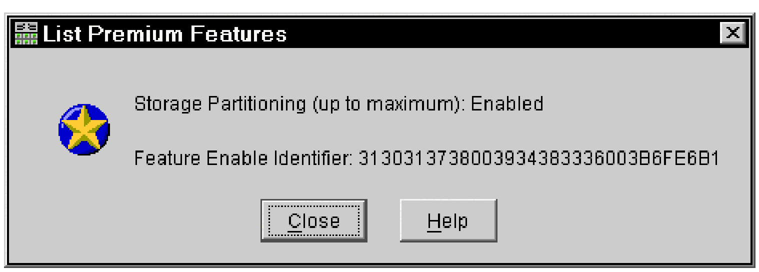

Figure 7-5 List Premium Features

Now you must download the new firmware and NVSRAM as explained here:

1. Download the firmware:

   a. To start the download process for the firmware, in the Subsystem Management window, select the **Storage Subsystem** menu and click **Download-> Firmware**. See Figure 7-6.

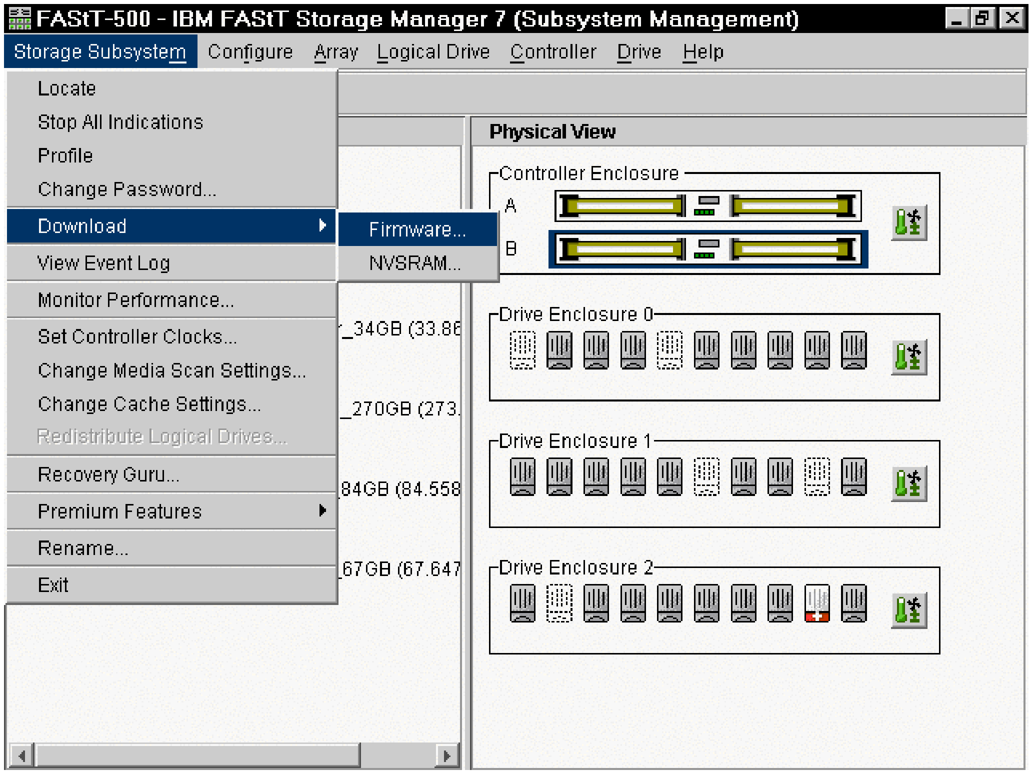

Figure 7-6 Download Firmware

> **Note:** The title bar shows Storage Manager 7 even though we're using Storage Manager 8. This is because the version of firmware is detected and the correct management window loaded that supports only the features in that firmware.

b. You are now prompted for the file containing the firmware you want to download. In the top part of the window, you can see which version of firmware and NVSRAM are currently loaded. Underneath it shows which firmware revision is in the file you selected. See Figure 7-7.

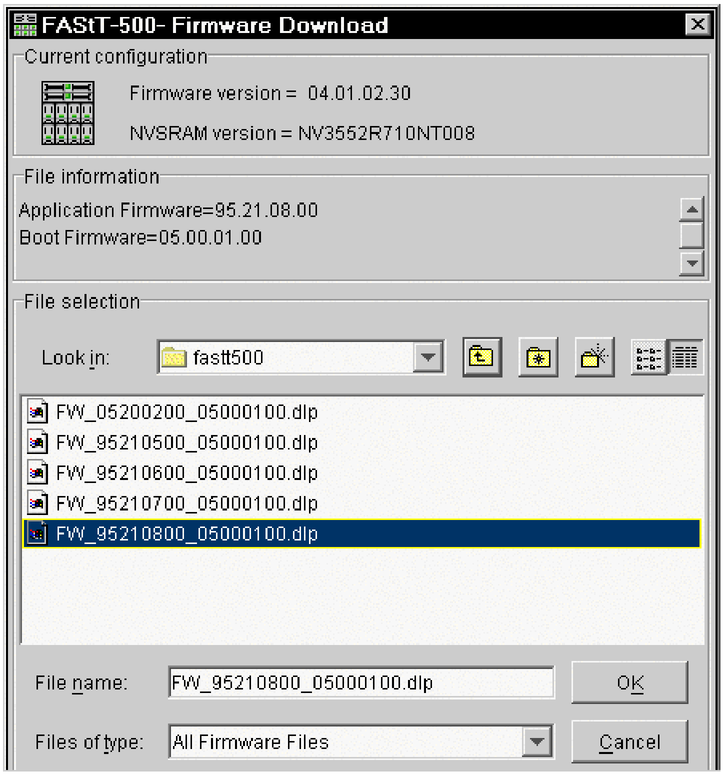

Figure 7-7 Selecting the firmware

c. Click **OK**.

d. You see a window (Figure 7-8) where you have to confirm that all the requirements are met. Click **Yes** to confirm the download.

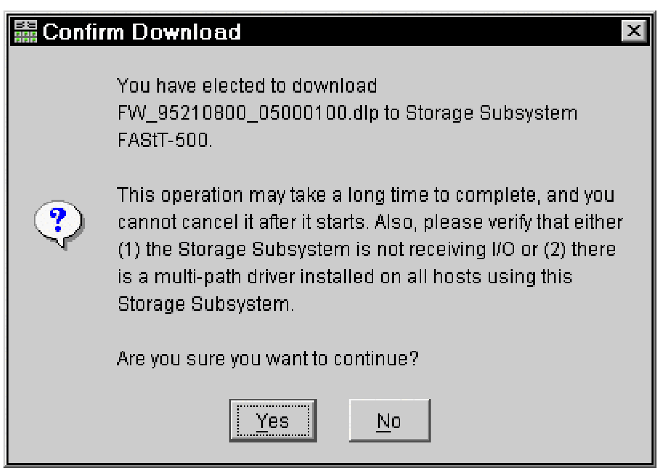

Figure 7-8 Confirm Download

If you are trying to update Storage Manager v7 to Storage Manager v8 without buying the upgrade, for example, by using the publicly downloadable code, you see a file selection error message (Figure 7-9).

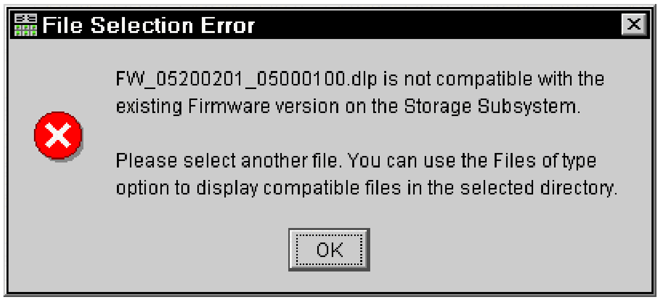

Figure 7-9 Firmware selection error

Otherwise, you see the downloading status window (Figure 7-10).

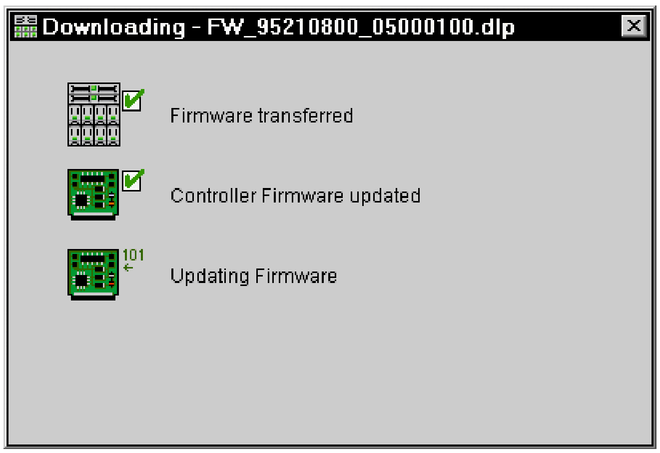

Figure 7-10 Downloading firmware

2. When the download is finished, you receive a message (Figure 7-11) telling you that new firmware has been detected and that you should relaunch the Subsystem Management window.

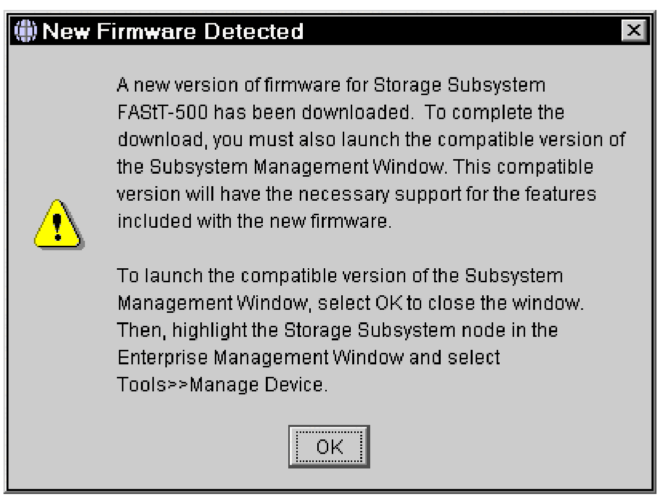

Figure 7-11 New Firmware Detected

3. A new window opens that is compatible with the version of firmware you downloaded. You must now update the NVSRAM.
   a. Click **Storage Subsystem-> Download-> NVSRAM**. See Figure 7-12.

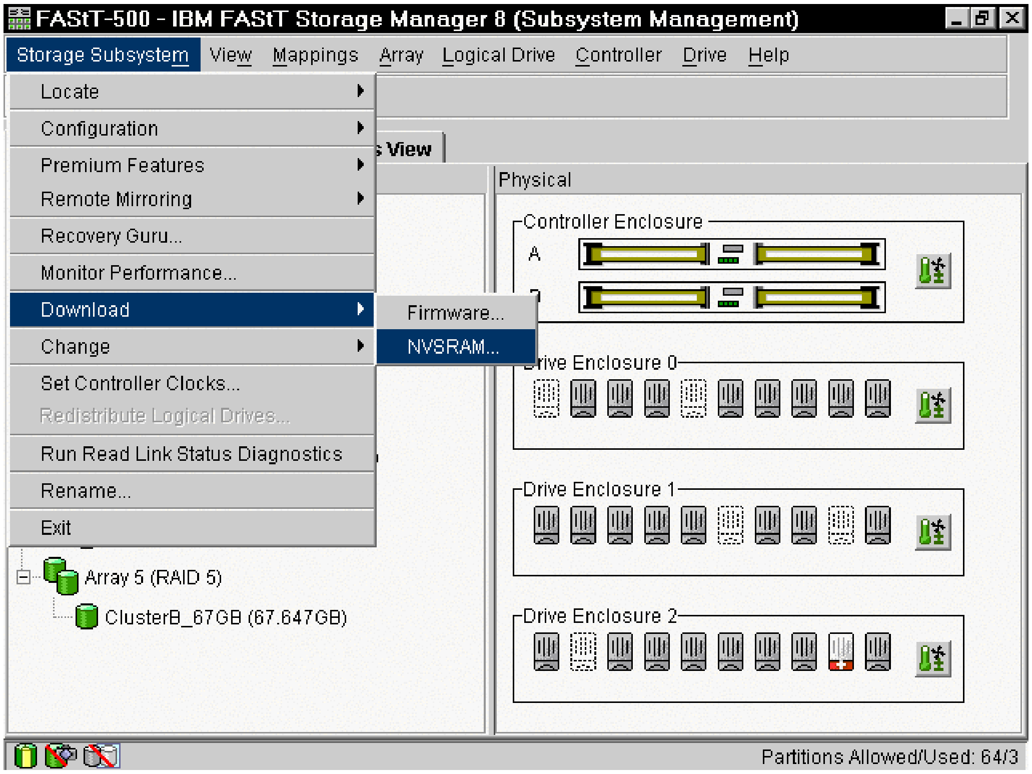

Figure 7-12 Downloading NVSRAM

> **Note:** The title bar now shows Storage Manager 8. Note that the partitions allowed display in the right-hand side of the status bar. The symbols for FlashCopy (disabled) and Remote Volume Mirroring (disabled) appear in the left-hand side of the status bar.

b. As for the firmware download, you are now prompted for the file containing the NVSRAM you want to download (Figure 7-13). The window has the same structure as the firmware download window (Figure 7-7 on page 241) and shows the current configuration and details of the file you selected.

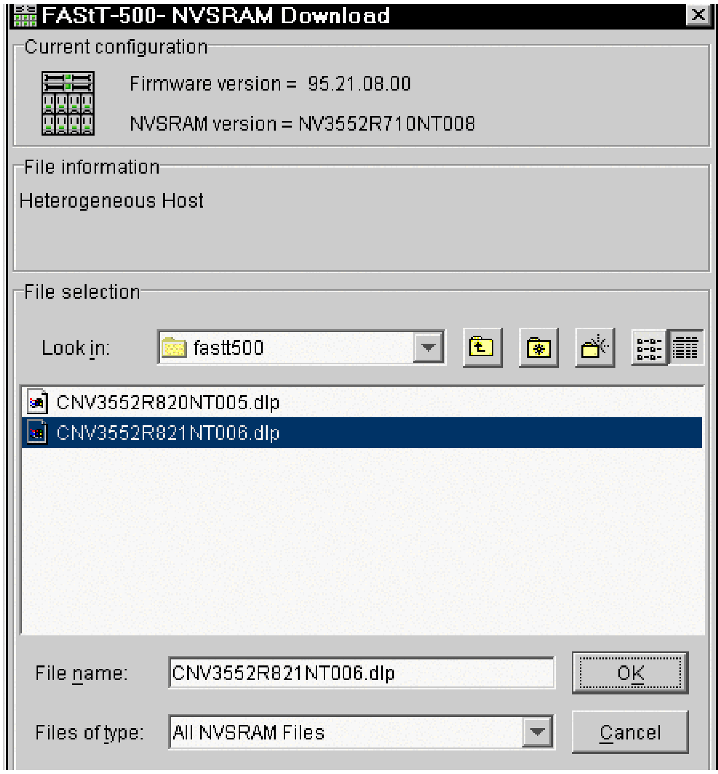

Figure 7-13 Selecting NVSRAM

c. A confirmation window appears. It reminds you to check that the NVSRAM version you are downloading is compatible with the firmware you downloaded (Figure 7-14).

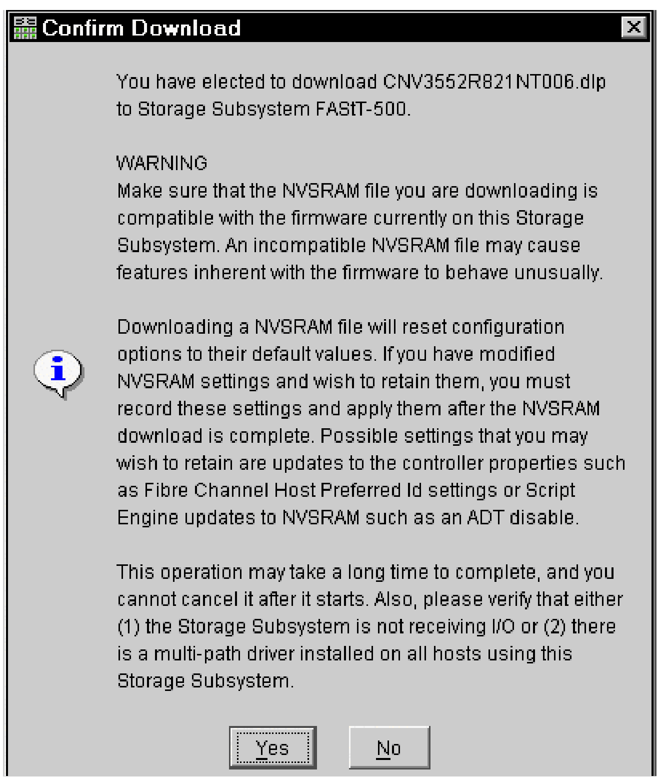

Figure 7-14 NVSRAM Confirm Download

Since the NVSRAM package is much smaller, it doesn’t take as long to download. The status window for downloading NVSRAM appears as shown in Figure 7-15.

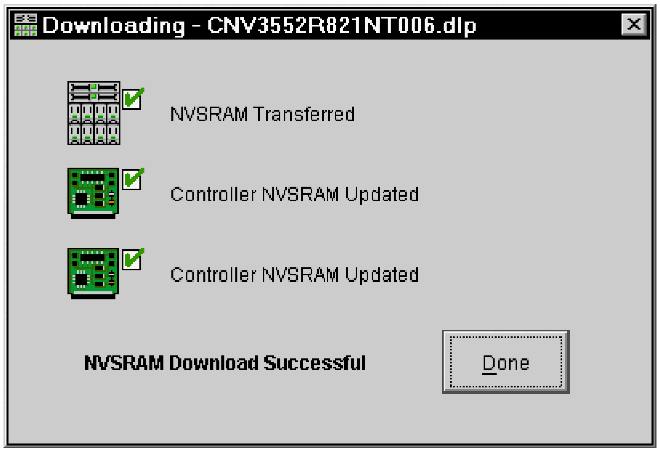

Figure 7-15 NVSRAM download done

4. After the upgrade procedure, close the FAStT Storage Manager and reboot the entire FAStT Storage Server by power cycling the controller unit. Be sure to switch off both power supplies.

   If the FAStT Storage Server is not recognized or unresponsive after the upgrade, initiate a new discovery in the Enterprise Management window. If the FAStT Storage Server is still unresponsive, you must also reboot the host system and initiate a discovery when the system is up again. This is due to the FAStT Storage Agent as explained in 8.4.2, “FAStT Client cannot connect to the FAStT Storage Server” on page 263.

5. After the update is complete, verify the new firmware and NVSRAM by right-clicking each controller in the physical view and selecting **Properties** (Figure 7-16).

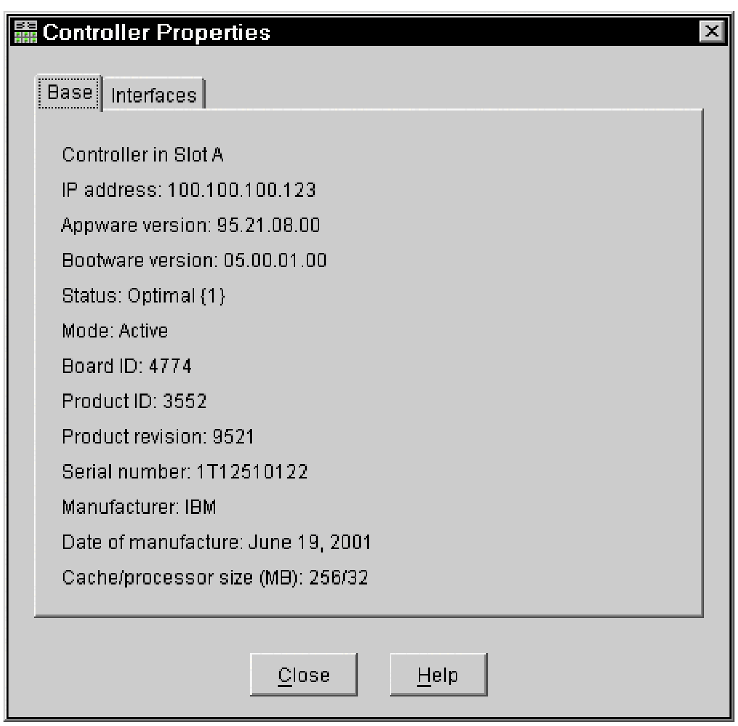

Figure 7-16 Controller A properties after the update

6. Verify that your configuration was successfully maintained. Open the Storage Subsystem Management window. Under the **Logical/Physical View** tab (Figure 7-17), visually check to see if the arrays look OK.

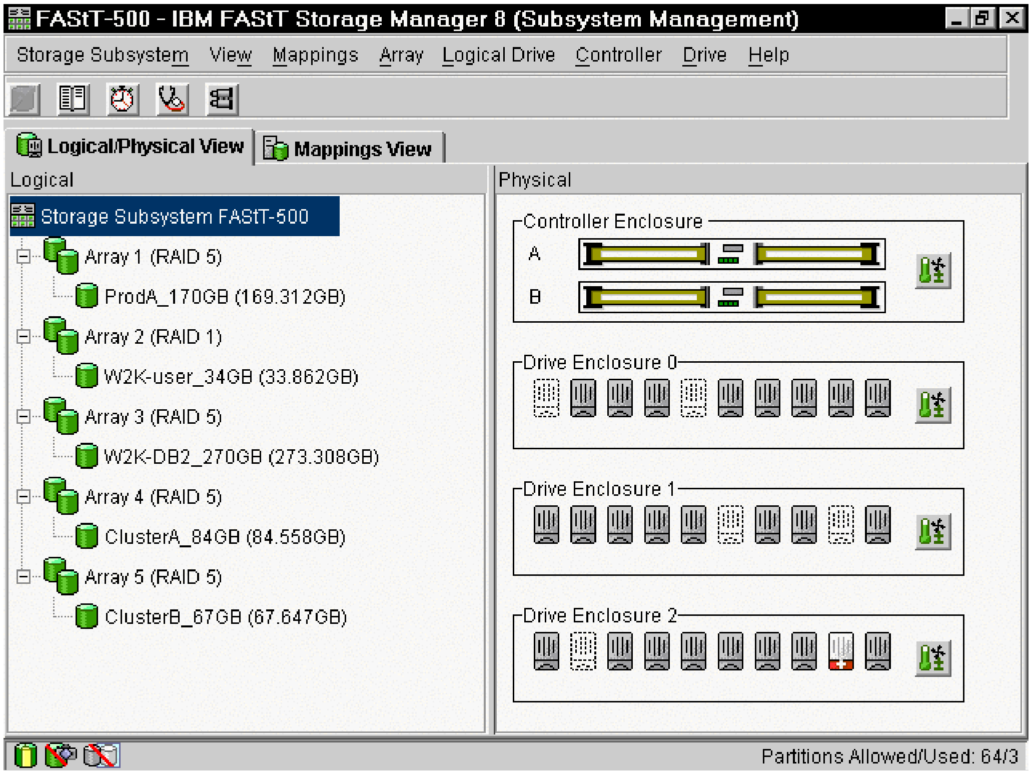

Figure 7-17 Logical/Physical View after the update

7. Click the **Mappings View** tab and see if the LUN mappings stayed the same (Figure 7-18).

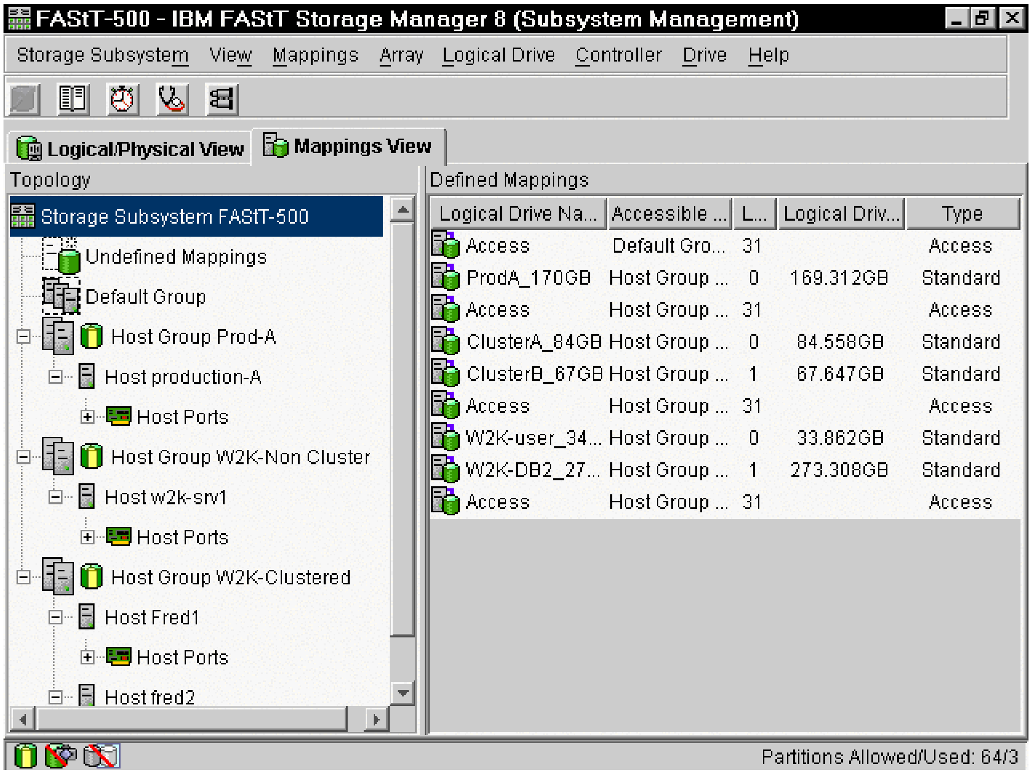

Figure 7-18 Mappings View after the update

8. Compare the new profile with the old one. Click **View-> Profile**. Note that this is now under a different menu. A new interface opens (Figure 7-19).

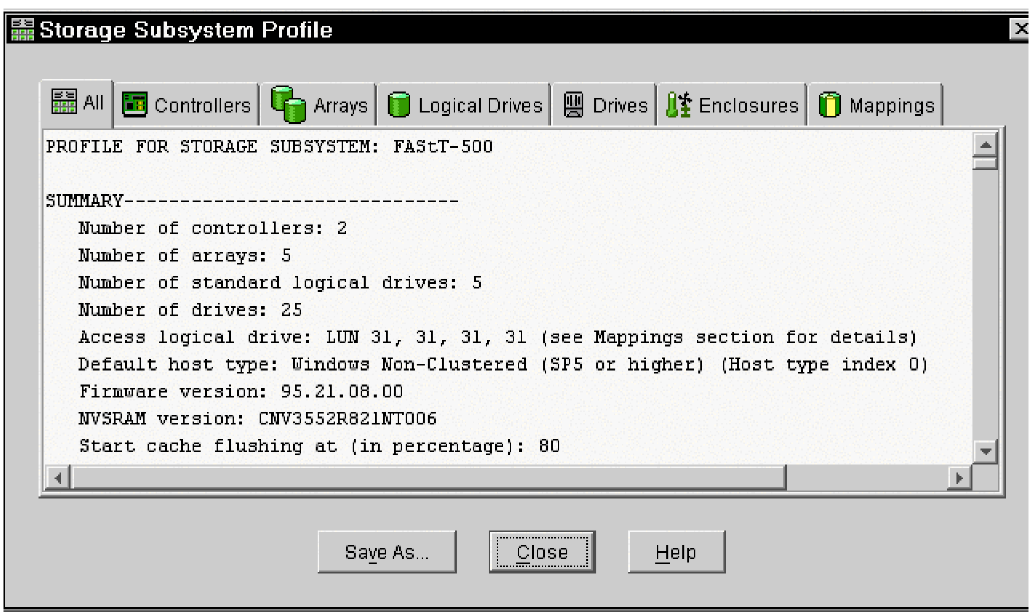

Figure 7-19 Storage Subsystem profile after the update

9. This allows you to view and save the entire profile or subsections of it. If you are satisfied that your update was successful, save your new profile. Be sure not to store the profile on the FAStT Storage Server itself!

Downgrading the FAStT Storage Server that is going from Storage Manager v8 back to Storage Manager v7 or from Storage Manager v8.2 back to Storage Manager v8.0 is not supported by the FAStT Storage Manager. If you try this, you receive a message like the one shown in Figure 7-20.

Figure 7-20 Downgrade not supported

However we found that, with the IBM FAStT Storage Manager Field Tool, as used to update drive and ESM firmware, you can downgrade within the same point release. For example, you can downgrade from Storage Manager v8.2 to Storage Manager v8.0, but you can't downgrade from Storage Manager v8 to Storage Manager v7. If you try this, you receive a message like the example in Figure 7-21.

Figure 7-21 Downgrading with the Field Tool

We couldn't see any reason you would want to downgrade, but it may be useful in a training environment. Do *not* attempt this on production equipment because you may lose functionality.

8

FAStT700 maintenance

This chapter describes the various maintenance functions of FAStT700, such as performance monitoring, error reporting, and diagnostics. It also covers how to use performance monitor data to make decisions regarding the tuning of the storage subsystem.

> **Attention:** Some of the operating systems that are discussed in this section, at the time this redbook was written, are not formally supported. You must check the FAStT supported servers Web site prior to connecting hosts to the FAStT Storage Server or using the advanced functions. You can locate the FAStT700 supported servers Web site at:
>
> `http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm`

8.1 Performance monitoring and tuning

This section describes the Performance Monitor and how the data it provides can be used to tune various parameters. It also looks at what settings can be adjusted and what the expected outcome should be. See 4.6, “Maintenance, tuning, diagnostics, and troubleshooting” on page 114, for step-by-step guide to some of the following topics.

8.1.1 Performance Monitor

FAStT Storage Manager includes a Performance Monitor as shown in Figure 8-1. To use the Performance Monitor, you must start it manually. Its data can help make storage subsystem tuning decisions as explained in the following sections.

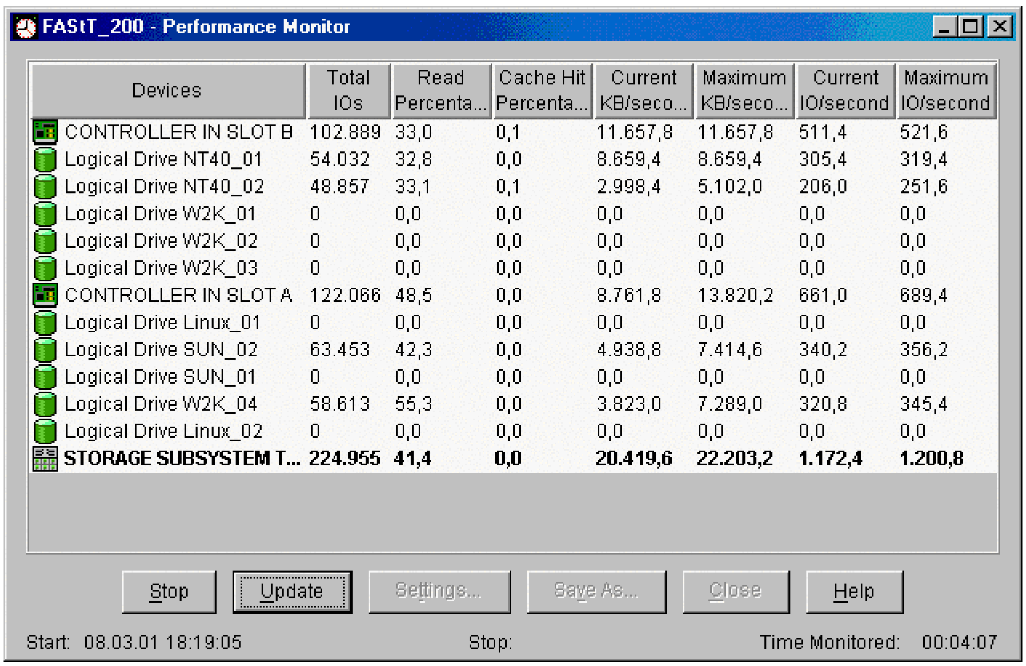

Devices	Total IOs	Read Percenta...	Cache Hit Percenta...	Current KB/seco...	Maximum KB/seco...	Current IO/second	Maximum IO/second
CONTROLLER IN SLOT B	102.889	33,0	0,1	11.657,8	11.657,8	511,4	521,6
Logical Drive NT40_01	54.032	32,8	0,0	8.659,4	8.659,4	305,4	319,4
Logical Drive NT40_02	48.857	33,1	0,1	2.998,4	5.102,0	206,0	251,6
Logical Drive W2K_01	0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive W2K_02	0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive W2K_03	0	0,0	0,0	0,0	0,0	0,0	0,0
CONTROLLER IN SLOT A	122.066	48,5	0,0	8.761,8	13.820,2	661,0	689,4
Logical Drive Linux_01	0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive SUN_02	63.453	42,3	0,0	4.938,8	7.414,6	340,2	356,2
Logical Drive SUN_01	0	0,0	0,0	0,0	0,0	0,0	0,0
Logical Drive W2K_04	58.613	55,3	0,0	3.823,0	7.289,0	320,8	345,4
Logical Drive Linux_02	0	0,0	0,0	0,0	0,0	0,0	0,0
STORAGE SUBSYSTEM T...	**224.955**	**41,4**	**0,0**	**20.419,6**	**22.203,2**	**1.172,4**	**1.200,8**

Figure 8-1 Performance Monitor

Total I/Os

This data is useful for monitoring the I/O activity of a specific controller and a specific logical drive, which can help identify possible high-traffic I/O areas. If I/O rate is slow on a logical drive, try increasing the array size by clicking **Array-> Add Free Capacity** (Drives).

You may notice a disparity in the Total I/Os (workload) of controllers. For example, the workload of one controller is heavy or is increasing over time, while that of the other controller is lighter or more stable. In this case, consider changing the controller ownership of one or more logical drives to the controller with the lighter workload. Use the logical drive Total I/O statistics to determine which logical drives to move.

You may notice the workload across the storage subsystem (Storage Subsystem Totals Total I/O statistic) continues to increase over time while application performance decreases. This may indicate the need to add more storage subsystems to your enterprise so that you can continue to meet application needs at an acceptable performance level.

Read Percentage

Use the Read Percentage for a logical drive to determine the actual application behavior. If there is a low percentage of read activity relative to write activity, consider changing the RAID level of an array from RAID 5 to RAID 1 for faster performance.

Cache Hit Percentage

A higher percentage is desirable for optimal application performance. There is a positive correlation between the cache hit percentage and I/O rates.

The cache hit percentage of all logical drives may be low or trending downward. This may indicate inherent randomness in access patterns. Or, at the storage subsystem or controller level, this can indicate the need to install more controller cache memory if you do not have the maximum amount of memory installed.

If an individual logical drive is experiencing a low cache hit percentage, consider enabling the *cache read ahead* for that logical drive. Cache read ahead can increase the cache hit percentage for a sequential I/O workload.

To determine if your I/O has sequential characteristics, try enabling a conservative cache read-ahead multiplier (four, for example) by clicking **Logical Drive-> Change-> Cache Settings**. Then, examine the logical drive cache hit percentage to see if it has improved. If it has, indicating that your I/O has a sequential pattern, enable a more aggressive cache read-ahead multiplier (8 for example). Continue to customize logical drive cache read-ahead to arrive at the optimal multiplier (in the case of a random I/O pattern, the optimal multiplier is 0).

Current KB/sec and Maximum KB/sec

The transfer rates of the controller are determined by the application I/O size and the I/O rate. Generally, small application I/O requests result in a lower transfer rate but provide a faster I/O rate and shorter response time. With larger application I/O requests, higher throughput rates are possible. Understanding your typical application I/O patterns can help you determine the maximum I/O transfer rates for a given storage subsystem.

Consider a storage subsystem, equipped with Fibre Channel controllers, that supports a maximum transfer rate of 100 MB/second (100,000 KB per second). Your storage subsystem typically achieves an average transfer rate of 20,000 KB/second. (The typical I/O size for your applications is 4K, with 5,000 I/Os transferred per second for an average rate of 20,000 KB/second.) In this case, I/O size is small. Because there is system overhead associated with each I/O, the transfer rates will not approach 100,000 KB per second. However, if your typical I/O size is large, a transfer rate within a range of 80,000 to 90,000 KB per second might be achieved.

Wide Ultra SCSI supports sustained data transfer rates of up to 40 MB/second with large host I/O sizes. The maximum transfer rates with relatively small I/O sizes, 4096 bytes, is approximately 17 MB/second.

Current I/O per second and Maximum I/O per second

Factors that affect I/Os per second include access pattern (random or sequential), I/O size, RAID level, segment size, and the number of drives in the arrays or storage subsystem. The higher the cache hit rate is, the higher I/O rates will be.

Performance improvements caused by changing the segment size can be seen in the I/Os per second statistics for a logical drive. Experiment to determine the optimal segment size, or use the file system or database block size.

Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual logical drive, consider the current and maximum I/Os per second. You should expect to see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, we recommend that you enable write caching to maximize the I/O rate and shorten application response time.

8.1.2 Cache parameters

Efficient use of the RAID controller cache is essential for good performance of the FAStT Storage Server. The FAStT Storage Manager utility enables you to configure various cache settings:

- Start and stop cache flushing levels
- Cache block size
- Enabling or disabling write cache mirroring
- Cache read-ahead multiplier
- Write-back and write-through mode

These settings have a large impact on performance of the FAStT Storage Server and on the availability of data. Be aware that performance and availability often conflict with each other. If you want to achieve maximum performance, in most cases, you must sacrifice system availability and vice versa. Most customers are usually more motivated in the highest availability they can get, rather than the maximum performance.

Starting and stopping cache flushing levels

These two settings affect the way the cache controller handles unwritten cache entries. They are obviously only effective when you configure the write-back cache policy. Writing the unwritten cache entries to the disk drives is called *flushing*. You can configure the start and stop flushing level values. They are expressed as percentages of the entire cache capacity. When the number of unwritten cache entries reaches the start flushing value, the controller begins to flush the cache (write the entries to the disk drives). The flushing stops when the number of unwritten entries drops below the stop flush value. The controller always flushes the oldest cache entries first. Unwritten cache entries older than 20 seconds are flushed automatically.

A typical start flushing level is 80%. Very often the stop flushing level is set to 80%, too. This means the cache controller will not allow more than 80% of the entire cache size for write-back cache, but it will also try to keep as much of it as possible for this purpose. If you use such settings, you can expect a high amount of unwritten entries in the cache. This is good for writing performance, but be aware that it offers less data protection.

If you are concerned for data protection, you might want to use lower start and stop values. With these two parameters, you can actually tune your cache for either reading or writing performance.

Performance tests have shown that it is a good idea to use similar values for start and stop flushing levels. If the stop level value is significantly lower than the start value, this causes a high amount of disk traffic when flushing the cache. This interferes with normal disk access and is usually not even needed. If the values are similar, then the controller only flushes the amount needed to stay within limits.

Cache block size

This is the size of the cache memory allocation unit and can be either 4K or 16K. By selecting the proper value for your particular situation, you can significantly improve the caching efficiency and performance. For example, if applications mostly access the data in small blocks up to 8K, but you use 16K for cache block size. Then each cache entry block is only partially populated. You will always occupy 16K in cache to store 8K (or less) of data. This means only up to 50% of cache capacity is effectively used to store the data. You can obviously expect lower performance. For random workloads and small data transfer sizes, 4K is better.

On the other hand, if the workload is sequential and you use large segment size, it is a good idea to use larger cache block size 16K. A larger block size means a lower number of cache blocks and reduces cache overhead delays. In addition, a larger cache block size requires fewer cache data transfers to handle the same amount of data.

Write cache mirroring

FAStT write cache mirroring provides the integrity of cached data if a RAID controller fails. This is excellent from a high availability perspective, but it decreases performance. The data is mirrored between controllers across the drive-side FC loop. This competes with normal data transfers on the loop.

Read-ahead multiplier

This parameter affects the reading performance and an incorrect setting can have a large negative impact. It controls how many additional sequential data blocks will be stored into cache after a read request.

Obviously, if the workload is random, this value should be 0. Otherwise each read request will unnecessarily prefetch additional data blocks. Since these data blocks will rarely be needed, the performance is going to be impacted.

For sequential workloads, a good value would be between 1 and 4, depending on the particular environment. When using such setting, a read request causes prefetching of several sequential data blocks into the cache; this speeds up consequent disk access. This leads to a fewer number of I/O transfers required to handle the same amount of data, which is good for performance in sequential environment.

Write-back and write-through caching

If you configure write-through caching, the cache only increases performance of reading operations. Writing operations do not use cache at all. The data is always going to be written directly to the disk drives.

Write-back caching can also increase performance of write operations. The data is not written straight to the disk drives; it is only written to the cache. From the application perspective, this is much faster than waiting for the disk write operation to complete. Therefore, you can expect a significant gain in application writing performance. It is the responsibility of the cache controller to eventually flush the unwritten cache entries to the disk drives.

Write-back mode appears to be faster than write-through mode, since it increases performance of both reads and writes. But is not always true because it depends on the disk access pattern and workload.

A lightly loaded disk subsystem usually works faster in write-back mode, but when the workload is high, the write cache may become inefficient. As soon as the data is written to the cache, it has to be flushed to the disks in order to make room for new data arriving into cache. The controller would actually perform faster if the data went directly to the disks. In this case, writing the data to the cache is an unnecessary step that actually decreases throughput.

You can use the FAStT Storage Manager Client to set the appropriate caching mode for each logical drive independently.

8.1.3 Controller ownership

The Total I/O data provided by the Performance Monitor is useful for monitoring the I/O activity of a specific controller and a specific logical drive, which can help identify possible high-traffic I/O areas.

Identify actual I/O patterns to the individual logical drives and compare those with the expectations based on the application. If a particular controller has considerably more I/O activity, consider moving logical drives to the other controller in the storage subsystem.

You may notice a disparity in the Total I/Os (workload) of controllers. For example, the workload of one controller is heavy or is increasing over time, while that of the other controller is lighter or more stable. In this case, consider changing the controller ownership of one or more logical drives to the controller with the lighter workload. Use the logical drive Total I/O statistics to determine which logical drives to move.

8.1.4 Segment size

The RAID controllers read and write data to and from the physical disk drives in fixed-sized segments, before writing data on the next physical drive. A *segment* is the amount of data, in kilobytes, that the controller writes on a single drive in a logical drive before writing data on the next drive. Data blocks store 512 bytes of data and are the smallest units of storage. The size of a segment determines how many data blocks it contains. For example, an 8 KB segment holds 16 data blocks, and a 64 KB segment holds 128 data blocks. The segment size is specified per each logical drive. You can use the following values for the segment size:

- 8K
- 16K
- 32K
- 64K
- 128K
- 256K

You should choose the segment size carefully, because it can greatly impact performance. A large segment size (relative to the average request size) increases the request rate by allowing multiple disk drives to respond to multiple requests. If one disk drive contains all of the data for one request, the other disk drives in the storage set are available to handle other requests. In principle, separate I/O requests can be handled in parallel, increasing the request rate.

A small segment size (relative to the average request size) increases the data transfer rate by allowing multiple disk drives to participate in one I/O request. Sequential write and read requests should use a small segment size relative to the I/O size to increase performance.

To determine the optimal segment size, you often have to go through performance measurement and tuning tasks. This may mean the storage server has to run in production for some time to give you an idea of data access patterns. Once the best value for the segment size is determined, you can use the FAStT Storage Manager to change it to this value. This process is again non-disruptive. No downtime is needed.

8.1.5 RAID levels

Performance varies based on the RAID level that is set. Use the Performance Monitor to obtain logical drive application read and write percentages. Applications with a high read percentage do very well using RAID level 5 logical drives because of the outstanding read performance of the RAID 5 configuration.

However, applications with a low read percentage (write-intensive) do not perform as well on RAID 5 logical drives because of the way a controller writes data and redundancy data to the drives in a RAID 5 array. If there is a low percentage of read activity relative to write activity, consider changing the RAID level of an array for faster performance.

Based on the respective level, RAID offers the following performance results:

- RAID 0 offers high performance, but does not provide any data redundancy.
- RAID 1 offers high performance for write-intensive applications.
- RAID 3 is good for large data transfers in applications, such as multimedia or medical imaging, that write and read large sequential chunks of data.
- RAID 5 is good for multi-user environments, such as database or file system storage, where the typical I/O size is small and there is a high proportion of read activity.

8.1.6 Logical drive modification priority

The modification priority defines how much processing time is allocated for logical drive modification operations relative to system performance. You can increase the logical drive modification priority, although this may affect system performance.

Operations affected by the modification priority include:

- Copyback
- Reconstruction
- Initialization
- Changing segment size
- Defragmentation of an array
- Adding free capacity to an array
- Dynamic logical drive expansion
- Changing the RAID level of an array

The following modification priority rates are available:

- Lowest
- Low
- Medium
- High
- Highest

Note: The lowest priority rate favors system performance, but the modification operation takes longer. The highest priority rate favors the modification operation, but system performance may be compromised.

The progress bar at the bottom of the Logical Drive Properties dialog displays the progress of a modification operation.

8.1.7 Remote Volume Mirroring

When a storage subsystem logical drive is a primary logical drive and a full synchronization is necessary, the controller owner performs the full synchronization in the background while processing local I/O writes to the Primary Logical Drive and associated remote writes to the secondary logical drive. Because the full synchronization diverts controller processing resources from I/O activity, it can impact performance on the host application. The synchronization priority defines how much processing time is allocated for synchronization activities relative to system performance.

The following guidelines may help you determine how long a synchronization priority can take and how much various synchronization priorities can affect system performance.

The following synchronization priority rates are available:

- Lowest
- Low
- Medium
- High
- Highest

Note: The lowest priority rate favors system performance, but the full synchronization takes longer. The highest priority rate favors full synchronization, but system performance may be compromised. The following guidelines roughly approximate the differences between the five priorities. Logical drive size and Host I/O rate loads affect the synchronization time comparisons.

A full synchronization at the *lowest* synchronization priority rate takes approximately eight times as long as a full synchronization at the highest synchronization priority rate.

A full synchronization at the *low* synchronization priority rate takes approximately six times as long as a full synchronization at the highest synchronization priority rate.

A full synchronization at the *medium* synchronization priority rate takes approximately three and a half times as long as a full synchronization at the highest synchronization priority rate.

A full synchronization at the *high* synchronization priority rate takes approximately twice as long as a full synchronization at the highest synchronization priority rate.

The synchronization progress bar at the bottom of the Mirroring tab of the Logical Drive Properties dialog displays the progress of a full synchronization.

8.2 Error reporting

The FAStT700 and FAStT Storage Manager have ways to log faults, notify users of faults, and guide them through the necessary recovery steps. The following sections explain how this works.

8.2.1 Problem notification

Indicators of a storage subsystem problem include:

- A Needs Attention storage subsystem icon displayed in the Overall Health Status area, Device Tree View, or Device Table of the Enterprise Management window or the Subsystem Management window Logical View
- The Recovery Guru Optimal toolbar button in the Subsystem Management window changes to a flashing Recovery Guru Needs Attention toolbar button.
- Non-optimal component icons displayed in the Subsystem Management window Logical and Physical View
- The receipt of critical Simple Network Management Protocol (SNMP) or e-mail error messages
- The display of hardware fault lights

Use the Recovery Guru to help troubleshoot storage subsystem problems. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

8.2.2 Event Viewer

The event log is a detailed record of events that occur on the storage subsystem. You can use the event log as a supplementary diagnostic tool to Recovery Guru for tracing storage subsystem events. However, always refer to the Recovery Guru first when recovering from storage subsystem component failures.

Examples of events logged to the event log include:

- Configuration events, such as logical drive creation or the assignment of a drive as a hot spare
- Failures of storage subsystem components

The event log is stored in reserved areas on the disks in the storage subsystem.

The Storage Management software provides a way to view and sort the data in the event log, refresh the display to retrieve any new events, display detailed information about a selected event, and save selected event log data to a file.

> **Important:** It can take up to several minutes for an event to be logged and become visible in the Event Log Viewer.

The event log contains two levels of events:

- Critical events
- Informational events

Critical events are sent to any Network Management Station (through SNMP traps) or e-mail recipient you configured by clicking **Edit-> Alert** in the Enterprise Management window. Figure 8-2 shows the Event Viewer listing critical events.

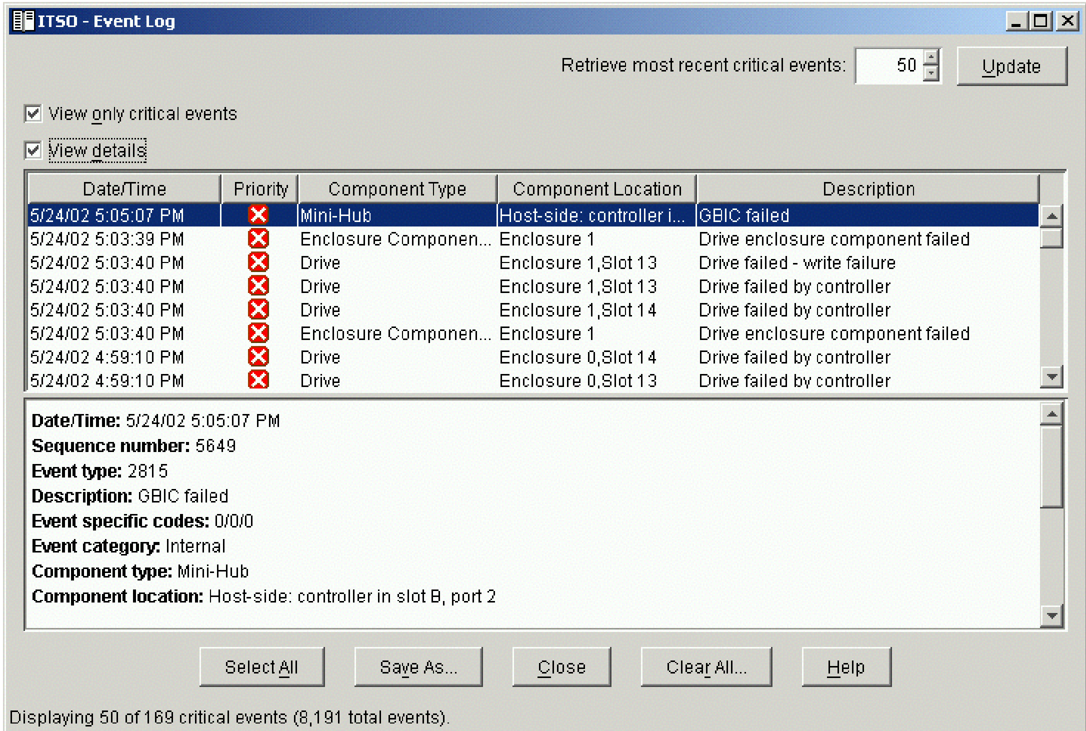

Figure 8-2 Event Viewer showing critical events

8.2.3 Recovery Guru

The Recovery Guru is a component of the Subsystem Management window that diagnoses storage subsystem problems and recommends recovery procedures to fix the problems.

To display Recovery Guru, select the Recovery Guru toolbar button in the Subsystem Management window. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

After you fix a failed component, note the following points:

- The Components button in the controller enclosure returns to Optimal (Physical View).
- The storage subsystem icon in the Logical View returns to Optimal.
- The storage subsystem icon in the Enterprise Management window changes from Needs Attention status to Optimal status.

Recovering from storage subsystem problems

To recover from storage subsystem problems, follow these steps:

1. From the Subsystem Management window, perform one of the following tasks:
   – Click the **Recovery Guru** toolbar button.
   – Click **Storage Subsystem-> Recovery Guru**.
   – Right-click the storage subsystem in the Logical View to select it. Then, select **Recovery Guru**.

   The Recovery Guru window opens. It is divided into three views, Summary, Detail, and Recovery Procedures, which are described in Recovery Guru Views in this help topic.

2. Select the first problem shown in the Summary View and follow the recovery procedure to correct it. Repeat this for each listed problem.

   When all problems are corrected, the storage subsystem icon eventually transitions from Needs Attention to Optimal. There are certain problems where a Fixing icon appears while an operation, such as reconstruction, is in progress.

3. Select **Recheck** to verify the success of the completed recovery procedure.

Figure 8-3 shows the Recovery Guru.

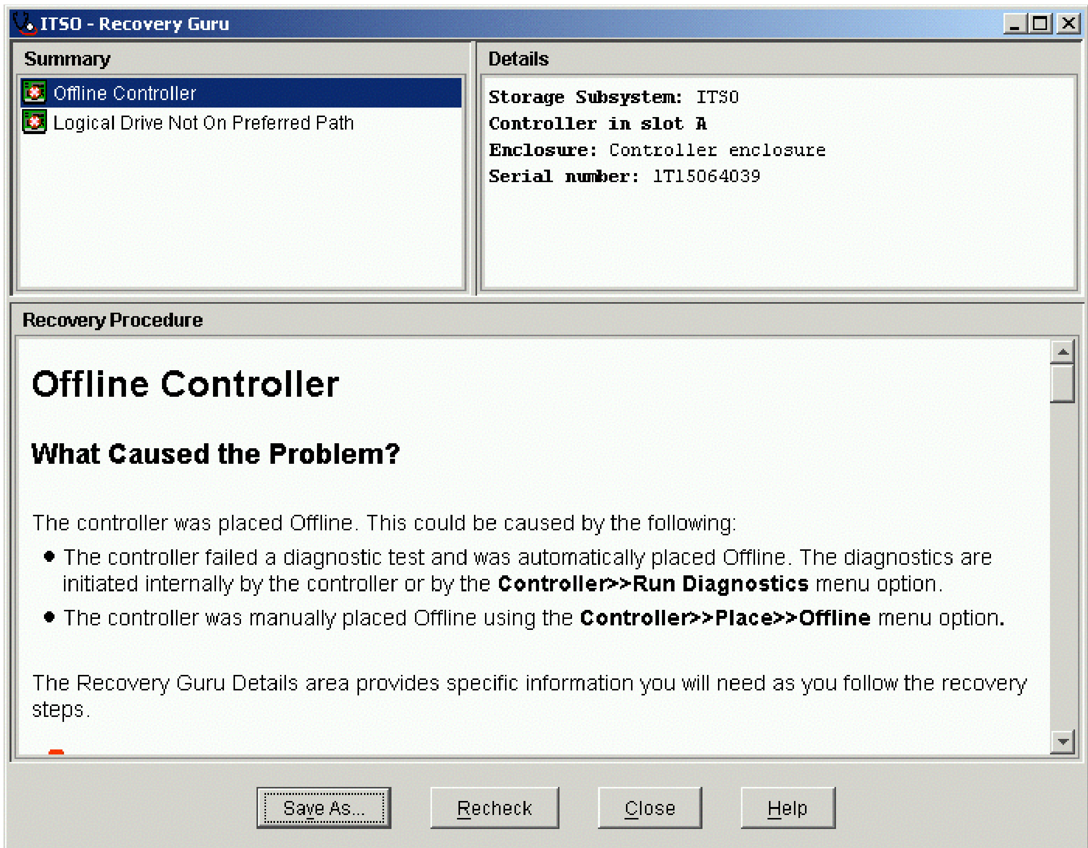

Figure 8-3 Recovery Guru

8.2.4 Missing logical drives

A missing logical drive is a placeholder node displayed in the Logical View. It indicates that the storage subsystem has detected inaccessible drives associated with a logical drive. Typically, this is the result of removing drives associated with an array, or when a loss of power to one or more drive enclosures has occurred.

Missing logical drives are only displayed in the Logical View if they are standard logical drives or repository logical drives. In addition, one of the following conditions must exist:

- The logical drive has an existing *logical drive-to-LUN mapping*, and drives associated with the logical drive are no longer accessible.
- The logical drive is participating in a *Remote Volume Mirror* as either a primary logical drive or secondary logical drive, and drives associated with the logical drive are no longer accessible.

- The logical drive is a *mirror repository logical drive*, and drives associated with the logical drive are no longer accessible. The Recovery Guru has a special recovery procedure for this case. Two mirror repository logical drives are created together on the same array when the Remote Volume Mirroring premium feature is activated and one is used for each controller in the storage subsystem. If drives associated with the array are no longer accessible, then both mirror repository logical drives are missing, and all Remote Volume Mirrors are in an Unsynchronized state.
- The logical drive is a *base logical drive* with associated FlashCopy logical drives, and drives associated with the logical drive are no longer accessible.
- The logical drive is a *FlashCopy repository logical drive*, and drives associated with the logical drive are no longer accessible.

If missing logical drives are detected by the storage subsystem, a Missing Logical Drives group is created in the Logical View of the Subsystem Management window. Each missing logical drive is shown and identified by its world-wide name and logical drive type. Missing logical drives are identified as being either a standard logical drive, base logical drive, FlashCopy repository logical drive, primary logical drive, secondary logical drive, or mirror repository logical drive.

Note: Missing logical drives, in most cases, are recoverable. Do not delete missing logical drives without confirming that the logical drives are no longer required, because they will be permanently removed from the configuration.

If missing logical drives are detected because drives are accidentally removed, or are detected as missing due to a loss of power to the drive enclosures, recovery of these logical drives is possible by:

- Re-inserting the drives back into the drive enclosure.
- Ensuring that the drive enclosure's power supplies are properly connected to an operating power source and have an Optimal status

8.3 Diagnostics

This section explains the different diagnostic aids that are available to help you with problem resolution.

8.3.1 Controller runtime diagnostics

The Controller Diagnostics option allows a user to verify that a controller is functioning properly, through the use of various internal tests. The Diagnostics use a combination of three different tests:

- Read Test
- Write Test
- Data Loopback Test

You should run all three tests at initial installation and anytime there are changes to the storage subsystem or components that are connected to the storage subsystem (such as hubs, switches, and host adapters).

Important: During the diagnostics, the controller on which the tests are run is *not* available for I/O.

Read Test

The Read Test initiates a read command as it would be sent over an I/O data path. It compares data with a known, specific data pattern, checking for data integrity and redundancy errors. If the read command is unsuccessful or the data compared is not correct, the controller is considered to be in error and is failed.

Write Test

A Write Test initiates a write command as it would be sent over an I/O data path (to the Diagnostics region on a specified drive). This Diagnostics region is then read and compared to a specific data pattern. If the write fails or the data compared is not correct, the controller is considered to be in error and is failed and placed offline. (Use the Recovery Guru to replace the controller.)

Data Loopback Test

The Data Loopback Test passes data through each controller's drive-side channel, mini-hub, out onto the loop, and then back again. Enough data is transferred to determine error conditions on the channel. If the test fails on any channel, then this status is saved so that it can be returned if all other tests pass.

Notes:

- All test results are displayed in the Diagnostics dialog status area.
- Events are written to the event log when Diagnostics is started and has completed testing. These events help you to evaluate whether diagnostics testing was successful or failed, as well as the reason for the failure. To view the event log, click **View-> Event Log** in the Subsystem Management window.

8.3.2 Read Link Status diagnostics

A *Fibre Channel loop* is an interconnection topology used to connect storage subsystem components and devices. The Storage Management software uses the connection between the host machine and each controller in the storage subsystem to communicate with each component and device on the loop.

During communication between devices, Read Link Status (RLS) error counts are detected within the traffic flow of the loop. Error count information is accumulated over a period of time for every component and device including:

- Drives
- ESMs
- Fibre Channel ports

Error counts are calculated from a baseline, which describes the error count values for each type of device in the Fibre Channel loop. Calculation occurs from the time when the baseline was established to the time at which the error count information is requested. The baseline is automatically set by the controller. However, a new baseline may be set manually through the Read Link Status Diagnostics dialog.

Read Link Status error counts refer to link errors that have been detected in the traffic flow of a Fibre Channel loop. The errors detected are represented as a count (32 bit field) of error occurrences accumulated over time and help to provide a coarse measure of the integrity of the components and devices on the loop.

By analyzing the error counts retrieved, it is possible to determine the components or devices within the Fibre Channel loop that may be experiencing problems communicating with the other devices on the loop. A high error count for a particular component or device indicates that it may be experiencing problems and should be given immediate attention.

Analyzing Read Link Status results

Analysis of the RLS error count data is based on the principle that the device immediately “downstream” of the problematic component should see the largest number of Invalid Transmission Word (ITW) error counts.

> **Important:** Because the current error counting standard is vague about when the ITW count is calculated, different vendor devices calculate errors at different rates. Take this into account in your analysis of the data.

The analysis process involves obtaining an ITW error count for every component and device on the loop, viewing the data in loop order, and then identifying any large jumps in the ITW error counts. Figure 8-4 shows the typical error count information displayed in the Read Link Status Diagnostics window.

Devices	Baseline Time	Elapsed Time	ITW	LF	LOS	LOSG	PSP	ICRC
Drive Channel 1								
Controller in slot A	3/19/01 8:32:56 AM	02:19:24	31	45	26	42	31	38
Drive [10,1]	3/19/01 8:32:56 AM	02:19:24	44	33	36	44	32	42
Drive [10,2]	3/19/01 8:32:56 AM	02:19:24	39	33	19	39	41	52
Drive [10,3]	3/19/01 8:32:56 AM	02:19:24	25	33	46	24	30	49
Drive [10,4]	3/19/01 8:32:56 AM	02:19:24	27	38	36	29	27	24
Drive [10,5]	3/19/01 8:32:56 AM	02:19:24	47	32	33	33	29	40
Drive [10,6]	3/19/01 8:32:56 AM	02:19:24	29	45	24	19	33	43
Drive [10,7]	3/19/01 8:32:56 AM	02:19:24	41	40	34	46	27	22
Drive [10,8]	3/19/01 8:32:56 AM	02:19:24	27	40	39	39	39	28
Drive [10,9]	3/19/01 8:32:56 AM	02:19:24	28	21	30	40	27	42
Drive [10,10]	3/19/01 8:32:56 AM	02:19:24	32	38	22	20	39	44
Drive Enclosure 10 (...	3/19/01 8:32:56 AM	02:19:24	41	33	41	40	42	56
Drive Enclosure 3 (...	3/19/01 8:32:56 AM	02:19:24	33	29	47	38	47	34
Drive Enclosure 2 (...	3/19/01 8:32:56 AM	02:19:24	47	39	30	44	32	39
Drive Enclosure 1 (...	3/19/01 8:32:56 AM	02:19:24	42	30	21	56	39	31

Figure 8-4 Read Link Status

Each of the columns in the window are explained in the following list:

- **Devices**: A list of all the devices on the Fibre Channel loop. The devices are displayed in channel order. Within each channel, they are sorted according to the devices position within the loop.
- **Baseline Time**: The date and time of when the baseline was last set.
- **Elapsed Time**: The elapsed time between when the Baseline Time was set and when the read link status data was gathered using the Run option.
- **Link Failure (LF)**: The total number of link failure errors detected on the Fibre Channel loop from the baseline time to the current date and time. When detected, link failures indicate that there has been a failure within the media module laser operation. Link failures may also be caused by a link fault signal, a loss of signal or a loss of synchronization.
- **Loss of Synchronization (LOS)**: The total number of loss of synchronization errors detected on the Fibre Channel loop from the baseline time to the current date and time. This indicates that the receiver cannot acquire symbol lock with the incoming data stream, due to a degraded input signal. If this condition persists, the number of loss of signal errors increases.
- **Loss of Signal (LOSG)**: The total number of loss of signal errors detected on the Fibre Channel loop from the baseline date to the current date and time. This indicates a loss of signal from the transmitting node or physical component within the Fibre Channel loop.

Physical components where a loss of signal typically occurs include the gigabit interface connectors and the Fibre Channel fibre optic cable.

- **Primitive Sequence Protocol (PSP)**: The total number of primitive sequence protocol errors detected on the Fibre Channel loop from the baseline date to the current date and time. This refers to the number of N_Port protocol errors detected and Link Reset Response (LRR) primitive sequences received while the link is up. A Link Reset Response (LRR) is issued by another N_Port in response to a link reset.
- **Invalid Cyclic Redundancy Check (ICRC)**: The total number of invalid cyclic redundancy check errors detected on the Fibre Channel loop, from the baseline date to the current date and time. This indicates that a frame has been received with an invalid cyclic redundancy check value. A cyclic redundancy check is performed by reading the data, calculating the cyclic redundancy check character, and then comparing its value to the cyclic check character already present in the data. If they are equal, the new data is presumed to be the same as the old data.

8.4 Common problems

There are some common problems that sometimes occur during the installation and configuration of the FAStT Storage Server. This section reviews those problems.

8.4.1 Out of compliance error message in Recovery Guru

There are cases where the Recovery Guru reports the Out of Compliance error. This can be caused by an improper firmware and NVSRAM upgrade. The procedure to recover from this error condition is to disable the premium feature storage partitioning and enable it with a new key.

The key file must be generated for your FAStT Storage Server since it is unique for each storage subsystem. Contact the IBM Help Center to have them generate the key for you. They need the Feature Enable Identifier of your FAStT Storage Server to do so. See 4.6.10, “Handling premium features” on page 127, for the exact procedure.

You can use SMdevices to obtain a list of all logical devices including their properties that can be accessed through this server. If all logical drives are listed but not the access logical drive, you may have a problem with storage partitioning.

8.4.2 FAStT Client cannot connect to the FAStT Storage Server

If the FAStT Client cannot connect to the FAStT Storage Server, there are different possible causes depending on whether you are using in-band or out-of-band management. Each is explained in the following sections.

In-band

If you are attempting to manage your Storage Server by the in-band method, it may be that you did not install all the Storage Manager components correctly as described in 4.1, “Driver and host software installation” on page 66. Or it may be that the FAStT host agent is not started or cannot see the FAStT Storage Server.

When the FAStT Host Agent is started by the operating system, it scans the fibre paths to find the attached storage subsystem. When the service (or daemon) cannot find a storage subsystem, it stops itself. On Windows systems, you receive a message that the host agent stopped with error code 100.

The reason is that the host system has no connection to the access logical drive. There can be several reasons for this error, which we discuss in the following sections.

No physical connection

No connection between the host bus adapter and the FAStT Storage Server can be established because the physical connection is broken or not connected properly. All cables must be seated properly in the GBIC ports and the link status must be online. Check the link status LEDs on switches, hubs, and the storage subsystem. Also, the cable itself may be broken. In this case, swap cables and watch if the error moves.

Within the operating system, check whether the host bus adapter itself sees any devices of the FAStT Storage Server. For the host agent, the access logical drive must be visible.

Zoning in the fabric

Even if the connections are physically online, it may be still impossible for the operating system to establish a link if there are zones defined in the fabric.

You need one connection between the host bus adapter and the controller in the FAStT Storage Server. There are different combinations possible depending on the number of host bus adapters, controllers, and whether the operating systems use redundant paths.

In general, each host bus adapter should see only one controller of the FAStT Storage Server. If you are attaching one server with two host bus adapters to one switch and connect the storage subsystem with two paths, you have to define two zones.

If the host is attached with one host bus adapter and the FAStT Storage Server is connected with two paths to the switch, all three ports should be in the same zone to allow the host agent to access both controllers. In this case, you still have multiple paths to the same logical drives in the storage subsystem.

For a host with redundant paths, there must be only one path per host bus adapter to one of the controller. This gives you two completely separate paths to the storage subsystem. If the host does not support redundant paths, you should ensure that the host systems only sees one controller of the storage subsystem. If the host sees both controllers, it depends on the operating system to access only one path.

Mapping does not allow to attach to the access logical drive

Another reason for this error is a mapping of the access logical drive to another group than the default host group. If there is a mapping defined for the access logical drive, you can only manage the storage subsystem through the fibre connection if you use one of the hosts inside this host group.

If you mapped the access logical drive to a host group containing only dummy entries for the host ports, or to a host group where the host operating system, you can only manage the FAStT Storage Server directly through the network ports.

Out-of-band management

You may be trying to use out-of-band management to manage your Storage Server over Ethernet and the FAStT Client program cannot automatically connect, because your management station is on a different subnet. In this case, you may have to add your Storage Server to the Enterprise Management window manually by clicking **Edit-> Add Device**.

If you still cannot connect, your storage server may not have a default gateway defined. Or if you have a FAStT500 that has been updated to Storage Manager 8.2, you may not be aware that the storage server now defaults to fixed IP addresses and not DHCP. In either case, you may have to edit the network settings through the serial ports as described in 4.6.11, "Network setup of the controllers" on page 129. This same information is published as RETAIN Tip H171389 on the support Web site. Although it was published for FAStT500 with SM7, the procedure for changing the network settings is still the same. See:

http://www.pc.ibm.com/qtechinfo/MIGR-4MMP53.html

8.4.3 Logical drives not on the preferred path

A very common but not severe problem is the message that not all logical drives are managed by their preferred controller. This happens for example if a host with redundant paths has a failure in one of the paths. The Auto Logical Drive (ADT) feature (if enabled) moves the affected logical drives to the intact path. If the failed path returns online, the logical drive moves back automatically as I/O requests are received.

8.4.4 Storage subsystem is unresponsive

This can happen in the Enterprise Management window if a defined subsystem has its name or IP address changed. See Figure 8-5. You can fix this by removing the subsystem from the Enterprise Management window and then adding it in again.

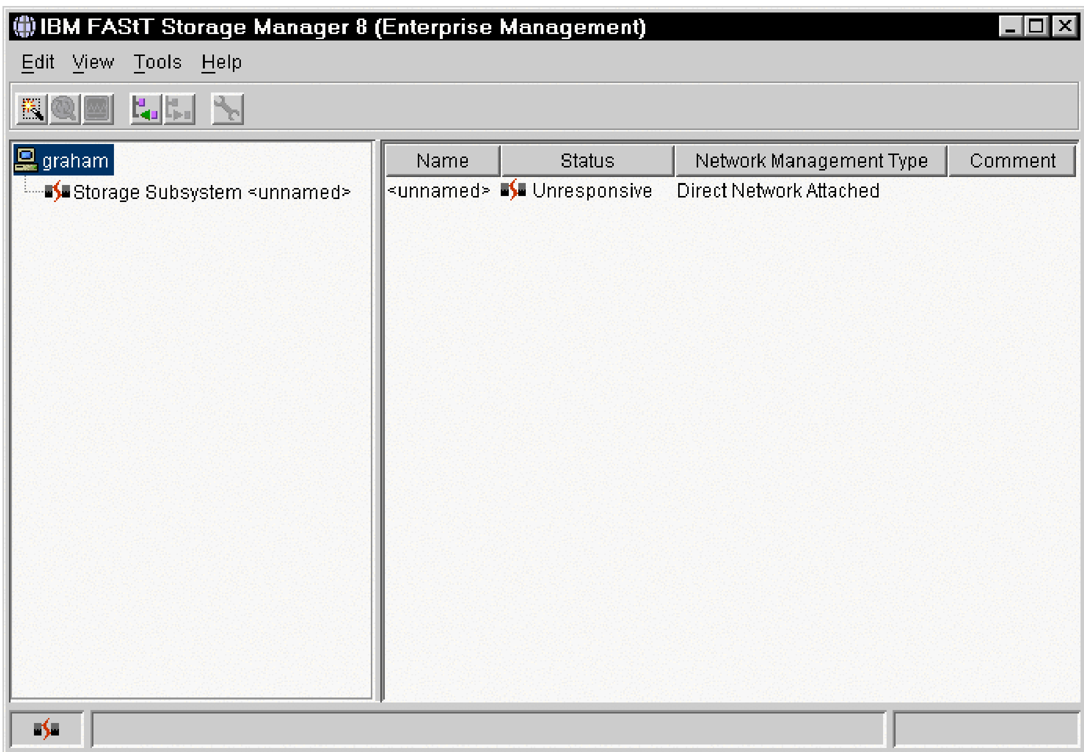

Figure 8-5 Storage subsystem unresponsive

8.4.5 Obtaining the world-wide name of the host bus adapter

The world-wide name (WWN) of the Fibre Channel host bus adapter is a unique identifier of the adapter, similar to the MAC address of a network interface card. Sometimes it is necessary or useful to have this WWN available. The WWN is needed as soon as you implement the storage partitioning feature of the FAStT Storage Server. It is also used in the Fibre Channel switch and hub for zoning the fabric. On a switch or hub level, it is also useful for problem determination because most switches and hubs allow monitoring of their ports, and the attached hosts are only identified by the adapter in use.

In general, on an Intel-based server, you can figure out the WWN of the adapter by entering the BIOS of the adapter while the server is booting up. For the QLogic QLA2x00 adapter, follow these steps:

1. Press Alt+Q.
2. Enter the BIOS.
3. Choose the host adapter.
4. In the option windows, click **Configuration Settings-> Host Adapter Settings**.

The WWN is the value of the entry Adapter Node Name.

Windows, NetWare, Linux

In a Windows NT4, Windows 2000, NetWare, or Linux environment, you can use the FAStT Management Suite Java (MSJ) application to obtain the WWN without having to reboot the server. Refer to 4.7, "The FAStT Management Suite Java" on page 133, for information on installing and using this application.

To obtain the WWN, start the FAStT MSJ and connect to the server you want to examine. When the connection is established, highlight the adapter you want. The WWN is listed under the Information tab. See Figure 8-6.

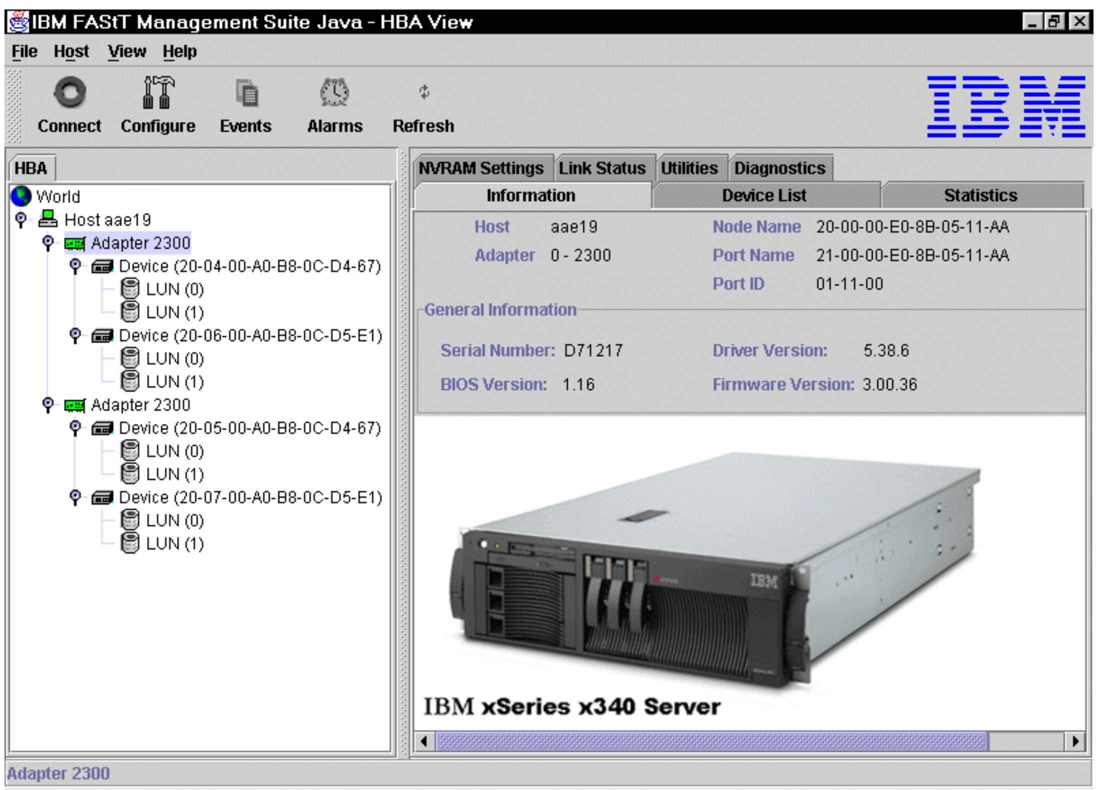

Figure 8-6 FAStT MSJ

Sun Solaris

The Solaris operating system provides the information in the loaded modules and drivers in the system log file, which is normally located in the `/var/adm` directory and called `messages`. After a certain period of time a new system log file is created, also called `messages`, and the old one is renamed to `messages.#`, where # is an increasing number.

Log on as `root` and open the system log file containing the last boot of the server, which may be `/var/adm/messages`. Open the file for reading by issuing less /var/adm/messages.

Search for the last system boot and the initialization of the Fibre Channel adapters. For each adapter, you see the startup procedure. Where the adapter is listed, the WWN is displayed. See Figure 8-7.

```
Mar 4 9:24:40 jnic: jnic0: Hba: fce  Model: FCE-1063  Type: scsi
Mar 4 9:24:40 jnic: jnic0: FCode: Version 2.6 [55E7]  XCode: 2000b [9053]
Mar 4 9:24:40 jnic: jnic0: Slot: 0  IRQ: 3  Mode: 64  Channel: 1  WWN: 200000E069A49497
Mar 4 9:24:40 jnic: jnic0: Configured as Private Loop port
Mar 4 9:24:41 jnic: jnic0: LIP detected (FF)
Mar 4 9:24:41 jnic: jnic0: JNIC v4.0.2.b.3 (00100203)
Mar 4 9:24:41 jnic: jnic0: Copyright(c) 1995-2000 JNI Corp, All Rights Reserved.
Mar 4 9:24:42 jnic: jnic0: Link Up
Mar 4 9:24:43 jnic: jnic0: Port 0000E4 (WWN 200B00A0B80B1148:200B00A0B80B1149) available.
Mar 4 9:24:46 rootnex: pcipsy2 at root: UPA 0xa 0x2000
Mar 4 9:24:46 genunix: pcipsy2 is /pci@a,2000
Mar 4 9:25:18 fca-pci: fca-pci0: JNI Fibre Channel Adapter model FCI-1063
Mar 4 9:25:18 fca-pci: fca-pci0: SCSI ID 125 / AL_PA 0x1
Mar 4 9:25:18 fca-pci: fca-pci0: Fibre Channel WWNN: 100000E069C05BC6  WWPN: 200000E069C05B66
Mar 4 9:25:18 fca-pci: fca-pci0: FCA SCSI/IP Driver Version 2.5, April 18, 2000 for Solaris 7
Mar 4 9:25:18 fca-pci: fca-pci0: All Rights Reserved.
Mar 4 9:25:18 fca-pci: fca-pci0: < Total IOPB space used: 1145024 bytes >
Mar 4 9:25:18 fca-pci: fca-pci0: < Total DMA space used: 8458269 bytes >
Mar 4 9:25:18 fca-pci: NOTICE: pci1242,46431 LOOP Initialization Complete, AL_PA=01
Mar 4 9:25:18 fca-pci: fca-pci0: Host: Port 000001 (100000E069C05BC6:200000E069C05B66)
Mar 4 9:25:19 fca-pci: fca-pci0: Port 0000EF (200A00A0B80B1148:200A00A0B80B1149) available.
```

Figure 8-7 Sun Solaris messages file

HP-UX

HP-UX delivers a small utility that displays the important information about the Fibre Channel adapter. To first determine the device name for the Fibre Channel adapters, log on as `root` and issue the command:

```
ioscan -fn
```

This displays a list of all I/O-related information. You see some entries regarding Fibre Channel adapter as shown in Figure 8-6.

```
Class    I   H/W Path           Driver    S/W State  H/W Type   Description
=====================================================================================

ba       3   0/3                lba       CLAIMED    BUS_NEXUS  Local PCI Bus Adapter
fc       0   0/3/0/0            td        CLAIMED    INTERFACE  HP Tachyon TL/TS Fibre
                                /dev/td0
fcp      4   0/3/0/0.2          fcp       CLAIMED    INTERFACE  FCP Domain
ext_bus  28  0/3/0/0.2.16.0.0   fcparray  CLAIMED    INTERFACE  FCP Array Interface
target   6   0/3/0/0.2.16.0.0.0 tgt       CLAIMED    DEVICE
```

Figure 8-8 HP-UX ioscan

As shown, the device name of the first Fibre Channel adapter is /dev/td0. Now use the following command to display the information on the specific adapter as shown in Figure 8-9:

```
fcsmutil /dev/td0
```

```
-> fcsmutil /dev/td0

                           Vendor ID is = 0x00103c
                           Device ID is = 0x001028
            PCI Sub-system Vendor ID is = 0x00103c
                   PCI Sub-system ID is = 0x000006
                               Topology = PTTOPT_FABRIC
                     Local N_Port_id is = 0x021c00
            N_Port Node World Wide Name = 0x50060b0000072dad
            N_Port Port World Wide Name = 0x50060b0000072dac
                           Driver state = ONLINE
                       Hardware Path is = 0/3/0/0
                 Number of Assisted IOs = 538248
        Number of Active Login Sessions = 1

->
```

Figure 8-9 HP-UX fcsmutil

The WWN is displayed in the output. You can also see the link status and the kind of link. Here the fabric protocol is used.

AIX

On an IBM @server pSeries server, the WWN is written on a label attached to the faceplate of the host bus adapter. You can also obtain it via the command line as explained here.

To obtain the number of Fibre Channel Host Bus Adapters adapters in your server, enter the following command:

```
lsdev -Cc adapter | grep fcs
```

The output should look similar to this example:

```
fcs0 Available 14-08 FC Adapter
fcs1 Available 2A-08 FC Adapter
```

To obtain the WWPN of the Fibre Channel Host Bus Adapters, enter the following command:

```
lscfg -vl fcs0 | grep Net
```

The output looks like this:

```
Network Address.............10000000C9280712
```

The network address is the WWPN. Repeat the same process for the other host bus adapter.

A

Additional host-specific instructions for FlashCopy logical drives

This appendix provides an overview and detailed information of using FlashCopy logical drives with specific operating systems and disk file systems.

Attention: Some of the operating systems that are discussed in this section are not, at the time this redbook was written, formally supported. This includes UNIX, which is explained in “UNIX: Logical Drive Manager logical drives” on page 290. You must check the FAStT supported servers Web site prior to connecting hosts to the FAStT Storage Server or using the advanced functions. You can locate the FAStT700 supported servers Web site at:

http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm

Operating system resources for additional instructions

Refer to one of the following documents to view the additional instructions required for creating FlashCopy logical drives on your host operating system:

Attention: Failure to complete the steps listed for your host operating system may result in a loss of FlashCopy data consistency

- **Windows - Basic/Regular Disks**
  - Process Overview
  - Additional Instructions for Windows 2000 - Basic Disks
  - Additional Instructions for Windows NT - Regular Disks
- **Windows - Dynamic Disks**
  - Process Overview
  - Additional Instructions for Windows 2000 - Dynamic Disks
- **UNIX - Regular Disks**
  - Process Overview
  - Additional Instructions for UNIX - Regular Disks

    Note: Use these instructions for HP-UX, IRIX, Linux, and Solaris
- **UNIX - Logical Drive Manager Logical Drives**
  - Process Overview
  - Additional Instructions for AIX - LVM Logical Logical Drives
  - Additional Instructions for Solaris - Veritas Logical Drive Manager
  - Additional Instructions for HP-UX - Logical Logical Drive Manager
- **NetWare**
  - Process Overview
  - Additional Instructions for NetWare

Note: For detailed procedures of the FlashCopy logical drive creation process, see 4.3, “Creating arrays and logical drives” on page 100.

Windows: Basic/regular disks

This section discusses the procedures for Windows basic/regular disks.

Process overview

The following process overviews outline the key steps required to:

- Create a FlashCopy logical drive on Windows 2000 or Windows NT (using basic or regular disks). See Figure A-1.
- Reuse a disabled FlashCopy logical drive on Windows 2000 or Windows NT. See Figure A-2.

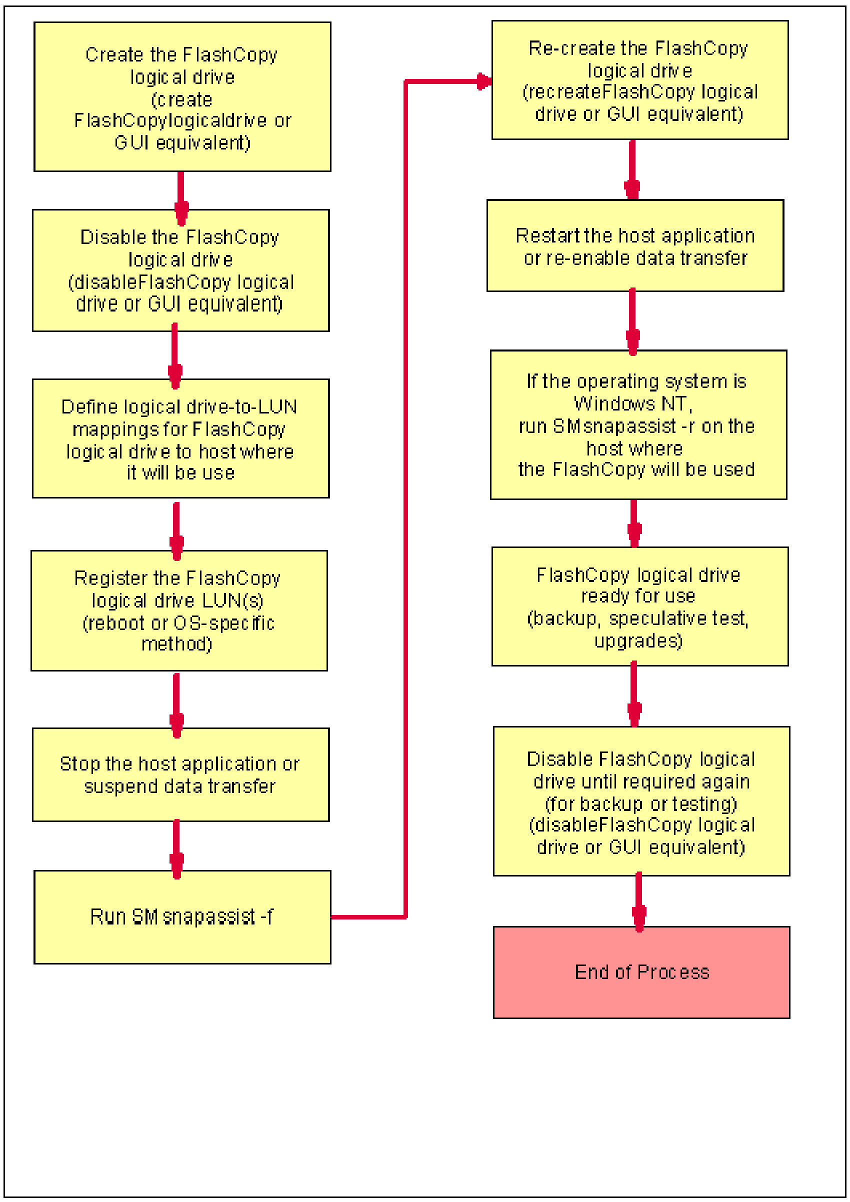

Figure A-1 Creating a FlashCopy logical drive: Windows basic/regular disks

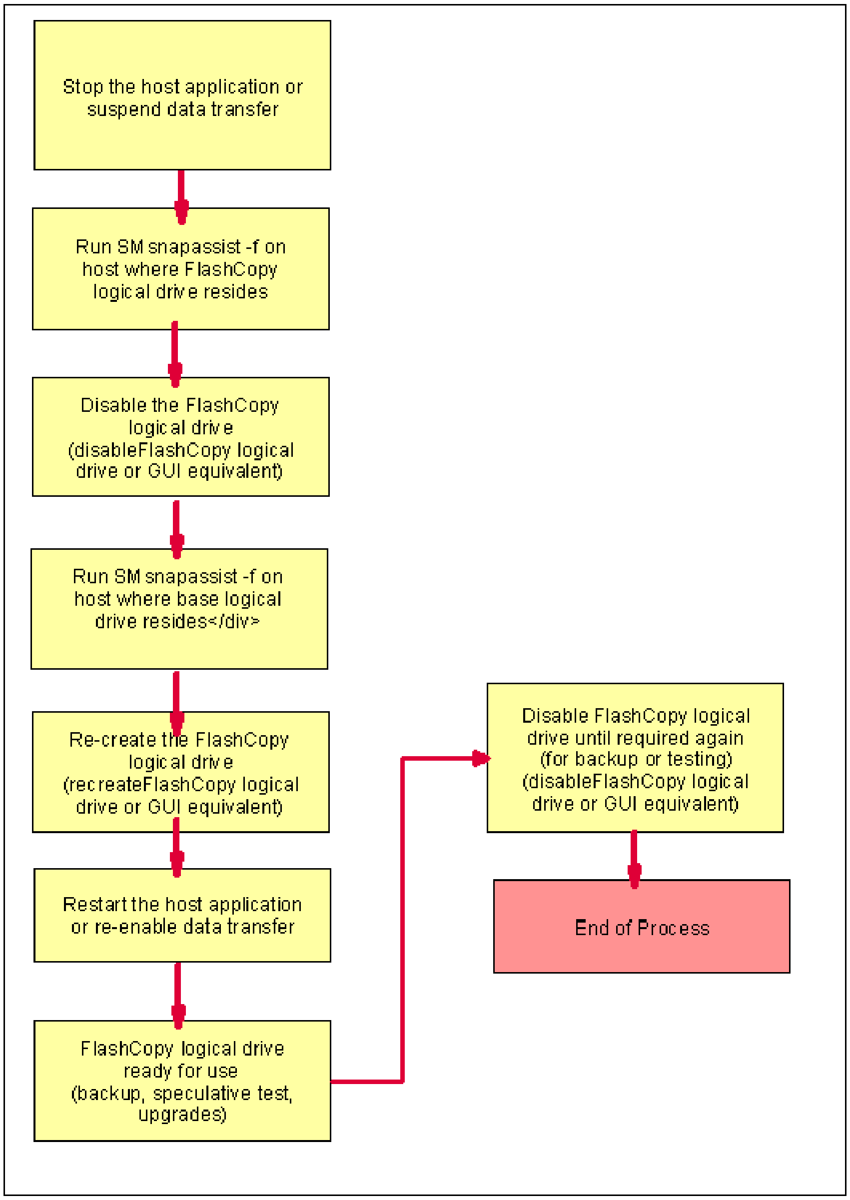

Figure A-2 Re-using disabled FlashCopy logical drives: Windows basic/regular disks

Additional instructions for Windows 2000 basic disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows 2000, using *basic disks*. Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (speculative change or upgrade testing). For instructions on how to reuse a disabled FlashCopy logical drive, see “Reusing FlashCopy logical drives” on page 275.

Tips:

- For command reference information, see Appendix C, “Command line reference” on page 317.
- For information on the use of the script editor and CLI, see 3.3.4, “Script Editor and command line interface” on page 58.

Creating a FlashCopy logical drive

To create a FlashCopy logical drive, follow these steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window, using one of the following methods:
   - Select the storage subsystem in either the Device Tree View or the Device Table. Then click the **Manage Device** toolbar button, or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table, and then select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in either the Device Tree View or the Device Table, and then press Enter.

   The Subsystem Management window opens in a separate window.
3. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   - Create FlashCopy Logical Drive Wizard, which you can access by using the Subsystem Management window.
   - Create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```
4. Disable the FlashCopy logical drive using one of the following methods:
   - Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.
   - Disable a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     disableFlashCopy logical drive
     ```
5. Assign a logical drive-to-Logical Unit Numbers (LUN) mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive. Logical drive-to-LUN mappings can be defined using one of the following methods:
   - Storage Partitioning Wizard, which helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive and the associated LUNs.
   - Create a logical drive-to-LUN mapping using either the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     create mapping logical drive
     ```

   Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows Explorer) to ensure all I/O activity is stopped.

6. Run the SMsnapassist utility to flush all the write buffers from the new disk drive. At a host prompt, type the following command and then press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive.

   The write buffers for the disk drive are flushed, which you can see in this example. If the new disk drive (for the base logical drive) is assigned drive letter "E" using the Create Partition Wizard, you type the following command and press Enter:

   ```
   SMsnapassist -f e:
   ```

7. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

   – Re-create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     recreateFlashCopy logical drive
     ```

   Important: If I/O activity to the base logical drive stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable a data transfer).

8. Run the hot_add utility (or operating system-specific utility) or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.

9. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the SMdevices utility is run to ensure that the logical drive name, and the operating system device name (assigned by the operating system) correlate.

10. Open **Disk Management** using one of the following methods:

    – Click **Start-> Settings-> Control Panel**. Double-click the **Administrative Tools** icon and then double-click **Computer Management**. In the console tree under **Storage**, select **Disk Management**.

    – Select **Start-> Run** option from the desktop. Then type the following command and press Enter:

      ```
      compmgmt.msc
      ```

    – At a host prompt, type the following command and press Enter:

      ```
      compmgmt.msc
      ```

    – In the console tree, under **Storage**, select **Disk Management**.

    Disk Management is displayed with a graphical representation of all the physical disks connected to the host and their associated partitions.

11. In the Disk Management dialog, locate the disk and logical drive definition that represents the FlashCopy logical drive you re-created and ensure that a new drive letter is automatically assigned.

12. Use the FlashCopy logical drive in conjunction with your backup application (reusing a FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).
13. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

    If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

> **Important:** All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows Explorer) to ensure all I/O activity is stopped.

1. Run the SMsnapassist utility on the host, where the FlashCopy logical drive is mounted, to flush all the write buffers from the new disk drive. At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the FlashCopy logical drive. The write buffers for the disk drive are flushed.
2. Disable the FlashCopy logical drive using one of the following methods:
   - Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.
   - Disable a FlashCopy logical drive using either the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     disableFlashCopy logical drive
     ```
3. Run the SMsnapassist utility on the host, where the base logical drive is mounted, to flush all the write buffers from the new disk drive. At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive. The write buffers for the disk drive are flushed.
4. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   - Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.
   - Recreate a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     recreateFlashCopy logical drive
     ```

> **Important:** If I/O activity to the base logical drive stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable the data transfer).

5. Use the FlashCopy logical drive in conjunction with your backup application (or with another application).
6. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Additional instructions for Windows NT regular disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows NT, using regular disks. Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (for speculative change or upgrade testing).

Creating a FlashCopy logical drive

Follow these steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch an Subsystem Management window, using one of the following methods:
   - Select the storage subsystem in either the Device Tree View or the Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table, and select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in the Device Tree View or the Device Table and then press Enter.

   The Subsystem Management window opens in a separate window.
3. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   - Create FlashCopy Logical Drive Wizard, accessed using an Subsystem Management window.
   - Create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```
4. Disable the FlashCopy logical drive. FlashCopy logical drives may be disabled using one of the following methods:
   - Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.

– Disable a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

```
disableFlashCopy logical drive
```

5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive.

   Logical drive-to-LUN mappings can be defined using one of the following methods:

   – Storage Partitioning Wizard: Helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive and the associated LUNs.

   – Create a logical drive-to-LUN mapping using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   create mapping logical drive
   ```

   Important: All I/O activity to the base logical drive must be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows NT Explorer) to ensure all I/O activity is stopped.

6. Run the SMsnapassist utility to flush all the write buffers from the disk drive.

   Note: Running this utility is only required when the FlashCopy logical drive is mapped to the same host as the base logical drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive.

   The write buffers for the disk drive are flushed. This is shown in the following example. If the new disk drive (for the base logical drive) is assigned drive letter "E" using the Create Partition Wizard, you type the following command and press Enter:

   ```
   SMsnapassist -f e
   ```

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

7. Run the hot_add utility (or operating system-specific utility), or reboot the host where the FlashCopy will be used, to ensure that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.

8. Run the SMdevices utility to associate the logical drive-to-LUN mappings with the host operating system.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the SMdevices utility is run to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

9. Run the SMsnapassist utility to resolve duplicate signatures and partition table information.

At the host prompt, type the following command and press Enter:

```
SMsnapassist -r
```

The utility reports back that a new signature is written for the FlashCopy logical drive.

10. Open the Disk Administrator using one of the following methods:

    – Click **Start-> Programs-> Administrative Tools** and select **Disk Administrator**.

    – Click **Start-> Run**. Then type the following command and press Enter:

    ```
    windisk
    ```

    – At the host prompt, type the following command and press Enter:

    ```
    windisk
    ```

    The Disk Administrator window opens, with a graphical representation of all the physical disks connected to the host and their associated partitions.

11. Locate the disk that represents the FlashCopy logical drive in the Disk Administrator window.

12. Select the disk, and click **Tools-> Assign Drive Letter**.

13. The Assign Drive Letter dialog opens. Choose a drive letter to be assigned (if not automatically assigned) and select **OK**.

    The disk drive representing the FlashCopy logical drive is assigned a new drive letter. You are not required to partition or format this drive.

14. Use the FlashCopy logical drive in conjunction with your backup application (reusing a FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).

15. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

    If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows NT Explorer) to ensure all I/O activity is stopped.

1. Run the SMsnapassist utility on the host, where the FlashCopy logical drive is mounted, to flush all the write buffers from the disk drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the FlashCopy logical drive. The write buffers for the disk drive are flushed.

2. Disable the FlashCopy logical drive. FlashCopy logical drives may be disabled using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.

   – Disable a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   disableFlashCopy logical drive
   ```

3. Run the SMsnapassist utility on the host where the base logical drive is mounted to flush all the write buffers from the disk drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive. The write buffers for the disk drive are flushed.

4. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

   – Re-create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

5. Use the FlashCopy logical drive in conjunction with your backup application (or with another application).

6. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive. It also stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Windows: Dynamic disks

This section discusses the procedures for Windows dynamic disks.

Process overview

The following process overviews outline the key steps required to:

- Create a FlashCopy logical drive on Windows 2000 (using dynamic disks). See Figure A-3.
- Reuse a disabled FlashCopy logical drive on Windows 2000 (using dynamic disks). See Figure A-4.

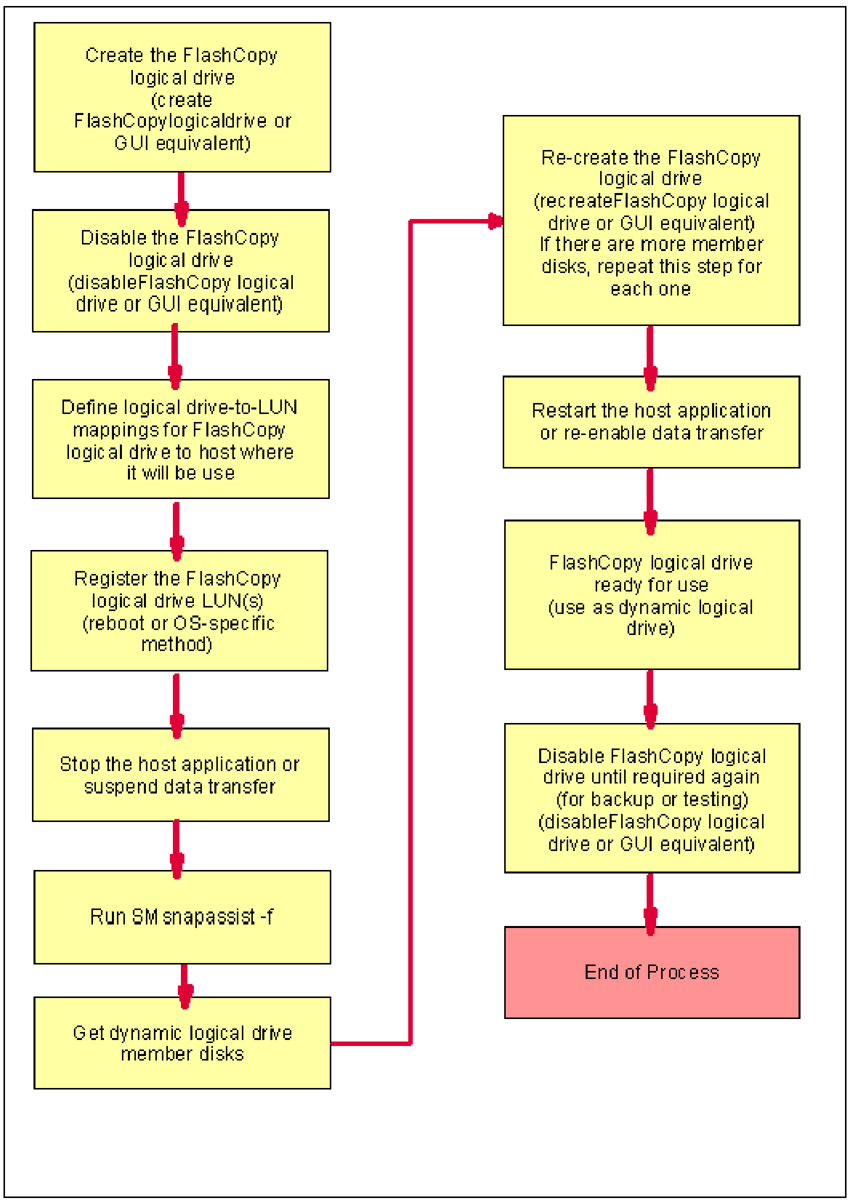

Figure A-3 Creating a FlashCopy logical drive: Windows 2000 dynamic disks

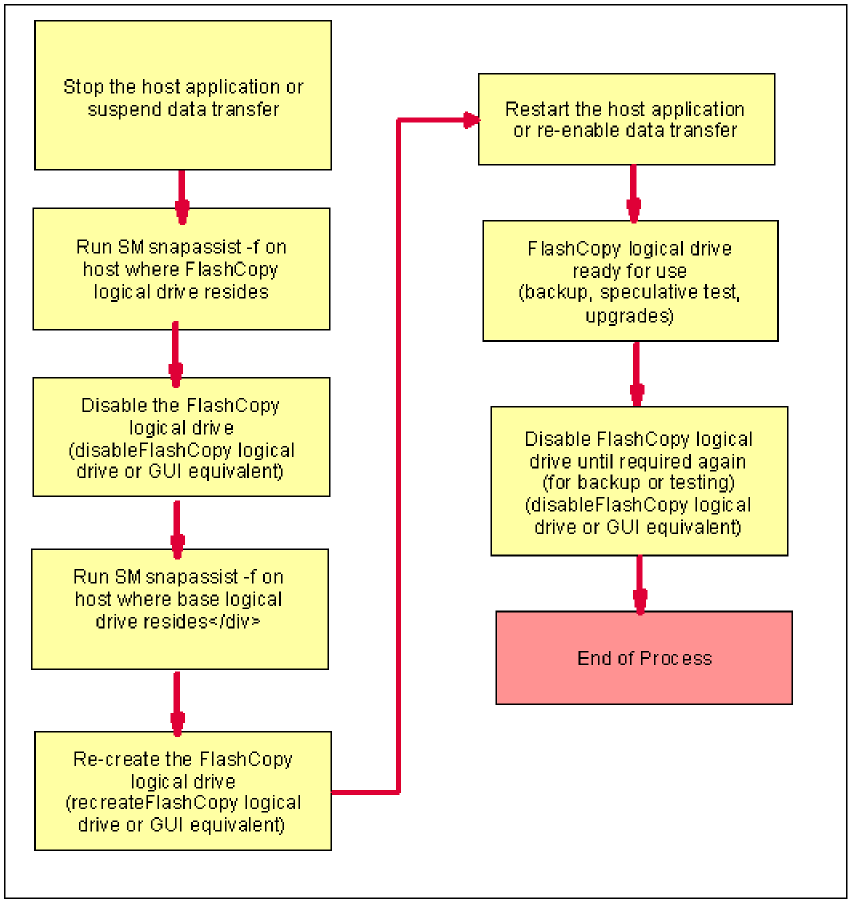

Figure A-4 Reusing a disabled FlashCopy logical drive: Windows 2000 dynamic disks

Additional instructions for Windows 2000 dynamic disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows 2000, using *dynamic disks*. Failure to complete the steps may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (for speculative change or upgrade testing).

Important: FlashCopy logical drives created on a host running Windows 2000, using dynamic disks, may not be mapped on the same host as the base logical drives.

Creating a FlashCopy logical drive

To create a FlashCopy logical drive on a host running Windows 2000, with dynamic disks, complete these steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch an Subsystem Management window, using one of the following methods:
   - Select the storage subsystem in the Device Tree View or Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table and select **Manage Device**.

– Double-click a storage subsystem node in the Device Table.
– Select the storage subsystem in the Device Tree View or Device Table and press Enter.

The Subsystem Management window opens in a separate window.

3. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   – Create FlashCopy Logical Drive Wizard, accessed using an Subsystem Management window.
   – Create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```
4. Disable the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.
   – Disable a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     disableFlashCopy logical drive
     ```
5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that accesses the FlashCopy logical drive. Use one of the following methods to define a logical drive-to-LUN mappings:

   Important: FlashCopies of Windows 2000 dynamic disks may not be mapped on the same host as the base logical drives.

   – Storage Partitioning Wizard: Helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive and the associated LUNs.
   – Create a logical drive-to-LUN mapping using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

     ```
     create mapping logical drive
     ```

   Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows Explorer) to ensure all I/O activity is stopped.

6. Run the SMsnapassist utility to flush all the write buffers from the new disk drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive.

   The write buffers for the disk drive are flushed. This is shown in this example. If the new disk drive (for the base logical drive) is assigned drive letter E:, enter the following command and press Enter:

   ```
   SMsnapassist -f e:
   ```
7. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

– Recreate a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

```
recreateFlashCopy logical drive
```

Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

8. Run the hot_add utility (or operating system-specific utility) or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.

9. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the SMdevices utility is run to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

10. Open Disk Management using one of the following methods:

   – Click **Start-> Settings-> Control Panel**. Double-click the **Administrative Tools** icon, and then double-click **Computer Management**. In the console tree under **Storage**, select **Disk Management.**

   – Click **Start-> Run**. Then type the following command and press Enter:

   ```
   compmgmt.msc
   ```

   – At a host prompt, type the following command and then press Enter:

   ```
   compmgmt.msc
   ```

   – In the console tree under **Storage**, select **Disk Management**.

   The Disk Management windows opens with a graphical representation of all the physical disks connected to the host and their associated partitions.

11. Locate the disk and logical drive definition that represents the FlashCopy logical drive.

   Note: The FlashCopy logical drives LUNs are displayed with the disk type “foreign”.

12. Select the FlashCopy logical drive LUNs, right-click, and select **Import Foreign Disks**. The Import Foreign Disks dialog opens.

13. Select the appropriate disk and click **OK**. The FlashCopy logical drives and LUNs are imported.

14. Perform a manual re-scan to verify that the disk information is correct. From the Disk Management main menu, click **Action-> Rescan Disks**.

15. In Disk Management, locate the disk and logical drive definition that represents the FlashCopy logical drive you re-created. Ensure that a new drive letter is automatically assigned.

16. Use the FlashCopy logical drive in conjunction with your backup application (reusing a FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).

17. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured. Close all applications (including Windows Explorer) to ensure all I/O activity is stopped.

1. Run the SMsnapassist utility on the host where the FlashCopy logical drive is mounted to flush all the write buffers from the new disk drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the FlashCopy logical drive. The write buffers for the disk drive are flushed.

2. Disable the FlashCopy logical drive using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.

   – Disable a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   disableFlashCopy logical drive
   ```

3. Run the SMsnapassist utility on the host where the base logical drive is mounted to flush all the write buffers from the new disk drive.

   At the host prompt, type the following command and press Enter:

   ```
   SMsnapassist -f <filesystem-identifier>
   ```

   Here *<filesystem-identifier>* is the drive letter assigned to the base logical drive. The write buffers for the disk drive are flushed.

4. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

   – Recreate a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

5. Use the FlashCopy logical drive in conjunction with your backup application (or another application).
6. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.
   – Recreate a FlashCopy logical drive using either the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     recreateFlashCopy logical drive
     ```

UNIX: Regular disks

This section explains the procedures for UNIX regular disks.

Process overview

The following process overview outlines the key steps required to create and reuse a FlashCopy logical drive on the following UNIX-based operating systems, using regular disks:

- HP-UX 11.0 (or higher)
- Linux RedHat 6.2 or 7.x
- Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8)

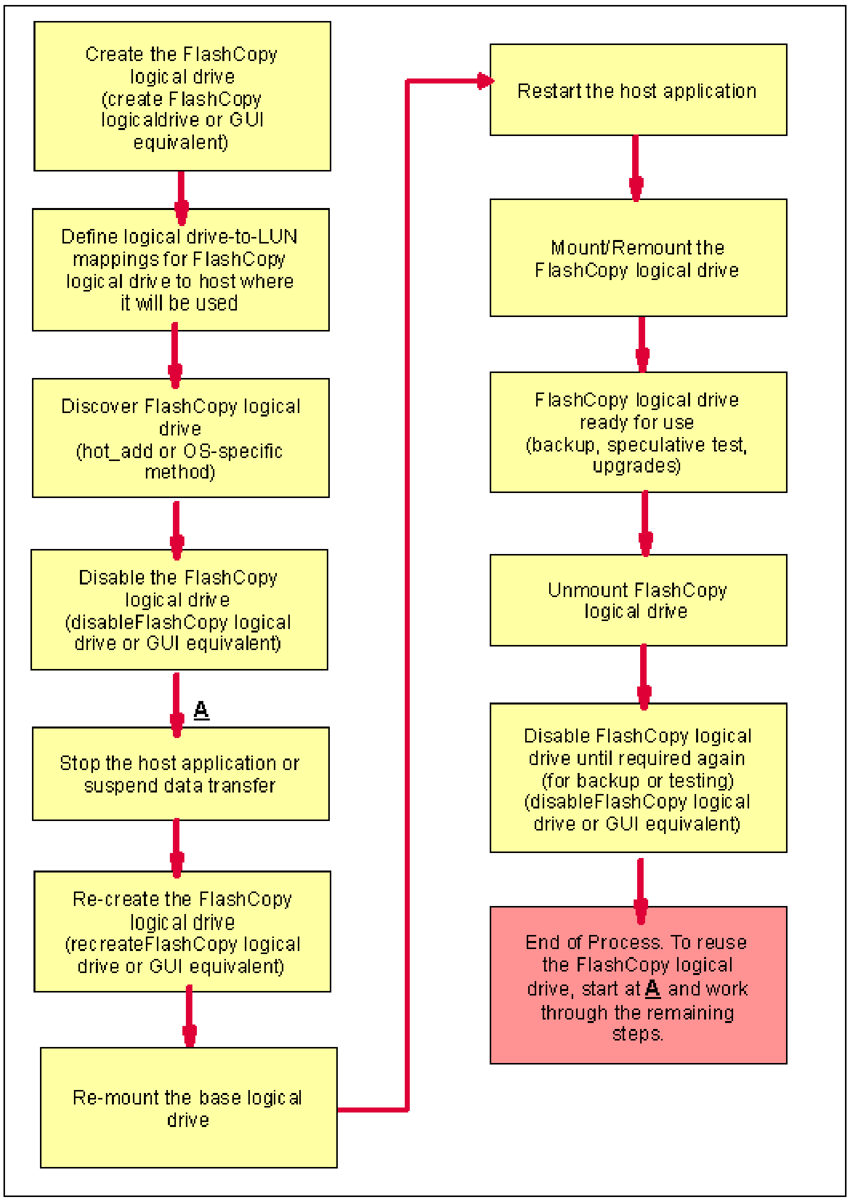

Figure A-5 FlashCopy logical drive on UNIX-based operating systems

Additional instructions for UNIX using regular disks

Use the following procedure when creating FlashCopy logical drives on a host running one of the following UNIX-based operating systems:

- HP-UX 11.0 (or higher)
- Linux RedHat 6.2 or 7.1
- Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8)

Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (speculative change or upgrade testing).

Creating a FlashCopy logical drive

Before you create a FlashCopy logical drive, consider the following points:

- Is the FlashCopy logical drive being created for immediate use? If the FlashCopy logical drive is to be created and then used immediately, all I/O activity to the base logical drive should be stopped. Also unmount the base logical drives from the host to which they are currently mounted. This ensures that an accurate point-in-time image of the data on the base logical drive is captured.
- Is the FlashCopy logical drive being created, but to be used at a later date? If the FlashCopy logical drive is to be created and used at a later date, do not stop I/O activity to the base logical drive or unmount at this time. This needs to be carried out the first time you want to use the FlashCopy logical drive.

To create a FlashCopy logical drive, complete the following steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window using one of the following methods:
   - Select the storage subsystem in either the Device Tree View or Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table, and select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in the Device Tree View or the Device Table, and then press Enter.

   The Subsystem Management window opens in a separate window.
3. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   - Create FlashCopy Logical Drive Wizard, accessed using a Subsystem Management window.
   - Create a FlashCopy logical drive using either the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```
4. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   - Create FlashCopy Logical Drive Wizard, accessed using an Subsystem Management window.

   – Create a FlashCopy logical drive using either the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```

5. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host. For information on which operating systems support the hot_add utility, see the *Storage Manager Software Installation Guide*, which is shipped on CD with the product. You can also download it from the Web at:

   http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700

   For the *Solaris* operating system, at the host prompt, type the following command and press Enter:

   ```
   /etc/raid/bin/hot_add
   ```

   Once logical drives are created and logical drive-to-LUN mappings are defined, this step ensures that the operating system is aware of the newly created logical drives, without having to reboot the host.

6. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the SMdevices utility is run to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

7. If the FlashCopy logical drive is being created for immediate use, go to step 8. If the FlashCopy logical drive is being created for use at a later date, disable the FlashCopy logical drive now using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.

   – Disable a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     disableFlashCopy logical drive
     ```

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

8. Mount the FlashCopy logical drive to its intended host.

9. Use the FlashCopy logical drive in conjunction with your backup application, for speculative testing, or with another application.

10. Unmount the FlashCopy logical drive.

11. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

    If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the base logical drive.
2. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   - Click **Logical Drive-> FlashCopy-> Re-create** from the menus in the Subsystem Management window.
   - Re-create a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     recreateFlashCopy logical drive
     ```
3. Remount the base logical drive (to its original host).
4. Mount the FlashCopy logical drive to its intended host.

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).
5. Use the FlashCopy logical drive in conjunction with your backup application (or with another application).
6. Unmount the FlashCopy logical drive.
7. When use of the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

UNIX: Logical Drive Manager logical drives

This section outlines the procedures when a UNIX Logical Drive Manager is used to manager the FAStT logical drives.

Process overview

The following process overview outlines the key steps required to create and reuse a FlashCopy logical drive on the following UNIX operating systems, using Logical Drive Manager logical drives:

- AIX with Logical Logical Drive Manager
- Solaris with Veritas Logical Drive Manager
- HP-UX with Logical Logical Drive Manager

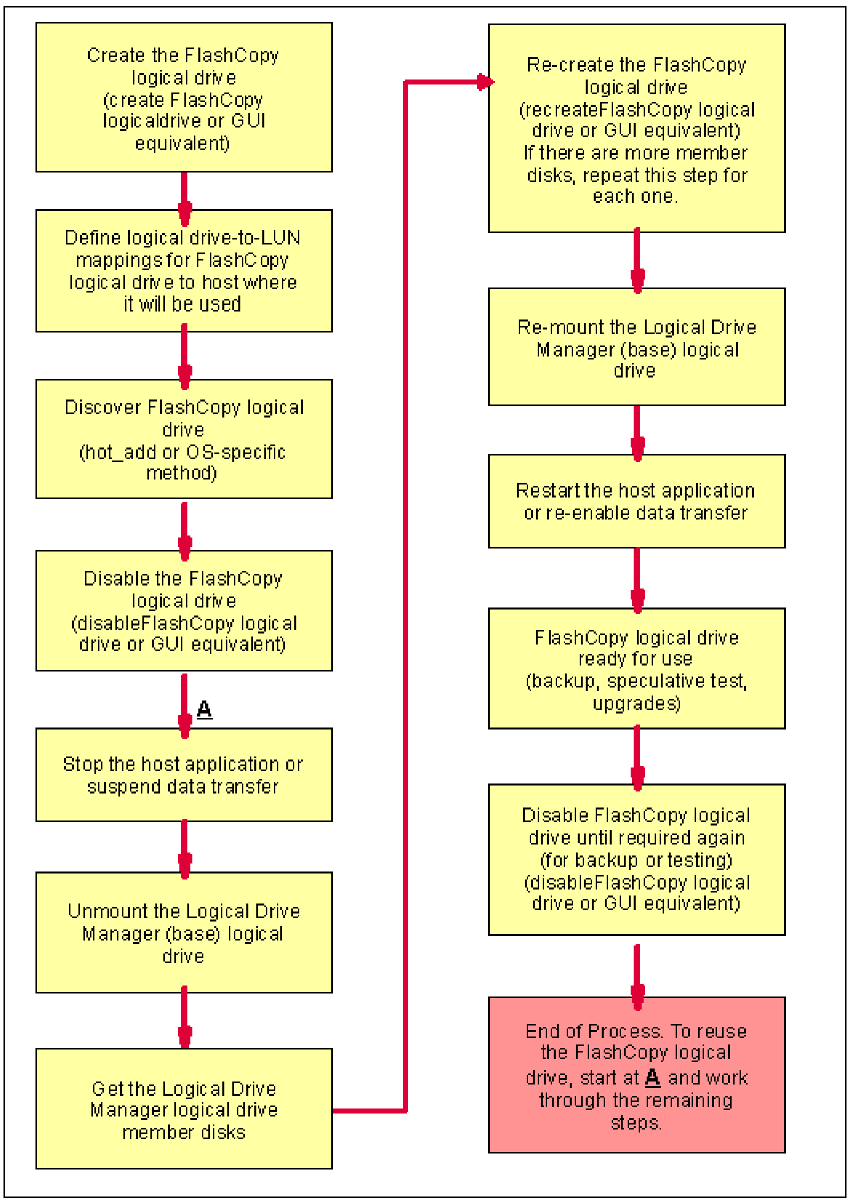

Figure A-6 FlashCopy logical drive UNIX operating systems

Attention: At the time this redbook was written, use of the RVM and FlashCopy advanced functions is not formally supported on AIX, Solaris, and HP-UX. You must check the supported server Web site prior to using this function:

http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm

Additional instructions for AIX: LVM Logical Logical Drives

Use the following procedure when creating FlashCopy logical drives on a host running AIX 4.3.3 (or higher), using LVM Logical Logical Drives. Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

Restriction: The following restrictions apply when attempting to create FlashCopy logical drives on AIX:

- To map FlashCopy logical drives to the same host as the base logical drive, you must use:
  - AIX version 4.3.3 with the AIX version 4.3.3 Recommended Maintenance Level 06 (AIX 4330-06) maintenance package
  - AIX version 5.1
- FlashCopy logical drives may only be created for AIX arrays. If the array has more than one logical drive, FlashCopy logical drives must be created for each logical drive in the array.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (for speculative change or upgrade testing).

Creating a FlashCopy logical drive

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Locate the array (on which the FlashCopy logical drive is to be based) and unmount its associated filesystems.

   At the host prompt, type the following command and press Enter:

   ```
   umount mount-point
   ```

   Here *mount-point* is the name of the filesystem being unmounted.
2. Start the Storage Management software. The Enterprise Management window opens.
3. Launch an Subsystem Management window using one of the following methods:
   - Select the storage subsystem in the Device Tree View or the Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table, and select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in the Device Tree View or Device Table and press Enter.

   The Subsystem Management window opens in a separate window.

> **Important:** If an AIX array has more than one logical drive, create FlashCopy logical drives for each logical drive in the array.

4. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using one of the following methods:
   - Create FlashCopy Logical Drive Wizard, accessed using an Subsystem Management window.
   - Create a FlashCopy logical drive using either the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     create FlashCopylogicaldrive
     ```
5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drives and the host that will access the FlashCopy logical drives. Logical drive-to-LUN mappings can be defined using one of the following methods:
   - Storage Partitioning Wizard: Helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive and the associated LUNs.
   - Create a logical drive-to-LUN mapping using either the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

     ```
     create mapping logical drive
     ```
6. Login to the host as `root`.
7. Run the hot_add utility (or an operating-system specific utility) to ensure that the host operating system recognizes the FlashCopy logical drives.

   To use hot_add, at the host prompt, type the following command and then press Enter:

   ```
   hot_add
   ```

   Several minutes may pass while the computer accesses the drives. When the program is finished, a window opens with the following message:

   ```
   Device nodes have been updated
   ```

   The new logical drives should now be available to you through the operating system.
8. Run the SMdevices utility to associate the LUNs with a host operating system device and to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

   Look for the newly created FlashCopy logical drive names and note the names of the associated operating system device name. For example, you created a FlashCopy logical drive named *accountingMay14* and its associated operating system device name is *hdisk4*.
9. At the host prompt, type the following command and press Enter:

   ```
   lspv
   ```

   A list of all the physical logical drives recognized by the host operating system appears.
10. Look for the operating system device name of your FlashCopy logical drive or logical drives in the list. The listing shows a Physical Logical Drive ID (PVID) for this logical drive that will be the same as the PVID for the associated base logical drive. This is because the FlashCopy logical drive contains the same array data structures as the base logical drive.
11. Clear the PVID for the FlashCopy logical drives. At the host prompt, type the following command and press Enter:

    ```
    chdev - l os_device_name -a pv=clear
    ```

Here *os_device_name* is the operating system device name of the FlashCopy logical drive. Repeat this step for each FlashCopy logical drive in the AIX array.

12. Recreate a new array. The **recreatevg** command, available in usr/sbin/, reads the array data structure inside a logical drive and reconstructs it. The command allocates new physical logical drive identifiers (PIDs) to the FlashCopy logical drives and enables access to the FlashCopy logical drive for the selected host.

> **Important:** If this command is not available in AIX 4.3.3., install the AIX version 4.3.3 Recommended Maintenance Level 06 (AIX 4330-06) maintenance package.

At the host prompt, type the following command and press Enter:

```
recreatevg -y logical drivegroupname -L /directoryname os_device_name
```

Note the following explanations, where:

– *logical drivegroupname* is the user-defined name to be assigned to the FlashCopy array.

– *directoryname* is the name of the directory where you want to mount the FlashCopy logical drive.

– *os_device_name* is the operating system device name of the FlashCopy logical drive. If your AIX array is made up of more than one FlashCopy logical drive, add an os_device_name for each logical drive.

The array is recreated and contains the FlashCopy logical drive or logical drives.

13. Mount the filesystem to its intended host. At the host prompt, type the following command and then press Enter:

```
mount mount-point
```

Here *mount-point* is the name of the filesystem being mounted. Include the directoryname used in step 12.

14. Ensure that the logical logical drives are back online. At the host prompt, type the following command and press Enter:

```
df -k
```

A list of the mounted disks appears.

15. Use the FlashCopy logical drive in conjunction with your backup application, for speculative testing, or with another application.

16. When the FlashCopy logical drive is no longer required, unmount the filesystem. At the host prompt, type the following command and press Enter:

```
umount mount-point
```

Here *mount-point* is the name of the filesystem being unmounted.

17. Delete the array created in step 12 that contains the FlashCopy logical drives. At the host prompt, type the following command and press Enter:

```
varyoffvg logical drivegroupname
exportvg logical drivegroupname
```

Here *logical drivegroupname* is the name of the FlashCopy array.

18. Disable or delete the FlashCopy logical drive or logical drives.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

> **Important:** All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the filesystems in the array on which the FlashCopy logical drive is to be based.

   At the host prompt, type the following command and press Enter:

   `umount` *`mount-point`*

   Here *mount-point* is the name of the filesystem being unmounted.

2. Run the hot_add utility (or an operating-system specific utility) to ensure that the host operating system recognizes the FlashCopy logical drive. At the host prompt, type the following command and press Enter:

   ```
   hot_add
   ```

   Several minutes may pass while the computer accesses the drives. When the program is finished, a window displays the following message:

   ```
   Device nodes have been updated
   ```

   The new logical drives should now be available to you through the operating system.

3. In the Storage Management software, re-create the FlashCopy logical drives using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

   – Re-create a FlashCopy logical drive using either the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```

4. Clear the PVID for the FlashCopy logical drives. At the host prompt, type the following command and press Enter:

   `chdev - l` *`os_device_name`* `-a pv=clear`

   Here *os_device_name* is the operating system device name of the FlashCopy logical drive. Repeat this step for each FlashCopy logical drive in the AIX array.

5. Recreate a new array. The **`recreatevg`** command, available in usr/sbin/, reads the array data structure inside a logical drive and reconstructs it. The command allocates new physical logical drive identifiers (PIDs) to the FlashCopy logical drives and enables access to the FlashCopy logical drive for the selected host.

   > **Important:** If this command is not available in AIX 4.3.3., install the AIX version 4.3.3 Recommended Maintenance Level 06 (AIX 4330-06) maintenance package.

   At the host prompt, type the following command and press Enter:

   `recreatevg -y` *`logical drivegroupname`* `-L /`*`directoryname os_device_name`*

Note the following explanations:

– *logical drivegroupname* is the user-defined name to be assigned to the FlashCopy array.

– *directoryname* is the name of the directory where you want to mount the FlashCopy logical drive.

– *os_device_name* is the operating system device name of the FlashCopy logical drive. If your AIX array is made up of more than one FlashCopy logical drive, add an *os_device_name* for each logical drive.

The array is recreated and contains the FlashCopy logical drive or logical drives.

6. Mount the filesystem to its intended host. At the host prompt, type the following command and then press Enter:

   ```
   mount mount-point
   ```

   Here *mount-point* is the name of the file system being mounted. Include the *directoryname* used in step 5.

7. Ensure that the logical logical drives are back online. At the host prompt, type the following command and press Enter:

   ```
   df -k
   ```

   A list of the mounted disks is displayed.

8. Use the FlashCopy logical drive in conjunction with your backup application (or another application).

9. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Additional instructions for Solaris: Veritas Logical Drive Manager

Use the following procedure when creating FlashCopy logical drives on a host running Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8) using Veritas Logical Drive Manager logical drives. Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (for speculative change or upgrade testing).

Attention: FlashCopy logical drives created on a host running Solaris (where the base logical drive is under Veritas Logical Drive Manager control) may not be mapped to the same host as the base logical drive.

Creating a FlashCopy logical drive

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the disk representing the base logical drive. At the host prompt, type the following command and press Enter:

   ```
   umount mount-point
   ```

   Here *mount-point* is the name of the disk being unmounted.

2. Start the Storage Management software. The Enterprise Management window opens.
3. Launch an Subsystem Management window, using one of the following methods:

   > **Note:** If your configuration spans across a number of storage subsystems, ensure that the procedure is repeated for each storage subsystem.

   – Select the storage subsystem in either the Device Tree View or the Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.

   – Right-click the storage subsystem in the Device Tree View or Device Table and select **Manage Device**.

   – Double-click a storage subsystem node in the Device Table.

   – Select the storage subsystem in the Device Tree View or Device Table, and then press Enter.

   The Subsystem Management window opens in a separate window.

4. Perform a **sync** to ensure that all previously unwritten system buffers are flushed out to disk. This ensures that all file modifications up to that point are saved. At the host prompt, type the following command and press Enter:

   ```
   sync
   ```

   All unwritten system buffers are flushed.

5. Within the Logical View of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive. Create a FlashCopy logical drive using the Script Editor or a command line interface on the host operating system. Type the following command and press Enter:

   ```
   create FlashCopylogicaldrive
   ```

6. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive.

   > **Important:** FlashCopy logical drives created on a host running Solaris (where the base logical drive is under Veritas Logical Drive Manager control) may not be mapped to the same host as the base logical drive.

   Logical drive-to-LUN mappings can be defined using one of the following methods:

   – Storage Partitioning Wizard: Helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive and the associated LUNs.

   – Create a logical drive-to-LUN mapping using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   create mapping logical drive
   ```

7. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.

8. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the SMdevices utility is run to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.

9. Open Veritas Logical Drive Manager Storage Administrator and scan all the mounted disks. To perform a scan, use one of the following methods:

   – Select the host where the FlashCopy logical drive resides, and click **Host-> Scan Disks**. Or you can right-click and select **Scan Disks**.

   – Select **Disk Scan** from the Command Launcher.

   A scan of all the mounted disks is performed.

10. Import the disk group that will enable access to a disk group for the selected host.

    a. At the host prompt, type the following command and press Enter:

    ```
    vxdiskadm
    ```

    b. To ensure that FlashCopy logical drives are available to be imported, at the *vxdiskadm* main menu, type the following command and press Enter:

    ```
    list
    ```

    c. Select the **Enable access** (import) menu option for a disk group from the main menu.

    d. Enter the name of the disk group to be imported and press Enter.

    e. Select **No** to choose to import another disk group.

    Note: If importing the disk group fails using the above method (or using the Veritas Logical Drive Manager System Administrator main screen), at the host prompt, type the following command and press Enter:

    `vxdg -C import` *`disk group`*

    Here *disk group* is the name of the disk group to be imported. All import locks are cleared and the disk group is imported.

11. Start the logical drive to make it available for use. At the host prompt, type the following command and then press Enter:

    `vxvol start` *`logical drive`*

    Here *logical drive* is the name of the FlashCopy logical drive.

    The defined logical drive changes state from Disabled to Enabled, and is now ready for use.

12. In the Veritas Logical Drive Manager System Administrator, mount the filesystem associated with the disk group.

    Select the filesystem associated with the disk group, right-click, and select **Filesystem-> Mount**. The disk groups associated with the filesystem are mounted.

13. On the host where the FlashCopy logical drive resides, ensure that the filesystem was mounted correctly. At the host prompt, type the following command and press Enter:

    ```
    df-k
    ```

14. At the host prompt, type the following command and press Enter:

```
cd mount-point
```

Here *mount-point* is the directory where the FlashCopy logical drive is mounted.

Locate the directory where the FlashCopy logical drive was mounted, and ensure that the FlashCopy logical drives contents match the original contents of the base logical drive.

15. Use the FlashCopy logical drive in conjunction with your backup application (reusing a FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).
16. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the disk representing the base logical drive and the disk representing the FlashCopy logical drive. At the host prompt, type the following command and press Enter:

```
umount moint-point
```

Here *mount-point* is the name of the disk being umounted.

2. Perform a **sync** to ensure that all previously unwritten system buffers are flushed out to disk. This ensures that all file modifications up to that point are saved. At the host prompt, type the following command and press Enter:

```
sync
```

All unwritten system buffers are flushed.

3. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:

Note: If your configuration environment spans across a number of storage subsystems, ensure that the procedure is repeated for each storage subsystem.

– Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.

– Re-create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

```
recreateFlashCopy logical drive
```

4. In the Veritas Logical Drive Manager System Administrator, mount the filesystem associated with the disk group:

Select the filesystem associated with the disk group. Then right-click and select **Filesystem-> Mount**. The disk groups associated with the filesystem are mounted.

5. On the host where the FlashCopy logical drive resides, ensure that the filesystem was mounted correctly. At the host prompt, type the following command and press Enter:

   ```
   df-k
   ```

6. Use the FlashCopy logical drive in conjunction with your backup application (or another application).
7. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Additional instructions for HP-UX: Logical Logical Drive Manager

Use the following procedure when creating FlashCopy logical drives on a host running HP-UX 11.0 (or higher) using LVM logical logical drives. Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (speculative change or upgrade testing).

Creating a FlashCopy logical drive

To create a FlashCopy logical drive, follow these steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch an Subsystem Management window, using one of the following methods:
   - Select the storage subsystem in the Device Tree View or Device Table. Then click the **Manage Device** toolbar button, or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table and select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in either the Device Tree View or the Device Table and press Enter.

   The Subsystem Management window opens in a separate window.

   Important: Stop the host application accessing the *base logical drive* and unmount the base logical drive. Unmounting the base logical drive does not apply when the base logical drive is the root disk of the host operating system.

3. When creating a FlashCopy logical drive based on a mounted filesystem, always perform a **`sync`** to flush the filesystem cache immediately prior to creating a FlashCopy logical drive.

   At the host prompt, type the following command and press Enter:

   ```
   sync
   ```

   All unwritten filesystem buffers are flushed.

4. Within the Logical View of the Subsystem Management window, select a *standard logical drive*, and create a FlashCopy logical drive using one of the following methods:
   – Create FlashCopy Logical Drive Wizard, accessed using an Subsystem Management window.
   – Create a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:
     ```
     create FlashCopylogicaldrive
     ```
5. Assign a *logical drive-to-LUN mapping* between the FlashCopy logical drive and the host that will access the FlashCopy logical drive. Logical drive-to-LUN mappings can be defined using one of the following methods:
   – Storage Partitioning Wizard: Helps you to quickly define a single storage partition. It guides you through the major steps required to specify which host will access a logical drive, and the associated LUNs.
   – Create a logical drive-to-LUN mapping using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:
     ```
     create mapping logical drive
     ```
6. Disable the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Disable** in the Subsystem Management window.
   – Disable a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:
     ```
     disableFlashCopy logical drive
     ```
7. At the host prompt, type the following command and press Enter:
   ```
   ioscan -fn
   ```
   A list of the mapped devices recognized by the host is displayed.

   > **Note:** If the required device names are not displayed using this command, at the host prompt, type the following command and then press Enter:
   > ```
   > insf
   > ```
8. Unmount the base logical drive.
9. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.
   – Re-create a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:
     ```
     recreateFlashCopy logical drive
     ```
10. Remount the base logical drive (to its original host).

    > **Important:** If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).
11. Complete the following steps to import the FlashCopy logical drives into the Logical Logical Drive Manager:

a. Create a new directory for the new array as shown in the following example:

```
mkdir /dev/vg02
```

b. Create a group node for the new array as shown in the following example:

```
mknod /dev/vg02/group c 64 -0x020000
```

c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter:

```
vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2 ...
```

Consider the following example:

```
vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
```

> **Note**: The */dev/dsk* device files must be verified to be the FlashCopy logical drive and to exist, using the SMdevices utility or the HP-UX **ioscan** utility.

A warning is displayed indicating that a backup of the array being imported may not exist on the host. This message is only a warning and can be ignored. The import continues and completes successfully.

The backup for this array is created when it is exported later.

d. Activate the new array as shown in the following example:

```
vgchange -a y /dev/vg02
```

12. If a filesystem existed on the base logical drive, then it also exists on the FlashCopy logical drive. However, before the FlashCopy logical drive can be mounted, run a filesystem check on it. A filesystem check is performed to ensure the filesystem is consistent, for example:

```
fsck /dev/vg02/lvol01
```

13. Mount the FlashCopy logical drive to its intended host.
14. Use the FlashCopy logical drive in conjunction with your backup application, for speculative testing, or with another application.
15. Unmount the FlashCopy logical drive.
16. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

> **Important:** All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the base logical drive.

2. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:

   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management

   – Re-create a FlashCopy logical drive using the Script Editor or a command line shell on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```

3. Remount the base logical drive (to its original host).

> **Important:** If I/O activity to the base logical drive was stopped or a data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

4. Complete the following steps to import the FlashCopy logical drives into the Logical Logical Drive Manager:

   a. Create a new directory for the new array as shown in the following example:

   ```
   mkdir /dev/vg02
   ```

   b. Create a group node for the new array as shown in the following example:

   ```
   mknod /dev/vg02/group c 64 -0x020000
   ```

   c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter:

   ```
   vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2 ...
   ```

   Consider the following example:

   ```
   vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
   ```

   > **Note**: The /dev/dsk device files must be verified to be the FlashCopy logical drive and to exist, using the SMdevices utility or the HP-UX **ioscan** utility.

   A warning is displayed indicating that a backup of the array being imported may not exist on the host. This message is only a warning and can be ignored. The import continues and completes successfully.

   The backup for this array is created when it is exported later.

   d. Activate the new array as shown in the following example:

   ```
   vgchange -a y /dev/vg02
   ```

5. If a filesystem existed on the base logical drive, then it also exists on the FlashCopy logical drive. However, before the FlashCopy logical drive can be mounted, run a filesystem check on it.

   A filesystem check is performed to ensure the filesystem is consistent, for example:

   ```
   fsck /dev/vg02/lvol01
   ```

6. Mount the FlashCopy logical drive to its intended host.

7. Use the FlashCopy logical drive in conjunction with your backup application (or other application).

8. Unmount the FlashCopy logical drive.

9. If the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

NetWare

This section explains the specific procedures that are used in a Netware environment.

Process overview

Figure A-7 shows the process flow.

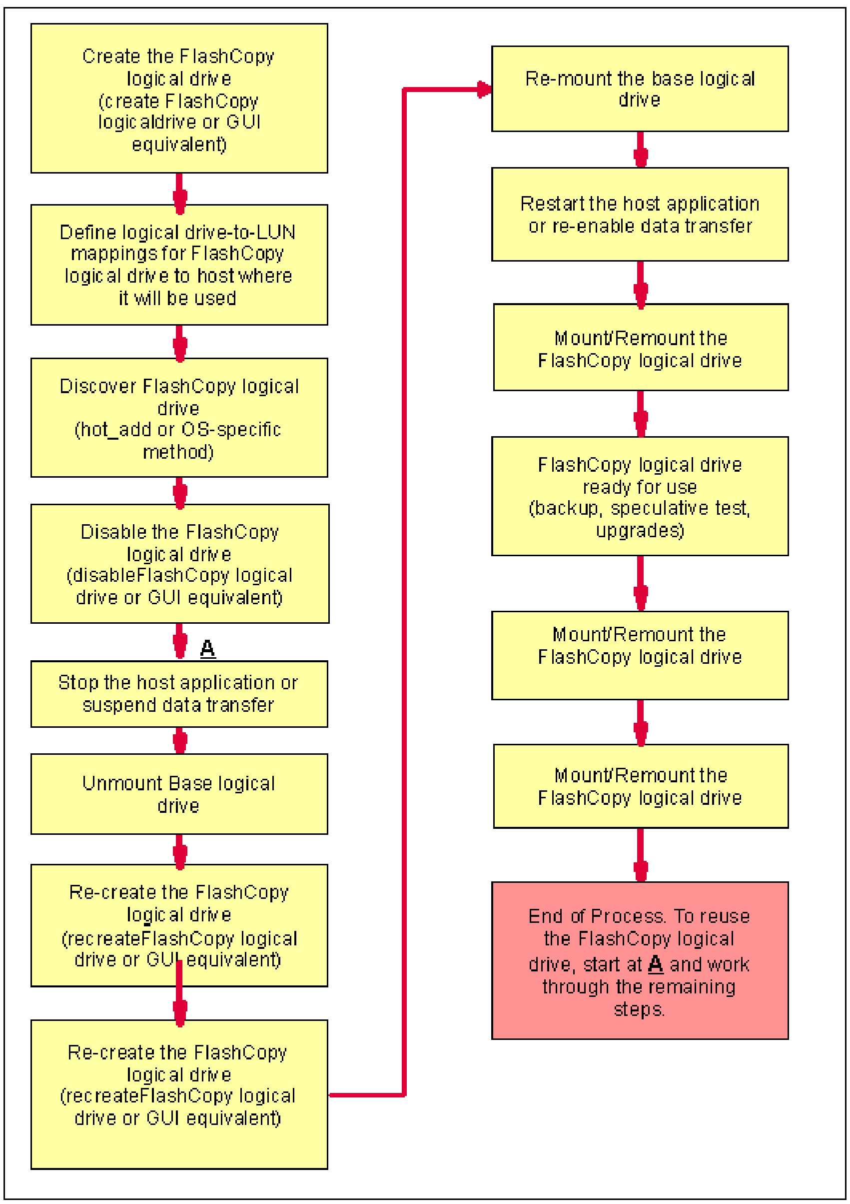

Figure A-7 Creating FlashCopy logical drives on NetWare

Additional instructions for NetWare

Use the following procedure when creating FlashCopy logical drives on a host running NetWare 5.1 (or higher). Failure to complete the steps listed may result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives may be reused (for frequent or nightly backups) or created for one-time usage (speculative change or upgrade testing).

Restriction: When attempting to create FlashCopy logical drives on NetWare, FlashCopy logical drives on NetWare may not be mapped on the same host as the base logical drives.

Creating a FlashCopy logical drive

To create a FlashCopy logical drive, follow these steps:

1. Start the Storage Management software. The Enterprise Management window opens.
2. Launch an Subsystem Management window using one of the following methods:
   - Select the storage subsystem in either the Device Tree View or the Device Table. Then click the **Manage Device** toolbar button or click **Tools-> Manage Device**.
   - Right-click the storage subsystem in the Device Tree View or Device Table and select **Manage Device**.
   - Double-click a storage subsystem node in the Device Table.
   - Select the storage subsystem in either the Device Tree View or the Device Table and press Enter.

   The Subsystem Management window opens in a separate window.
3. Within the Logical pane of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive. Using either the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   create FlashCopylogicaldrive
   ```
4. Press Ctrl+Esc and select **Server Console**. The server console prompt appears.
5. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility is run to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.

   Important: FlashCopy logical drives on NetWare may not be mapped on the same host as the base logical drive.
6. Complete the following steps to mount the FlashCopy logical drive to its intended host:
   a. At the host prompt, type the following command and press Enter:

   ```
   nwconfig
   ```

   The Configuration Options menu appears.

   b. From the Configuration Options menu, click **Standard Disk Options-> NetWare Logical Drive Options**.

   c. Select the logical drive you want to mount and press Enter. The Logical Drive Information window opens.

d. Using the arrow keys, click in the **Status** field. Depending on the status of the logical drive, this field displays *Mounted*, *Not Mounted*, or *New*.
e. Press Enter to display a menu of available actions.
f. Select **Mount** and press Enter. NetWare mounts the selected logical drive.

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

7. Complete the following steps to dismount the base logical drive:
   a. From the Configuration Options menu, click **Standard Disk Options-> NetWare Logical Drive Options**.
   b. Select the logical drive you want to dismount and press Enter. The Logical Drive Information window opens.
   c. Using the arrow keys, select the **Status** field. Depending on the status of the logical drive, this field displays *Mounted*, *Not Mounted*, or *New*.
   d. Press Enter to display a menu of available actions.
   e. Select **Dismount** and press Enter. NetWare dismounts the selected logical drive.
8. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.
   – Re-create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```
9. Return to the server console and remount the base logical drive.
10. Exit all applications and reboot the host operating system.

Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

11. Use the FlashCopy logical drive in conjunction with your backup application (reuse of FlashCopy logical drive) or for speculative change and upgrade testing (one-time usage).
12. Once the FlashCopy logical drive is no longer required, disable or delete the FlashCopy logical drive.

    If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives

Typically, once a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is required. To create a new point-in-time image of the same base logical drive, complete the following steps:

Important: All I/O activity to the base logical drive should be stopped at this point (or data transfer suspended). This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Complete the following steps to dismount the base logical drive:
   a. From the Configuration Options menu, click **Standard Disk Options-> NetWare Logical Drive Options**.
   b. Select the logical drive you want to dismount and press Enter. The Logical Drive Information window opens.
   c. Using the arrow keys, select the **Status** field. Depending on the status of the logical drive, this field displays *Mounted*, *Not Mounted*, or *New*.
   d. Press Enter to display a menu of available actions.
   e. Select **Dismount** and press Enter. NetWare dismounts the selected logical
2. In the Storage Management software, re-create the FlashCopy logical drive using one of the following methods:
   – Click **Logical Drive-> FlashCopy-> Re-create** in the Subsystem Management window.
   – Re-create a FlashCopy logical drive using the Script Editor or a CLI on the host operating system. Type the following command and press Enter:

   ```
   recreateFlashCopy logical drive
   ```
3. Press Ctrl+Esc and select **Server Console**. The server console prompt is displayed.
4. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or reboot the host where the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

   Once logical drives are created and logical drive-to-LUN mappings are defined, the hot_add utility runs to ensure that the operating system is aware of the newly created logical drives, without having to reboot the host.
5. Remount the base logical drive.
6. Exit all applications and reboot the host operating system.

   Important: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive at this time (or re-enable data transfer).

7. Use the FlashCopy logical drive in conjunction with your backup application (or another application).
8. Once the FlashCopy logical drive is no longer required, disable the FlashCopy logical drive.

   If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that may occur if the FlashCopy logical drive remains available.

B

Critical event descriptions

This appendix provides more information about events with a critical priority in the event log. When a critical event occurs, it is logged in the event log. It is also sent to any e-mail and SNMP destinations that may have been configured using the Edit-> Alert Destinations menu option in the Enterprise Management window.

Critical event type and sense	Description and action
Event 1001 - Channel failed	**Description:** The controller failed a channel and will not access drives on this channel any more. The FRU Group Qualifier (byte 26) in the sense data will indicate the 1-relative channel number of the failed channel. This condition is typically caused by a drive ignoring SCSI protocol on one of the controller's destination channels. The controller typically fails a channel if it issued a reset on a channel, and it continued to see drives ignore the SCSI Bus Reset on this channel. **Action:** Select the **Recovery Guru** to obtain the “Failed Drive SCSI Channel” recovery procedure. Contact IBM Technical Support to complete this procedure.
Event 1010 - Impending drive failure (PFA) detected	**Description:** A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive may fail within 24 hours. **Action:** Select the **Recovery Guru** to obtain the “Impending Drive Failure” recovery procedure.
Event 1015 - Incorrect mode parameters set on drive	**Description:** The controller was unable to query the drive for its current critical mode page settings or was unable to change these to the correct setting. This indicates the Qerr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data. **Action:** The controller has not failed yet. Contact IBM Technical Support for instructions on recovering from this failure.

Event 1207 - Fibre Channel link errors - threshold exceeded	**Description:** Invalid characters have been detected in the Fibre Channel signal. Possible causes for the error are a degraded laser in a Gigabit Interface Converter (GBIC), Small Form-factor Pluggable (SFP) transceiver, or Media Interface Adapter (MIA); damaged or faulty Fibre Channel cables; or poor cable connections between components on the loop. **Action:** In the main Subsystem Management window, click **Help-> Recovery Procedures.** Select **Fibre Channel Link Errors - Threshold Exceeded** for more information about recovering from this failure.
Event 150E - Controller loopback diagnostics failed	**Description:** The controller cannot initialize the drive-side Fibre Channel loops. Diagnostics have been run to identify a controller problem, and the controller has been placed offline. This event only occurs on certain controller models. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 150F - Channel miswire	**Description:** Two or more drive channels are connected to the same Fibre Channel loop. This can cause the storage subsystem to behave unpredictably. **Action:** Select the **Recovery Guru** to obtain the “Channel Miswire” recovery procedure.
Event 1510 - ESM canister miswire	**Description:** Two ESM canisters in the same drive enclosure are connected to the same Fibre Channel loop. A level of redundancy has been lost, and the I/O performance for this drive enclosure is reduced. **Action:** Select the **Recovery Guru** to obtain the "ESM Canister Miswire" recovery procedure.
Event 200A - Data/parity mismatch detected on logical drive	**Description:** A media scan operation has detected inconsistencies between a portion of the logical drive's data blocks and associated parity blocks. User data in this portion of the logical drive may have been lost. **Action:** Select an application-specific tool (if available) to verify the correctness of the logical drive's data. If no such tool is available, or if problems with the user data are reported, the entire logical drive contents should be restored from the most recent backup, if it is critical data.
Event 202E - Read drive error during interrupted write	**Description:** A media error has occurred on a read operation during interrupted write processing. **Action:** Select the **Recovery Guru** to obtain the “Unrecovered Interrupted Write” recovery procedure. Contact IBM Technical Support to complete this procedure.
Event 2109 - Controller cache not enabled - cache sizes do not match	**Description:** The controller will not allow mirroring to be enabled if the alternate controller's cache size is different. **Action:** Ensure that both controllers have the same cache size. Contact IBM Technical Support for instructions on recovering from this failure.
Event 210C - Controller cache battery failed	**Description:** The controller has detected that the battery is either not physically present, is fully discharged, or has reached its expiration date. **Action:** Select the **Recovery Guru** to obtain the “Failed Battery Canister” recovery procedure.

Event 210E - Controller cache memory recovery failed after power cycle or reset	**Description:** Recovery from a Data Cache error was unsuccessful. User data may have been lost. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 2110 - Controller cache memory initialization failed	**Description:** The controller has detected the failure of an internal controller component (RAID Buffer). This failure may have been detected during operation as well as during an on-board diagnostic routine. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 2113 - Controller cache battery nearing expiration	**Description:** The cache battery is within the specified number of weeks before failing. **Action:** Select the **Recovery Guru** to obtain the “Battery Nearing Expiration” recovery procedure.
Event 2229 - Drive failed by controller	**Description:** The controller has failed a drive because of a problem with the drive. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 222D - Drive manually failed	**Description:** The drive was manually failed by a user. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 2247 - Data lost on logical drive during unrecovered interrupted write	**Description:** An error has occurred during interrupted write processing during the start-of-day routine causing the logical drive to transition to the Failed state. **Action:** Select the **Recovery Guru** to obtain the “Unrecovered Interrupted Write” recovery procedure. Contact IBM Technical Support to complete this procedure.
Event 2248 - Drive failed - write failure	**Description:** The drive failed a write command to it. The drive will be marked as failed. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 2249 - Drive capacity less than minimum	**Description:** During drive replacement, the capacity of the drive is not large enough to support all the logical drives that must be reconstructed on it. **Action:** Replace the drive with a larger capacity drive.
Event 224A - Drive has wrong block size	**Description:** The drive's block size does not match that of the other drives in the logical drive. The drive will be marked as failed. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 224B - Drive failed - initialization failure	**Description:** The drive failed either a Format Unit command or a Write operation (issued when a logical drive was initialized). The drive will be marked as failed. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.

Event 224D - Drive failed - no response at start of day	**Description:** The drive failed a Read Capacity or Read command during the start-of-day routine. The controller was unable to read the configuration information stored on it. The drive will be marked as failed. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 224E - Drive failed - initialization/reconstruction failure	**Description:** The previously-failed drive was marked failed because either the drive failed a Format Unit command issued to it, or the reconstruction on the drive failed due to the controller being unable to restore it. For example, an error occurring on another drive required for reconstruction. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 2250 - Logical Drive failure (3F E0)	**Description:** The controller has marked the logical drive failed. User data and redundancy (parity), or redundancy, can no longer be maintained to ensure availability. The most likely cause is the failure of a single drive in non-redundant configurations or a second drive in a configuration protected by one drive. **Action:** Select the **Recovery Guru** to obtain the “Failed Logical Drive - Drive Failure” recovery procedure.
Event 2251 - Drive failed - reconstruction failure	**Description:** A drive failed due to a reconstruction failure during the start-of-day routine. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 2252 - Drive marked offline during interrupted write	**Description:** An error has occurred during interrupted write processing causing the logical drive to be marked as failed. Drives in the array that did not experience the read error will transition to the Offline state and log this error. **Action:** Select the **Recovery Guru** to obtain the “Unrecovered Interrupted Write” recovery procedure. Contact IBM Technical Support to complete this procedure.
Event 2254 - Redundancy (parity) and data mismatch was detected	**Description:** The controller detected inconsistent redundancy (parity)/data during a parity verification. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 2255 - Logical Drive definition incompatible with ALT mode - ALT disabled	**Description:** Auto-LUN Transfer (ALT) works with arrays that have only one logical drive defined. There are currently arrays on the storage subsystem that have more than one logical drive defined; therefore, ALT mode has been disabled. The controller will operate in normal redundant controller mode and if there is a problem, it will transfer all logical drives on an array instead of transferring individual logical drives. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.

Event 2602 - Automatic controller firmware synchronization failed	**Description:** The versions of firmware on the redundant controllers are not the same because the automatic controller firmware synchronization failed. Controllers with an incompatible version of firmware may cause unexpected results. **Action:** Try the firmware download again. If the problem persists, contact IBM Technical Support.
Event 2801 - Storage Subsystem running on UPS battery	**Description:** The Uninterruptible Power Source (UPS) has indicated that AC power is no longer present and the UPS has switched to standby power. While there is no immediate cause for concern, you should save your data frequently, in case the battery is suddenly depleted. **Action:** Select the **Recovery Guru** to obtain the “Lost AC Power” recovery procedure.
Event 2803 - UPS battery - two minutes to failure	**Description:** The UPS has indicated that its standby power source is nearing depletion. **Action:** You should take actions to stop I/O activity to the controller. Normally, the controller will change from a write-back caching mode to a write-through mode.
Event 2804 - UPS battery failed	**Description:** The UPS battery has failed. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 2807 - ESM canister failed	**Description:** An ESM canister has failed. **Action:** Select the **Recovery Guru** to obtain the “Failed ESM Canister” recovery procedure.
Event 2808 - Enclosure ID not unique	**Description:** The controller has determined that there are multiple drive enclosures with the same ID selected. Ensure that each drive enclosure has a unique ID setting. **Action:** Select the **Recovery Guru** to obtain the “Enclosure ID Conflict” recovery procedure.
Event 280A - Controller enclosure component missing	**Description:** A component other than a controller is missing in the controller enclosure (for example, fan, GBIC or SFP transceiver, power supply, or battery). The FRU codes indicate the faulty component. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 280B - Controller enclosure component failed	**Description:** A component other than a controller has failed in the controller enclosure (for example, fan, GBIC or SFP transceiver, power supply, or battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component. **Action:** Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.
Event 280D - Drive enclosure component failed	**Description**: A component other than a drive has failed in the drive enclosure (for example, controller, GBIC, ESM, fan, power supply, battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component. **Action**: Select the **Recovery Guru** to obtain the recovery procedure and follow the instructions to correct the failure.

Event 280E - Standby power source not fully charged	**Description:** The Uninterruptible Power Source has indicated that its standby power source is not at full capacity. **Action**: Select the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.
Event 280F - ESM canister - loss of communication	**Description:** Communication has been lost to one of the dual ESM canisters in a drive enclosure. The drive enclosure has only one I/O path available. **Action:** Select the **Recovery Guru** to obtain the “ESM Canister - Loss of Communication” recovery procedure.
Event 2813 - Mini-hub canister failed	**Description:** Communications with the mini-hub canister has been lost. This may be the result of a mini-hub canister failure, a controller failure, or a failure in an internal backplane communications board. If there is only one mini-hub failure, the storage subsystem is still operational, but a second mini-hub failure could result in the failure of the affected enclosures. **Action:** Select the **Recovery Guru** to obtain the “Failed Mini-hub Canister” recovery procedure.
Event 2815 - GBIC/SFP failed	**Description:** A GBIC or SFP transceiver on either the controller enclosure or the drive enclosure has failed. If there is only one GBIC or SFP failure, the storage subsystem is still operational but a second GBIC or SFP failure could result in the failure of the affected enclosures. **Action:** Select the **Recovery Guru** to obtain the “Failed GBIC/SFP” recovery procedure.
Event 2816 - Enclosure ID conflict - duplicate IDs across drive enclosures	**Description:** Two or more drive enclosures are using the same enclosure identification number. **Action:** Select the **Recovery Guru** to obtain the “Enclosure ID Conflict” recovery procedure.
Event 2818 - Enclosure ID mismatch - duplicate IDs in the same drive enclosure	**Description:** A drive enclosure in the storage subsystem contains ESM canisters with different enclosure identification numbers. **Action:** Select the **Recovery Guru** to obtain the “Enclosure ID Mismatch” recovery procedure.
Event 281B - Nominal temperature exceeded	**Description:** The nominal temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. If the temperature of the enclosure continues to rise, the affected enclosure may automatically shut down. Fix the problem immediately, before it becomes more serious. The automatic shutdown conditions depend on the model of the enclosure. **Action:** Select the **Recovery Guru** to obtain the “Nominal Temperature Exceeded” recovery procedure.
Event 281C- Maximum temperature exceeded	**Description:** The maximum temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. This condition is critical and may cause the enclosure to shut down if you do not fix the problem immediately. The automatic shutdown conditions depend on the model of the enclosure. **Action:** Select the **Recovery Guru** to obtain the “Maximum Temperature Exceeded” recovery procedure.

Event 281D - Temperature sensor removed	**Description:** A fan canister containing a temperature sensor has been removed from the storage subsystem. **Action:** Replace the canister as soon as possible. Select the **Recovery Guru** to obtain the “Failed or Removed Fan Canister” recovery procedure.
Event 281E - ESM canister firmware mismatch	**Description:** A drive enclosure in the storage subsystem contains ESM canisters with different versions of firmware. ESM canisters in the same drive enclosure must have the same version firmware. **Action:** Select the **Recovery Guru** to obtain the “ESM Canister Firmware Version Mismatch” recovery procedure. If you do not have a replacement canister, call IBM Technical Support to perform the firmware download.
Event 2821 - Incompatible mini-hub	**Description:** An incompatible mini-hub canister has been detected in the controller enclosure. **Action:** Select the **Recovery Guru** to obtain the “Incompatible Mini-hub Canister” recovery procedure.
Event 3019 - Logical Drive ownership changed due to failover	**Description:** The multi-path driver software has changed ownership of the logical drives to the other controller, because it could not access the logical drives on the particular path. **Action:** Select the **Recovery Guru** to obtain the “Logical Drive Not on Preferred Path” recovery procedure.
Event 502F - Missing logical drive deleted	**Description:** The storage subsystem has detected that the drives associated with a logical drive are no longer accessible. This can be the result of removing all drives associated with an array or a loss of power to one or more drive enclosures. **Action:** Select the **Recovery Guru** to obtain the “Missing Logical Drive” recovery procedure.
Event 5005 - Place controller offline	**Description:** The controller was placed offline. This could be caused by the controller failing a diagnostic test. (The diagnostics are initiated internally by the controller or by using clicking **Controller-> Run Diagnostics** from the menu options.) Or the controller was manually placed Offline using the **Controller-> Place Offline** menu option. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 5602 - This controller's alternate failed - timeout waiting for results	**Description:** This controller initiated diagnostics on the alternate controller but did not receive a reply indicating that the diagnostics completed. The alternate controller in this pair has been placed offline. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 560B - ctlrDiag task cannot obtain Mode Select lock	**Description:** This controller was attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.

Event 560C - ctlrDiag task on controller's alternate cannot obtain Mode Select lock	**Description:** The alternate controller in this pair was attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 560D - Diagnostics read test failed on controller	**Description:** While running diagnostics, the controller has detected that the information received does not match the expected return for the test. This could indicate that I/O is not completing or that there is a mismatch in the data being read. The controller was placed offline as result of this failure. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 560E - This controller's alternate failed diagnostics read test	**Description:** While running diagnostics, this controller's alternate detected that the information received does not match the expected return for the test. This could indicate that I/O is not completing or that there is a mismatch in the data being read. The alternate controller in this pair was placed offline as result of this failure. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 560F - Diagnostics write test failed on controller	**Description:** While running diagnostics, the controller was unable to write data to the test area. This could indicate that I/O is not completing or that there is a mismatch in the data being written. The controller was placed offline as result of this failure. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 5610 - This controller's alternate failed diagnostics write test	**Description:** While running diagnostics, this controller's alternate was unable to write data to the test area. This could indicate that I/O is not completing or that there is a mismatch in the data being written. The alternate controller in this pair was placed offline as result of this failure. **Action:** Select the **Recovery Guru** to obtain the “Offline Controller” recovery procedure. Use this procedure to replace the controller.
Event 5616 - Diagnostics rejected - configuration error on controller	**Description:** This controller was attempting to run diagnostics and could not create the test area necessary for the completion of the tests. The diagnostics were canceled. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.
Event 5617 - Diagnostics rejected - configuration error on controller's alternate	**Description:** This controller's alternate was attempting to run diagnostics and could not create the test area necessary for the completion of the tests. The diagnostics were canceled. **Action:** Contact IBM Technical Support for instructions on recovering from this failure.

Event 6101 - Internal configuration database full	**Description:** Because the amount of data required to store certain configuration data for the maximum number of logical drives has been underestimated, two types of data may have caused the internal configuration database to become full: - FlashCopy logical drive configuration data - Remote Volume Mirror configuration data **Action:** Recovery options include deleting one or more FlashCopy logical drives from your storage subsystem, and/or removing one or more mirror relationships.
Event 6200 - FlashCopy repository logical drive threshold exceeded	**Description:** The FlashCopy's FlashCopy repository logical drive capacity has exceeded a warning threshold level. If the FlashCopy repository logical drive's capacity becomes full, its associated FlashCopy logical drive can fail. If you do not resolve this problem, this is the last warning you will receive before the FlashCopy's FlashCopy repository logical drive becomes full. **Action:** Select the **Recovery Guru** to obtain the “FlashCopy Repository Logical Drive Threshold Exceeded” recovery procedure and follow the instructions to recover from this failure.
Event 6201 - FlashCopy repository logical drive full	**Description:** All of the FlashCopy repository logical drive's available capacity has been used. The FlashCopy repository logical drive's failure policy determines what happens when the FlashCopy's FlashCopy repository logical drive becomes full. The failure policy can be set to either fail the FlashCopy logical drive (default setting) or fail incoming I/Os to the FlashCopy's base logical drive. **Action:** Select the **Recovery Guru** to obtain the “FlashCopy Repository Logical Drive Capacity - Full” recovery procedure and follow the instructions to recover from this failure.
Event 6202 - Failed FlashCopy logical drive	**Description:** Either the FlashCopy logical drive's associated FlashCopy repository logical drive is full or its associated base or FlashCopy repository logical drives failed due to one or more drive failures on their respective arrays. **Action:** Select the **Recovery Guru** to obtain the “Failed FlashCopy Logical Drive” recovery procedure and follow the instructions to recover from this failure.
Event 6400 - Dual primary logical drive	**Description:** Both logical drives have been promoted to Primary after a forced role reversal. This could be reported at controller reset, or when a cable from an array to a Fibre Channel switch is reinserted after having been removed, and the other logical drive has been promoted to primary. **Action:** Select the **Recovery Guru** to obtain the “Dual Primary Logical Drive Conflict” recovery procedure, and follow the instructions to recover from this failure.
Event 6401 - Dual secondary logical drive	**Description:** Both logical drives in the Remote Volume Mirror have been demoted to Secondary after a forced role reversal. This could be reported at controller reset, or when a cable from an array to a Fibre Channel switch is reinserted after having been removed, and the other logical drive has been demoted to secondary. **Action:** Select the **Recovery Guru** to obtain the “Dual Secondary Logical Drive Conflict” recovery procedure, and follow the instructions to recover from this failure.

Event 6402 - Mirror data unsynchronized	**Description:** This may occur because of I/O errors, but there should be other events associated with it. One of them would be the root cause, that would contain the sense data. A Needs Attention icon appears on both the primary and secondary storage subsystems of the Remote Volume Mirror. **Actions:** Select the **Recovery Guru** to obtain the “Mirror Data Unsynchronized” recovery procedure, and follow the instructions to recover from this failure.
Event 6503 - Remote logical drive link down	**Description:** This event is triggered when either a cable between one array and its peer has been disconnected, the Fibre Channel switch has failed, or the peer array has reset. This error could result in the Mirror Data Unsynchronized event. The affected remote logical drive will show an Unresponsive icon, and this state will be selected in the logical drive's tooltip. **Action:** Select the **Recovery Guru** to obtain the “Mirror Communication Error - Unable to Contact Logical Drive” recovery procedure, and follow the instructions to recover from this failure.
Event 6505 - WWN change failed	**Description:** Mirroring causes a WWN change to be communicated between arrays. Failure of a WWN change is caused by non-I/O communication errors between one array, on which the WWN has changed, and a peer array. (The array WWN is the unique name used to locate an array on a fibre network. When both controllers in an array are replaced, the array WWN changes). The affected remote logical drive will show an Unresponsive icon, and this state will be selected in the logical drive's tooltip. **Action:** Select the **Recovery Guru** to obtain the “Unable to Update Remote Mirror” recovery procedure, and follow the instructions to recover from this failure. The only solution to this problem is to delete the Remote Volume Mirror and then to establish another one.

C

Command line reference

This appendix describes the command line interface (CLI) commands and their various parameters. For full descriptions and example uses, refer to the Enterprise Management window online help.

Command line interface parameters

To use the command line interface, follow these steps:

1. Go to the command line shell of your operating system. At the command prompt, enter SMcli, followed by either the controller name, host-agent name, world-wide name (WWN), or user-supplied name of the specific storage subsystems.
2. Enter one or more commands or the name of a script file as shown here:
   – One or more commands:
   ```
   -c "<command>;[<command2>;...]"
   ```
   – Script file:
   ```
   -f <scriptfile>
   ```
3. Specify the output file [-o <outputfile>], password [-p <password>, or [-e] to execute the script only. These arguments are optional.

The command line interface supports the command line parameters in the following table.

Command line parameter	Lets you...
<IP address> or <hostname>	Specify an IP address (xx.xx.xx.xx) or hostname (of host-agent or controller) of a storage subsystem managed through the host-agent or directly- managed method.
-a	Add an SNMP trap destination or e-mail alert destination. To add an SNMP trap destination, enter: -a trap:Community,HOST Here *COMMUNITY* is the SNMP Community Name set in the NMS configuration file by a Network Administrator. The default is *public*. *HOST* is the IP address or the host name of a station running an SNMP service. At a minimum, this is the Network Management station. **Important:** There is no space after the colon (;) or the comma (,). To add an e-mail alert destination, enter: -a email:MAILADDRESS Here *MAILADDRESS* is the fully qualified e-mail address to which the alert message should be sent.
-A	Specify a storage subsystem to add to the management domain. Specify an IP address (xx.xx.xx.xx) for each controller in the storage subsystem. **Important:** If you specify one IP address, the storage subsystem will be partially managed. If no IP address is specified, an automatic discovery is performed of storage subsystems attached to the local subnet.

-c	Specify the list of commands to be performed on the specified storage subsystem. **Important:** Note the following usage requirements: ▸ You cannot place multiple -c parameters on the same command line. However, you can include multiple commands after the -c parameter. ▸ Each command must be terminated with a semicolon (;). **Windows:** The entire command string must be enclosed in double quotes ("). Each command must be terminated with a semicolon (;). **UNIX:** The entire command string must be enclosed in single quotes ('). Each command must be terminated with a semicolon (;). **Note:** Any errors encountered when executing the list of commands, by default, cause the execution to stop. Use the `on error continue;` command first in the list of commands to override this behavior.
-d	Display the contents of the configuration file in the following format: <storagearrayname> <hostname> <hostname> The configuration file lists all known storage subsystems that are currently configured in the Enterprise Management window.
-e	Execute the commands only, without performing a syntax check first.
-f	Specify the name of a file containing script engine commands to be performed on the specified storage subsystem. Use the -f parameter in place of the -c parameter. **Note:** Any errors encountered when executing the list of commands, by default, cause the execution to stop. Use the `on error continue;` command in the script file to override this behavior.
-F	Specify the e-mail address that sends the alerts.
-i	When used with the -d parameter, display the contents of the configuration file in the following format: <storagearrayname> <IP address> <IP address>.
-m	Specify the IP address or host name of the mail/SNMP server that will send the alerts.
-n	Specify the storage subsystem name on which you want to perform the script commands. This name is optional when a <hostname or IP address> is used. However, if you are managing the storage subsystem using the host-agent management method, you must use the -n option if more than one storage subsystem is connected to the host at the specified address. This name is required when the <hostname or IP address> is not used. However, the storage subsystem name must be configured for use in the Enterprise Management window and must not be a duplicate of any other configured storage subsystem name.
-o	Specify a file name for all output text from the script engine. If this parameter is not used, the output goes to stdout.
-p	Specify the password for the storage subsystem on which you want to perform a command script. A password is not necessary if: ▸ A password has not been set on the storage subsystem. ▸ The password is specified with the use password command in the script file with the -f parameter. ▸ You specify the password with the use password command using the -c parameter.
-s	Displays the alert settings for the storage subsystems currently configured in the Enterprise Management window.

-w	Specify the storage subsystem, using its world-wide name (WWN), on which you want to perform the script commands. **Note**: The WWN is optional when a <hostname> is used or if the -n option is used to identify the storage subsystem with its <storagearrayname>. Use this option INSTEAD of the -n option.
-W	Display the WWNs of all known storage subsystems currently configured in the Enterprise Management window. **Note:** The -W is capitalized (it is lowercase when you use the option to specify the WWN) and is followed by the -d option.

-x	Delete an SNMP trap destination or e-mail alert destination. To delete an SNMP trap destination, enter:
	`-x trap:Community, HOST`
	
	Here *COMMUNITY* is the SNMP Community Name, and *HOST* is the IP address or the host name of a station running an SNMP service.
	
	To delete an e-mail alert destination, enter:
	`-x email:MAILADDRESS`
	
	Here *MAILADDRESS* is the fully qualified e-mail address to which the alert message should no longer be sent.

-?	Display usage information.

Commands and parameters

Command	Parameter	Value
create logicalDrive		
	*userLabel	
	*array=[array number] OR drive or drives=[drive list] OR driveCount =[number of drives]	
	*array [array number] OR drive or drives <drive list> OR driveCount [number of drives]	
	enforceSoftLimit	True or false
	capacity	
	owner	a OR b
	RAIDLevel	0, 1, 3, or 5
	readAheadMultiplier	
	segmentSize	
	usageHint	fileSystem, dataBase, or multiMedia
	mapping	Default or none
create FlashCopylogicaldrive		
	*baselogicaldrive	
	*repositoryarray=[array number] OR repositoryDriveCount = <number of drives> OR repositoryDrives = <drivelist>	
	userLabel	
	repositoryUserLabel	
	repositoryRAIDLevel	0, 1, 3, or 5
	warningThresholdPercent	
	repositoryPercentofBase	
	repositoryFullPolicy	failBaseWrites or failFlashCopy
create hostGroup		
	*userLabel	
create host		
	*userLabel	
	hostGroup	

Command	Parameter	Value
create hostPort		
	*host	
	*userLabel	
	*identifier	
	*hostType	
create mapping		
	*logicalDrive	
	*logicalUnitNumber	
	host	
	hostGroup	
delete array		
	*[array number]	
delete logicalDrive		
	*[“logical drive label”]	
delete hostGroup		
	*["group label"]	
delete host		
	*["host label"]	
delete hostPort		
	*["port label"]	
delete mapping		
	*logicalDrive	
	host	
	hostGroup	
disableFlashCopy logicalDrive		
	*["logical drive label"]	
download storagesubsystem		
	*file	
	*content	
		firmware
		NVSRAM
		featureKey
	firmwareFile	
	NVSRAMFile	

Command	Parameter	Value
download drive		
	*[tray, slot]	
	*file	
	*content	*firmware
download tray		
	*[trayID]	
	*file	
	*content	*firmware
on error		
	stop	
	continue	
recreateFlashCopy logicalDrive		
	*["logicaldrivelabel"]	
	warningThresholdPercent	
	repositoryFullPolicy	failBaseWrites or failFlashCopy
	userLabel	
reset controller		
	*[controller-spec]	a OR b
reset storagesubsystem		
	*logicalDriveDistribution	
set storagesubsystem		
	batteryInstallDate	true
	cacheBlockSize	
	cacheFlushStart	
	cacheFlushStop	
	mediaScanRate	
	userLabel	
	resetConfiguration	true
	timeOfDay	true
	clearEventLog	true
	password	
	RLSBaseline	currentTime
	defaultHostType	

Command	Parameter	Value
set controller		
	*[controller-spec]	a OR b
	mode	
		active
		passive
	availability	online OR offline
	batteryInstallDate	true
	GlobalNVSRAMByte[nvsram-offset]	
	HostNVSRAMByte [host-type, nvsram-offset]	
	NVSRAMByte [nvsram-offset]	
set array		
	*[array number]	
	owner	a OR b
set logicalDrive		
	*[logical drive label]	
	owner	a OR b
	readCacheEnabled	true OR false
	writeCacheEnabled	true OR false
	mirrorEnabled	true OR false
	cacheWithoutBatteryEnabled	true OR false
	cacheFlushModifier	
	readAheadMultiplier	
	mediaScanEnabled	true OR false
	redundancyCheckEnabled	true OR false
	modificationPriority	lowest, low, medium, high, or highest
	userLabel	
	warningThresholdPercent	
	repositoryFullPolicy	failBaseWrites or failFlashCopy
	role	
		primary OR secondary
		force=<true OR false>
set performanceMonitor		
	*interval	
	*iterations	

Command	Parameter	Value
set drive (OR set drives)		
	*[drive list]	
	hotspare	true OR false
	operationalState	failed OR optimal
show storagesubsystem		
	batteryAge	
	healthStatus	
	profile	
	hostTypeTable	
	defaultHostType	
	logicalDriveDistribution	
	preferredVolumeOwners	
show controller		
	*[controller-spec]	a OR b
	mode	
	GlobalNVSRAMByte[nvsram-offset]	
	HostNVSRAMByte[host-type, nvsram-offset]	
	NVSRAMByte [nvsram-offset]	
show logicalDrives		
	logicalDrive [“logicaldrivelabel”]	
show hostTopology		
show <string>		
start increaseVolCapacity		
	logicalDrive="logical drive label"	
	incrementalCapacity	
	drive or drives=[drive list]	
upload storagesubsystem		
	*file	
	*content	
		RLSCounts
		configuration
		performanceStats
		allEvents
		criticalEvents

Command	Parameter	Value
		stateDump
use password		
	*<”password”>	

D

FlashCopy example scripts

This appendix provides sample scripts that can be used for example in conjunction with backup software. Three scenarios are provided:

- Windows 2000 Basic Disks: FlashCopy mapped to a different host
- Windows 2000 Dynamic Disks: FlashCopy mapped to a different host
- Linux Disks: FlashCopy mapped to the same host

The scripts are in three parts:

- Initial script: Creates the FlashCopy logical drives from scratch (needed only once) (Windows only)
- Disable script: Disables the FlashCopy logical drives but keeps the definition (usually run after the backup process) (Windows and Linux)
- Recreate script: Recreates the FlashCopy logical drives (usually run just before the backup process) (Windows and Linux)

The usage of these scripts is explained in detail in 5.3, “Scenarios using scripts” on page 176.

Windows 2000 basic disks

This script covers the following scenario:

- Host AAE19 is using the logical drive named “FlashSource”. This is the source data drive that it is required for making a FlashCopy.
- The logical drive FlashSource is a basic disk containing one partition using the drive letter E:
- Host AAE18 is using the FlashCopy logical drive named “FlashTarget” for backup purposes.
- The FlashCopy logical drive FlashTarget is mounted under drive letter F:

See 5.3.1, “Windows 2000 basic disks” on page 177, for a detailed description of all tools and command line tools used in these scripts.

Initial flashcopy script

This script generates the initial FlashCopy logical drive including the repository, defines the mapping in storage partitioning, and brings the disk online in the operating system:

```
@echo off
rem
rem This script initiates a first flashcopy
rem

rem --- general Setting ----

rem Path to the SMflashcopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Path to hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Path to script for diskpart.
set DISKPART_SCRIPT="c:\scripts\rescan.dp"

rem Host name of the machine flashcopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
```

```
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem Name of the source logical drive to take a flashcopy of on FAStT.
set SOURCE_DRIVE_NAME=\"FlashSource\"

rem Name of the target logical drive hold the flashcopy of on FAStT.
set TARGET_DRIVE_NAME=\"FlashTarget\"

rem Label to give the new repository logical drive.
set REPOSITORY_NAME=\"FlashRepository-1\"

rem Raid level of the new repository logical drive.
set REPOSITORY_RAID_LEVEL=5

rem Array to create the repository logical drive on.
set REPOSITORY_ARRAY=1

rem LUN number to map the target drive to the host
set LUN_NUMBER=1

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Create Flashcopy Script with Repository on existing array.
set CREATE_FLASHCOPY_SCRIPT="on error stop; create FlashCopyLogicalDrive
baseLogicalDrive=%SOURCE_DRIVE_NAME% repositoryArray=%REPOSITORY_ARRAY%
userLabel=%TARGET_DRIVE_NAME% repositoryUserLabel=%REPOSITORY_NAME%
repositoryRaidLevel=%REPOSITORY_RAID_LEVEL%;"

rem Create mapping script to map target drive to host %HOST_NAME%.
set MAP_FLASHCOPY_SCRIPT="on error stop; create mapping logicalDrive=%TARGET_DRIVE_NAME%
logicalUnitNumber=%LUN_NUMBER% host=%HOST_NAME%;"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"

echo Creating new flashcopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -c %CREATE_FLASHCOPY_SCRIPT%
if errorlevel 1 goto create_flashcopy_fail_exit

echo Mapping FlashCoy drive to %HOST_NAME%.
call %SM_CLI% -n %FAStT_NAME% -c %MAP_FLASHCOPY_SCRIPT%
if errorlevel 1 goto map_flashcopy_fail_exit

echo Adding Flashcopy Target to operating system.
%HOTADD%

echo Bringing FlashCopy Target drive online with drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit
```

```
echo The initial flashcopy creation is complete.
goto done

rem --- Error Handling ---

:create_flashcopy_fail_exit
echo Creating the flashcopy logical drive failed.
goto done

:map_flashcopy_fail_exit
echo Mapping the flashcopy logical drive failed.
goto done

:diskpart_fail_exit
echo Diskpart failed.
goto done

:done

rem --- Post CleanUp ---

set REMOTE_SNAPASSIST=
set SM_CLI=
set RCMD_TOOL=
set NETSVC_TOOL=
set HOTADD=
set DISKPART_TOOL=
set DISKPART_SCRIPT=
set HOST_NAME=
set REMOTE_HOST_NAME=
set FAStT_NAME=
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set SOURCE_DRIVE_NAME=
set TARGET_DRIVE_NAME=
set REPOSITORY_NAME=
set REPOSITORY_RAID_LEVEL=
set REPOSITORY_ARRAY=
set LUN_NUMBER=
set REMOTE_SCRIPT=
set CREATE_FLASHCOPY_SCRIPT=
set MAP_FLASHCOPY_SCRIPT=
```

Diskpart script

This is the script `DISKPART_SCRIPT` called *rescan.dp* used for the diskpart tool in the initial flashcopy script:

```
rescan
```

Disable flashcopy script

This scripts unmounts, unmaps, and disables the FlashCopy logical drive to avoid error messages in Windows 2000:

```
@echo off
rem
rem This script disables a flashcopy logical drive
rem

rem --- general Setting ----
```

```
rem Path to the SMflashcopyassist utility.
set SNAPASSIST="c:\Program Files\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMflashcopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem Host name of the machine flashcopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem The drive letter of the flashcopy target drive.
set TARGET_FLASHCOPY_DRIVE_LETTER="F:"

rem Name of the target logical drive to hold the flashcopy of on FAStT.
set TARGET_DRIVE_NAME=\"FlashTarget\"

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Disable flashcopy script.
set DISABLE_FLASHCOPY_SCRIPT="on error stop; disableFlashCopy logicalDrive
[%TARGET_DRIVE_NAME%];"

rem Unmapping script to map target drive to host %HOST_NAME%.
set MAP_FLASHCOPY_SCRIPT="on error stop; delete mapping logicalDrive=%TARGET_DRIVE_NAME%
host=%HOST_NAME% ;"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"
```

```
echo Flushing target drive %TARGET_FLASHCOPY_DRIVE_LETTER%.
%SNAPASSIST% -f %TARGET_FLASHCOPY_DRIVE_LETTER%
if errorlevel 1 goto flush_target_fail_exit

echo Unmount target drive %TARGET_FLASHCOPY_DRIVE_LETTER%.
mountvol %TARGET_FLASHCOPY_DRIVE_LETTER% /d
if errorlevel 1 goto unmount_fail_exit

echo Unmapping flashcopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -c %MAP_FLASHCOPY_SCRIPT%
if errorlevel 1 goto mapping_fail_exit

echo Disabling flashcopy logical drive.
call %SM_CLI% -n %FAStT_NAME% -c %DISABLE_FLASHCOPY_SCRIPT%
if errorlevel 1 goto disable_flashcopy_fail_exit

echo Bus rescan to remove stale drives.
%HOTADD%

echo Disabling flashcopy is complete.
goto done

rem --- Error Handling ----

:flush_target_fail_exit
echo Flushing file systems for target logical drive containing drive
%TARGET_FLASHCOPY_DRIVE_LETTER% failed
goto done

:unmount_fail_exit
echo Unmounting the flashcopy logical drive failed
goto done

:disable_flashcopy_fail_exit
echo Disabling the flashcopy logical drive failed
goto done


:mapping_fail_exit
echo Unmapping the flashcopy logical drive failed
goto done

:done

rem --- Post CleanUp ---

set SNAPASSIST=
set REMOTE_SNAPASSIST=
set SM_CLI=
set HOTADD=
set RCMD_TOOL=
set NETSVC_TOOL=
set REMOTE_HOST_NAME=
set HOST_NAME=
set FAStT_NAME=
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set TARGET_FLASHCOPY_DRIVE_LETTER=
set TARGET_DRIVE_NAME=
set REMOTE_SCRIPT=
```

```
set DISABLE_FLASHCOPY_SCRIPT=
set MAP_FLASHCOPY_SCRIPT=
```

Recreate flashcopy script

This scripts recreates the FlashCopy logical drive, maps it back to the host, and brings it online with the old drive letter:

```
@echo off
rem
rem This script recreates a flashcopy logical drive
rem

rem --- general Setting ----

rem Path to the SMflashcopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path tho the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem Host name of the machine flashcopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem The drive letter of the flashcopy target drive.
set TARGET_FLASHCOPY_DRIVE_LETTER=F:

rem Name of the target logical drive to hold the flashcopy of on FAStT.
set TARGET_DRIVE_NAME=\"FlashTarget\"

rem Volume ID of the FlashCopy target drive.
set TARGET_FLASHCOPY_DRIVE_VOLUME_ID=\\?\Volume{062d708e-732b-11d6-ab27-0003477b412d}\

rem LUN number to map the target drive to the host.
set LUN_NUMBER=1

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Recreate flashcopy script.
```

```
set RECREATE_FLASHCOPY_SCRIPT="on error stop; recreateFlashCopy logicalDrive
[%TARGET_DRIVE_NAME%];"

rem Create mapping script to map target drive to host %HOST_NAME%.
set MAP_FLASHCOPY_SCRIPT="on error stop; create mapping logicalDrive=%TARGET_DRIVE_NAME%
logicalUnitNumber=%LUN_NUMBER% host=%HOST_NAME%;"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"

echo Recreating flashcopy logical drive %TARGET_FLASHCOPY_DRIVE%.
call %SM_CLI% -n %FAStT_NAME% -c %RECREATE_FLASHCOPY_SCRIPT%
if errorlevel 1 goto recreate_flashcopy_fail_exit

echo Mapping flashopy logical drive %TARGET_FLASHCOPY_DRIVE%.
call %SM_CLI% -n %FAStT_NAME% -c %MAP_FLASHCOPY_SCRIPT%
if errorlevel 1 goto mapping_fail_exit

echo Bus rescan using hot_add.
%HOTADD%

echo Wait 5 sconds.
sleep 5

echo Mount the volume under old drive letter %TARGET_FLASHCOPY_DRIVE_LETTER%.
mountvol %TARGET_FLASHCOPY_DRIVE_LETTER% %TARGET_FLASHCOPY_DRIVE_VOLUME_ID%

echo Recreating flashcopy is complete.
goto done

rem --- Error Handling ---

:recreate_flashcopy_fail_exit
echo Recreating the flashcopy logical drive failed
goto done

:mapping_fail_exit
echo Recreating the flashcopy logical drive failed
goto done

:done

rem --- Post CleanUp ---

set REMOTE_SNAPASSIST=
set SM_CLI=
set HOTADD=
set RCMD_TOOL=
set NETSVC_TOOL=
set REMOTE_HOST_NAME=
set HOST_NAME=
set FAStT_NAME=
```

```
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set TARGET_FLASHCOPY_DRIVE_LETTER=
set TARGET_DRIVE_NAME=
set TARGET_FLASHCOPY_DRIVE_VOLUME_ID=
set LUN_NUMBER=
set REMOTE_SCRIPT=
set REMOTE_SCRIPT=
set RECREATE_FLASHCOPY_SCRIPT=
set MAP_FLASHCOPY_SCRIPT=
```

Windows 2000 dynamic disks

This script covers the following scenario:

- Host AAE19 is using three logical drives – FlashSource-1, FlashSource-2 and FlashSource-3 – as dynamic disks. They are configured as a spanned volume with the drive letter E:, which is the source for the FlashCopy logical drive.
- Host AAE18 is using the FlashCopy logical drives – FlashTarget-1, FlashTarget-2, and FlashTarget-3 – as imported dynamic disks for backup purposes. The imported volume set will be mounted under drive letter F:

See 5.3.2, "Windows 2000 dynamic disks" on page 183, for a detailed description of all tools and command line tools used in these scripts.

Initial flashcopy script

This script generates the three initial FlashCopy logical drives including the repository, defines the mapping in storage partitioning, and imports the three dynamic disks as a spanned volume and assigns a drive letter:

```
@echo off
rem
rem This script initiates a first flashcopy
rem

rem --- general Setting ----

rem Path to the SMflashcopyassist utilityon remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Path to the script for diskpart.
set DISKPART_SCRIPT="c:\scripts\initial.dp"
```

```
rem Host name of the machine flashcopy target to map to on FAStT.
set HOST_NAME="AAE18"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Create Flashcopy Script with Repository on existing array.
set INITIAL_FLASHCOPY_SCRIPT="c:\temp\scripts\initial.cli"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"

echo Creating and mapping new flashcopy logical drives.
call %SM_CLI% -n %FAStT_NAME% -f %INITIAL_FLASHCOPY_SCRIPT%
if errorlevel 1 goto create_flashcopy_fail_exit

echo Adding Flashcopy Target to operating system.
%HOTADD%

echo Bringing FlashCopy Target drive online with drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit

echo The initial flashcopy creation is complete.
goto done

rem --- Error Handling ---

:create_flashcopy_fail_exit
echo Creating the flashcopy logical drive failed.
goto done

:diskpart_fail_exit
echo Diskpart failed.
goto done

:done

rem --- Post CleanUp ---

set REMOTE_SNAPASSIST=
```

```
set SM_CLI=
set RCMD_TOOL=
set NETSVC_TOOL=
set HOTADD=
set DISKPART_TOOL=
set DISKPART_SCRIPT=
set HOST_NAME=
set REMOTE_HOST_NAME=
set FAStT_NAME=
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set REMOTE_SCRIPT=
```

Diskpart script

This is the script DISKPART_SCRIPT called *initial.dp* used for the diskpart tool in the initial flashcopy script:

```
rescan
select disk 2
import
select volume 0
online
assign letter=f
```

SMcli script

This is the script INITIAL_FLASHCOPY_SCRIPT called *initial.cli* used for the Storage Manager command line tool in the initial FlashCopy script:

```
// Initial Creation of three FlashCopy logical Drives and their mappings
on error stop;
// Create the FlashCopy
create FlashCopyLogicalDrive baseLogicalDrive="FlashSource-1" repositoryArray=1
userLabel="FlashTarget-1" repositoryUserLabel="Repository-1" repositoryRaidLevel=5 ;
create FlashCopyLogicalDrive baseLogicalDrive="FlashSource-2" repositoryArray=1
userLabel="FlashTarget-2" repositoryUserLabel="Repository-2" repositoryRaidLevel=5 ;
create FlashCopyLogicalDrive baseLogicalDrive="FlashSource-3" repositoryArray=1
userLabel="FlashTarget-3" repositoryUserLabel="Repository-3" repositoryRaidLevel=5 ;
// Map drives to host system
create mapping logicalDrive="FlashTarget-1" logicalUnitNumber=0 host="AAE18" ;
create mapping logicalDrive="FlashTarget-2" logicalUnitNumber=1 host="AAE18" ;
create mapping logicalDrive="FlashTarget-3" logicalUnitNumber=2 host="AAE18" ;
```

Disable flashcopy script

This script unmounts, unmaps, and disables the FlashCopy logical drives to avoid error messages in Windows 2000:

```
@echo off
rem
rem This script disables a flashcopy logical drive
rem

rem --- general Setting ----

rem Path to the SMflashcopyassist utility.
set SNAPASSIST="c:\Program Files\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMflashcopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
```

```
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool
set RCMD_TOOL="C:\program files\resource kit\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Path to the diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Path to the script for diskpart.
set DISKPART_SCRIPT="c:\scripts\disable.dp"

rem Path to the resacn script for diskpart.
set DISKPART_RESCAN_SCRIPT="c:\temp\scripts\rescan.dp"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem The drive letter of the flashcopy target drive.
set TARGET_FLASHCOPY_DRIVE_LETTER="F:"

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Disable flashcopy script.
set DISABLE_FLASHCOPY_SCRIPT="c:\temp\scripts\disable.cli"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"

echo Flushing target drive %TARGET_FLASHCOPY_DRIVE_LETTER%.
%SNAPASSIST% -f %TARGET_FLASHCOPY_DRIVE_LETTER%
if errorlevel 1 goto flush_target_fail_exit

echo Removing FlashCopy Target drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit

echo Disabling and unmapping flashcopy logical drive.
```

```
call %SM_CLI% -n %FAStT_NAME% -f %DISABLE_FLASHCOPY_SCRIPT%
if errorlevel 1 goto disable_flashcopy_fail_exit

echo Bus rescan to remove stale devices.
%HOTADD%

echo Rescanning disks to remove stale drives and volumes.
%DISKPART_TOOL% -s %DISKPART_RESCAN_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit

echo Disabling flashcopy is complete.
goto done

rem --- Error Handling ----

:flush_target_fail_exit
echo Flushing file systems for target logical drive containing drive
%TARGET_FLASHCOPY_DRIVE_LETTER% failed
goto done

:disable_flashcopy_fail_exit
echo Disabling the flashcopy logical drive failed
goto done

:diskpart_fail_exit
echo Diskpart failed
goto done

:done

rem --- Post CleanUp ---

set SNAPASSIST=
set REMOTE_SNAPASSIST=
set SM_CLI=
set HOTADD=
set RCMD_TOOL=
set NETSVC_TOOL=
set DISKPART_TOOL=
set DISKPART_SCRIPT=
set DISKPART_RESCAN_SCRIPT=
set REMOTE_HOST_NAME=
set FAStT_NAME=
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set TARGET_FLASHCOPY_DRIVE_LETTER=
set REMOTE_SCRIPT=
set DISABLE_FLASHCOPY_SCRIPT=
```

Diskpart Script

This is the script DISKPART_SCRIPT called **disable.dp** used for the diskpart tool in the disable flashcopy script.

```
select volume 0
remove
```

This is the script DISKPART_RESCAN_SCRIPT called **rescan.dp** used for the diskpart tool in the disable flashcopy script.

```
rescan
```

SMcli Script

This is the script DISABLE_FLASHCOPY_SCRIPT named **disable.cli** used for the Storage Manager command line tool in the disable flashcopy script.

```
// Unmap and disable the three FlashCopy logical Drives
on error stop;
```

```
// Unamp drives
delete mapping logicalDrive="FlashTarget-1" host="AAE18" ;
delete mapping logicalDrive="FlashTarget-2" host="AAE18" ;
delete mapping logicalDrive="FlashTarget-3" host="AAE18" ;
// Disable the FlashCopy
disableFlashCopy logicalDrive ["FlashTarget-1"] ;
disableFlashCopy logicalDrive ["FlashTarget-2"] ;
disableFlashCopy logicalDrive ["FlashTarget-3"] ;
Recreate FlashCopy script
This scripts recreates the FlashCopy logical drives, maps them back to the host and
brings the dynamic disks and the spanned volume online with the old drive letter.
@echo off
rem
rem This script recreates a flashcopy logical drive
rem

rem --- general Setting ----

rem Path to the SMflashcopyassist utility on remote system.
set REMOTE_SNAPASSIST="c:\Progra~1\IBM_FAStT\Util\SMflashcopyassist"

rem Path to the SMClient command line script executable.
set SM_CLI="C:\Program Files\IBM_FAStT\client\SMcli"

rem Path to the hot_add utility.
set HOTADD="c:\program files\ibm_fastt\util\hot_add"

rem Path to the RCMD tool.
set RCMD_TOOL="c:\program files\resource kit\scripts\rcmd"

rem Path to Netsvc Tool from the Resource Kit.
set NETSVC_TOOL="c:\program files\resource kit\netsvc"

rem Path to the diskpart tool.
set DISKPART_TOOL="c:\program files\resource kit\diskpart"

rem Pathto the script for diskpart.
set DISKPART_SCRIPT="c:\scripts\recreate.dp"

rem Remote Host name running RCMD service.
set REMOTE_HOST_NAME="\\AAE19"

rem FAStT name of the machine executing this script.
set FAStT_NAME="ITSO"

rem The drive letter of the flashcopy source drive.
set SOURCE_FLASHCOPY_DRIVE_LETTER="E:"

rem --- Script Generation ---

rem Flushing buffers on remote system.
set REMOTE_SCRIPT="%REMOTE_SNAPASSIST% -f %SOURCE_FLASHCOPY_DRIVE_LETTER%"

rem Recreate flashcopy script.
set RECREATE_FLASHCOPY_SCRIPT="c:\temp\scripts\recreate.cli"

rem --- Script Execution ---

echo Starting Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /start %REMOTE_HOST_NAME% "Remote Command Service"
```

```
echo Flushing drive %SOURCE_FLASHCOPY_DRIVE_LETTER% on remote system %REMOTE_HOST_NAME%.
%RCMD_TOOL% %REMOTE_HOST_NAME% %REMOTE_SCRIPT%

echo Stopping Remote Command Service on remote host %REMOTE_HOST_NAME%.
%NETSVC_TOOL% /stop %REMOTE_HOST_NAME% "Remote Command Service"

echo Recreating flashcopy logical drive and remapping.
call %SM_CLI% -n %FAStT_NAME% -f %RECREATE_FLASHCOPY_SCRIPT%
if errorlevel 1 goto recreate_flashcopy_fail_exit

echo Bus rescan.
%HOTADD%

echo Wait 5 seconds.
sleep 5

echo Bringing FlashCopy Target drive online with drive letter.
%DISKPART_TOOL% -s %DISKPART_SCRIPT%
if errorlevel 1 goto diskpart_fail_exit

echo Recreating flashcopy is complete.
goto done

rem --- Error Handling ---

:recreate_flashcopy_fail_exit
echo Recreating the flashcopy logical drive failed
goto done

:diskpart_fail_exit
echo Diskpart failed.
goto done

:done

rem --- Post CleanUp ---

set REMOTE_SNAPASSIST=
set SM_CLI=
set HOTADD=
set RCMD_TOOL=
set NETSVC_TOOL=
set DISKPART_TOOL=
set DISKPART_SCRIPT=
set REMOTE_HOST_NAME=
set FAStT_NAME=
set SOURCE_FLASHCOPY_DRIVE_LETTER=
set REMOTE_SCRIPT=
set RECREATE_FLASHCOPY_SCRIPT=
```

Diskpart script

This is the script DISKPART_SCRIPT, called *recreate.dp*, used for the diskpart tool in the recreate flashcopy script:

```
rescan
select disk 2
online
select disk 3
online
```

```
select disk 4
online
select volume 0
online
assign letter=F
```

SMcli script

This is the script RECREATE_FLASHCOPY_SCRIPT, called *recreate.cli*, used for the Storage Manager command line tool in the recreate flashcopy script:

```
// Recreate and map the three FlashCopy logical Drives
on error stop;
// Create the FlashCopy
recreateFlashCopy logicalDrive ["FlashTarget-1"] ;
recreateFlashCopy logicalDrive ["FlashTarget-2"] ;
recreateFlashCopy logicalDrive ["FlashTarget-3"] ;
// Map drives
create mapping logicalDrive="FlashTarget-1" logicalUnitNumber=0 host="AAE18" ;
create mapping logicalDrive="FlashTarget-2" logicalUnitNumber=1 host="AAE18" ;
create mapping logicalDrive="FlashTarget-3" logicalUnitNumber=2 host="AAE18" ;
```

Linux disks

This script covers the following scenario:

- Host AAE19 is using the logical drive FlashSource with ext2 or ReiserFS as the filesystem. Mounted as /mnt/source, this is the logical drive to FlashCopy.
- The FlashCopy logical drive FlashTarget is mounted to the same host AAE19 again under /mnt/target to be used for backup purposes.

See 5.3.3, “Linux disks” on page 188, for a detailed description of all tools and command line tools used in these scripts.

Disable flashcopy script

This scripts unmounts and disables the FlashCopy logical drive:

```
#!/bin/bash
#
# Disable FlashCopy Logical Drive and unmount it
#
#

SMCLI_TOOL="/opt/IBM_FAStT/client/SMcli"

HOST_NAME="100.100.100.19"
FAStT_NAME=ITSO

TARGET_FLASHCOPY_DRIVE="FlashTarget"
TARGET_PARTITION="/dev/sdc1"

SMCLI_SCRIPT="disableFlashCopy logicalDrive [$TARGET_FLASHCOPY_DRIVE];"

echo Flushing all buffers

sync
sleep 3
sync
```

```
echo Unmounting target drive
umount $TARGET_PARTITION

echo Disabling FlashCopy Target on FAStT
$SMCLI_TOOL -n "``$FAStT_NAME" -e -c "``$SMCLI_SCRIPT"
```

Recreate flashcopy script

This script recreates the FlashCopy logical drive and mounts it back online:

```
#!/bin/bash
#
# Recreate FlashCopy Logical Drive and mount it
#
#

SMCLI_TOOL="/opt/IBM_FAStT/client/SMcli"

HOST_NAME="100.100.100.19"
FAStT_NAME=ITSO

TARGET_FLASHCOPY_DRIVE="FlashTarget"
TARGET_PARTITION="/dev/sdc1"
TARGET_MOUNT_POINT="/mnt/target/"
TARGET_FILE_SYSTEM="ext2"

SMCLI_SCRIPT="recreateFlashCopy logicalDrive [$TARGET_FLASHCOPY_DRIVE];"

echo Flushing all buffers
sync
sleep 3
sync

echo Recreating FlashCopy Target on FAStT
$SMCLI_TOOL -n "``$FAStT_NAME" -e -c "``$SMCLI_SCRIPT"

echo Mounting target drive
mount -t $TARGET_FILE_SYSTEM $TARGET_PARTITION $TARGET_MOUNT_POINT
```

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 346.

- *Designing an IBM Storage Area Network*, SG24-5758
- *Implementing an Open IBM SAN,* SG24-6116
- *IBM SAN Survival Guide Featuring the IBM 2109*, SG24-6127
- *Fibre Array Storage Technology: A FAStT Introduction*, SG24-6246
- *Implementing an Open IBM SAN Featuring the IBM 2109, 3534-1RU, 2103-H07,* SG24-6412

Other resources

These publications are also relevant as further information sources:

- *FAStT Storage Manager online help*
- *Storage Manager Software Installation Guide*, which is shipped on CD with the product. You can also download it from the Web at:

  `http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700`

Referenced Web sites

These Web sites are also relevant as further information sources:

- IBM TotalStorage FAStT Storage Servers home page:

  `http://www.storage.ibm.com/hardsoft/disk/fastt/`
- IBM TotalStorage FAStT700 home page:

  `http://www.storage.ibm.com/hardsoft/products/fast700/fast700.htm`
- IBM TotalStorage FAStT500 home page:

  `http://www.storage.ibm.com/hardsoft/products/fast500/fast500.htm`
- IBM TotalStorage FAStT200 home page:

  `http://www.storage.ibm.com/hardsoft/products/fast200/fast200.htm`
- IBM TotalStorage supported servers page:

  `http://www.storage.ibm.com/hardsoft/disk/fastt/supserver.htm`
- IBM FAStT Storage Manager Version 8.2 download files matrix:

  `http://www.pc.ibm.com/qtechinfo/MIGR-42081.html`
- IBM FAStT MSJ for Windows and NetWare download page:

  `http://www.pc.ibm.com/qtechinfo/MIGR-39194.html`

- IBM FAStT MSJ for Linux download page:

  http://www.pc.ibm.com/qtechinfo/MIGR-42065.html

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

Numerics
2048 logical drive support 7
2109-F08 13
2109-F16 13

A
access logical drive 37, 70
  mapping can't attach to 264
activating RVM 205
activating the Remote Volume Mirror option 204
ADT (Auto Logical Drive Transfer) 56
AIX 47, 268, 290
  creating a FlashCopy logical drive 291
  installing a host adapter driver 79
  installing the FAStT Client 79
  installing the host software 79
  LVM Logical Logical Drives 291
  reusing FlashCopy logical drives 294
AL_PA 3
Alert Destinations 111
alerting 111
array
  creating 100–101
  defragmentation 45
  defragmenting 124
  maintenance 123
Auto Logical Drive Transfer (ADT) 56
Automatic Discovery 86

B
base logical drive 7, 34, 260
basic disk 272
basic management 43
battery 15
battery module 15
battery module LED 19

C
cabling 26–27
cache 3, 14
  global settings 117
  logical drive settings 118
  memory 54
  parameter 57, 252
  settings 117
  size 14
cache block size 252
cache flushing level 252
Cache Hit Percentage 251
cache read ahead 251
CLI 7, 35, 140
  create 62
  delete 62
  disableFlashCopy 62
  download 62
  recreateFlashCopy 63
  reset 63
  set 63
  show 63
  upload 63
CLI (command line interface) 7, 35, 58, 60, 140
cluster 2–3
cluster setup 234
command line interface (CLI) 7, 35, 58, 60, 140
command reference 62
common installation problems 263
component properties 114
configuration
  FAStT Storage Server 85
  initial steps 96
  planning 85
connection loss to secondary drive 199
connectivity 27
controller fan module 15
controller microcode, updating 88, 236
controller owner 193
controller ownership 57, 254
Controller power LED 19
controller runtime diagnostics 260
controller unit 3
controllers
  network setup 129
  resetting 131
Copy Services 1
create command 62
Create FlashCopy Logical Drive Wizard 159
Create Logical Drive Wizard 43
creating a logical drive 100
critical event descriptions 307
current I/O per second 251
Current KB/sec and Maximum KB/sec 251

D
DAE/DVE 44
Data Loopback Test 261
data protection 51
data replication 49, 197
DC power good LED 26
Default Host Group 46
defragmenting an array 45, 124
delete command 62
DHCP/BOOTP 35
DHCP/BOOTP server 38
diagnostic LEDs 18, 24
diagnostics 64, 114, 260
  analyzing Read Link Status results 262

controller runtime 260
Data Loopback Test 261
Read Link Status 261
Read Test 261
Write Test 261
direct management method 38
disable flashcopy script 330, 342
disaster plan 219
disaster recovery 219
disaster recovery plan 219
disaster recovery toolkit 219
disk array
defragmentation 123
expansion 123
disk drive 2
disk drive LED, FC 25
Disk Management 274
disk signature 179
diskpart script 330
distance limitations of Fibre Channel inter-switch links 196
driver installation 66
drive-side connection 17
DVE (Dynamic logical volume expansion) 7
dynamic disk 183
Dynamic logical volume expansion (DVE) 7
Dynamic logical-drive expansion 7, 35
Dynamic Volume Expansion 4, 44

E

enabling multipath I/O 72
enabling the Remote Volume Mirror option 204
Enterprise Management 6
Enterprise Management window 39–40, 87
error reporting 256
ESM 23
ESM board 22
ESM board LEDs 25
ESM board microcode 94
updating 95
Ethernet 87
port 15
Ethernet link speed LED 20
Event 1001 307
Event 1010 307
Event 1015 307
Event 1207 308
Event 150E 308
Event 150F 308
Event 1510 308
Event 200A 308
Event 202E 308
Event 2109 308
Event 210C 308
Event 210E 309
Event 2110 309
Event 2113 309
Event 2229 309
Event 222D 309
Event 2247 309
Event 2248 309
Event 2249 309
Event 224A 309
Event 224B 309
Event 224D 310
Event 224E 310
Event 2250 310
Event 2251 310
Event 2252 310
Event 2254 310
Event 2255 310
Event 2602 311
Event 2801 311
Event 2803 311
Event 2804 311
Event 2807 311
Event 2808 311
Event 280A 311
Event 280B 311
Event 280D 311
Event 280E 312
Event 280F 312
Event 2813 312
Event 2815 312
Event 2816 312
Event 2818 312
Event 281B 312
Event 281C 312
Event 281D 313
Event 2821 313
Event 3019 313
Event 5005 313
Event 502F 313
Event 5602 313
Event 560B 313
Event 560C 314
Event 560D 314
Event 560E 314
Event 560F 314
Event 5610 314
Event 5616 314
Event 5617 314
Event 6101 315
Event 6200 315
Event 6201 315
Event 6202 315
Event 6400 315
Event 6401 315
Event 6402 316
Event 6503 316
Event 6505 316
event log 121, 257
saving 122
viewing 121
Event Monitor 6, 42, 111
Event Viewer 257
exclusive-or (XOR) operation 53
EXP500 3–4, 10, 13
EXP700 4, 10, 13
cabling 26

diagnostic LEDs 24
EXP700 expansion enclosure 21
expansion enclosure 21

F

fabric configuration 196
fabric zoning 264
fan and communications module LEDs 20
Fan fault LED 20
fan module LED 26
FAStT Agent
installing on Solaris 84
FAStT Client 66–67, 76
for AIX 79
installation 72
installing on HP-UX 80
installing on Novell 77
installing on Solaris 84
FAStT configuration, saving 121
FAStT Event Monitor 67
FAStT Field Tool 94
FAStT Host Agent 67–68, 76
installing on HP-UX 81
FAStT host software 232
FAStT installation problems 263
FAStT Management Suite Java (MSJ) 56, 65–66, 133, 266
installation 71
FAStT RDAC 67–68, 84, 236
FAStT Runtime Environment 66
installation 72
installing on HP-UX 80
installing on Solaris 84
FAStT Storage Agent 245
FAStT Storage Manager Client 86
FAStT Storage Server configuration 85
FAStT Utility 58, 67–68
installation 72
installing on HP-UX 81
installing on Novell 77
FAStT-2 Host Adapter 13
FAStT200 2–3, 13–14
FAStT500 2–3, 13–14
FAStT700 xi, 2, 4, 358
cabling 26
maintenance 249
managing 5
Fault controller LED 19
Fault LED 26
FC disk drive LED 25
FC-AL 51
Feature Enable Identifier 128
Fibre Channel 2
Fibre Channel inter-switch link 196, 229
Fibre Channel loop 261
FlashCopy xi, 1, 4–5, 7, 11, 34, 44, 48, 139, 201, 358
disabling 48
enabling 150
hardware 201
how it works 140
key generation 150
Linux disks 188
operating systems 201
recreate script 178
re-creating 48
repository logical drive capacity 145
scenarios with scripts 176
scripts 176
Windows 2000 basic disks 177
Windows 2000 dynamic disks 183
FlashCopy example scripts 327
FlashCopy logical drive 11, 34, 140
creating 140, 152
creating on AIX 291
creating on Solaris 295
creating on UNIX 287
creating on Windows 273
creating using NetWare 304
deleting 174
disabling 141
maintenance 142
mapping 161
re-creating 141, 171
reusing in AIX 294
reusing on HP-UX 301
reusing on Solaris 298
reusing on UNIX 289
reusing on Windows 275, 278
reusing using NetWare 305
FlashCopy logical drive name 158
FlashCopy Logical Drive Wizard 140, 151
FlashCopy parameters 142
FlashCopy repository
capacity 145
increasing capacity 147
life 146
resizing the drive 172
FlashCopy repository logical drive 48, 260
FlashCopy repository logical drive full 315
FlashCopy repository logical drive name 158
flushing 252
full synchronization 49

G

GenKey file 202
global cache settings 117
global hot spare 98

H

hardware for FlashCopy and RVM 201
Heartbeat LED 19
heterogeneous hosts 9, 47
heterogeneous operating systems 3
high availability 3
high performance 14
host 45
host adapter 13
host adapter drive 80
host adapter driver 76, 79

installing on Solaris 84
host bus adapter 67
host computer 6
host connection 4
host group 45
host port 45, 108
host software 80, 84
for AIX 79
installation 66, 72
installing on Novell 76
host system, mapping logical drives to 46
host type 101
host-agent management method 36
host-side connection 17
hot spare
assign 99
defining drives 98
disk drive 55
unassign 100
hot_add utility 7, 58, 110, 141
hot-swappable components 3
HP_UX with Logical Logical Drive Manager 290
HP-UX 34, 47, 79, 267, 290
creating a FlashCopy logical drive 299
enabling multipath I/O 81
installing FAStT Client 80
installing FAStT Runtime Environment 80
installing host adapter drive 80
installing the FAStT Host Agent 81
installing the FAStT Utility 81
installing the host software 80, 84
Logical Logical Drive Manager 299
reusing FlashCopy logical drives 301

I
I/O performance 54
IBM AIX 34
IBMSAN 66
ICRC (Invalid Cyclic Redundancy Check) 263
ID switch 23
in-band management method 36, 66
initial FlashCopy script 328
initial script 178
installing on FAStT Host Agent 77
installing on Solaris 84
internal cabling 17
Intersite with FlashCopy drives and tape backup 229
intersite with redundant fabric 228
inter-switch link 196, 229
Invalid Cyclic Redundancy Check (ICRC) 263
Invalid Transmission Word (ITW) error counts 262
ITW (Invalid Transmission Word) 262

J
Java Virtual Machine on Novell 77

L
latency 30
Link Failure 262
link speed selection 23
Linux 5, 47, 266
RedHat 34, 70
SuSE 34, 70
TurboLinux 34, 70
Linux disks 342
disable flashcopy script 342
recreate flashcopy script 343
Locate component 115
logical drive 5, 100, 259
changing ownership 122
creating 100
customize settings 103
expansion 126
limits 195
mapping to host system 46
mappings 105
modification priority 57, 255
naming 103
not on preferred path 265
segment size 125
Logical Drive Manager Logical 270
Logical Drive Manager Logical Drives 290
logical drive ownership 200
logical drive-to-LUN mapping 259
Logical Logical Drive Manager 299
Logical/Physical View 40
LUN masking 9
LVM Logical Logical Drives 291

M
maintenance 114, 249
major event log 121
management ports 15
management station 5
managing the FAStT700 Storage Server 5
mapping
access logical drive 264
definition 109
logical drives to host systems 46
secondary mirror drives 219
Mappings View 41
maximum I/O per second 251
maximum latency 30
media scan 55
media scanning 119
Microsoft Cluster (MSCS) environment 233
Microsoft Cluster Services 34
Microsoft Virtual Machine 6
Microsoft Windows 5
Microsoft Windows 2000 34
Microsoft Windows NT4 34
mini-hub 14, 16–17, 20, 196
mirror properties 215
mirror relationship 193, 212
recreating 221
mirror repository 193
mirror repository logical drive 49, 260
mirror synchronization priority 218

missing logical drive 259
modification priority 116, 255
modification priority rate 255
monitoring 111, 250
MSCS (Microsoft Cluster) 34, 233
MSJ (Management Suite Java) 65–66, 133
  GUI 135
  installation 134
multipath I/O 72, 78

N

NetWare 66, 266, 270, 303–304
  creating a FlashCopy logical drive 304
  reusing FlashCopy logical drives 305
network attached management 85
Network Management Station (NMS) 42
network setup of controllers 129
NMS (Network Management Station) 42
Novell
  enabling multipath I/O 78
  installing host adapter driver 76
  installing the FAStT Client 77
  installing the FAStT Host Agent 77
  installing the FAStT Utility 77
  installing the host software 76
  installing the Java Virtual Machine 77
Novell Cluster Services 34
Novell NetWare 47, 76
Novell NetWare 5.1 34
Novell NetWare 6 34
NVSRAM 236, 240, 242–244
NVSRAM package 88
NVSRAM settings 92

O

operating systems supported 201
Out of Compliance error condition 127
out-of-band management 38
output view 59
ownership of a logical drive 122

P

password 56, 98
  considerations 56
Performance Monitor 250
performance monitor data 57, 249
physical connection 264
Point-in-time 140
power supply LEDs 20, 26
preferred controller ownership 56
preferred path for logical drive 265
premium feature 127
  disabling 128
  enabling 128
  listing 127
primary controller 49
primary logical drive 51, 192–193
Primitive Sequence Protocol (PSP) 263
problem notification 256
Properties component 114
PSP (Primitive Sequence Protocol) 263
PV-Links 66

R

RAID 2, 5, 33–34
  level selection 102
RAID 0 52
RAID 1 52
RAID 10 52
RAID 5 53
RAID controller 5, 14–15, 17, 19
RAID controller unit 3
RAID level 4, 8, 52, 57, 254
  changing 123
RDAC (Redundant Disk Array Controller) 6, 56, 66
Read Caching parameter 118
Read Link Status 261
Read Link Status Diagnostic (RLS) 7, 35
Read Link Status error counts 261
Read Link Status results 262
Read Percentage 251
Read Test 261
read-ahead multiplier 119, 253
read-intensive environment 53
recovery from switch failure 220
Recovery Guru 64, 258, 263
recreate flashcopy script 343
recreate script 178
Redbooks Web site 346
  Contact us xiii
redundancy 17
Redundant Disk Array Controller (RDAC) 6, 56
redundant drive-side FC loop connection 14
redundant fabric 228
redundant power supplies 3
ReiserFS file system 44
Remote Mirroring Wizard 50
Remote Volume Mirror 191, 259
  creating relationships 208
  option 205
  premium feature 202
  status 203
Remote Volume Mirroring xi, 4, 11, 34–35, 58, 201, 255, 358
  solution design 226
Remote Volume Mirroring (RVM) 1, 7, 49
repository life 146
repository logical drive 48, 193
repositoryFullPolicy=failBaseWrites 178
repositoryFullPolicy=failFlashCopy 178
repositoryPercentofBase=n 178
repositoryTresholdPercent=n 178
reversing roles of primary/secondary drives 221
RLS (Read Link Status Diagnostic) 7, 35
role reversal of primary/secondary drives 221
rolling upgrade 234
RVM 7, 11, 34–35, 49, 191
  activating 204

changing secondary to primary logical drive 222
creating Remote Volume Mirror relationships 208
data replication 197
distance limitations 196
enabling 204
enabling and activating 204
enabling the premium feature option 202
Intersite with FlashCopy drives and tape backup 229
link loss 199
logical drive ownership 200
mapping secondary mirror drives 219
mirror properties 215
mirror relationships 212
mirror synchronization priority 218
performance considerations 198
recovery from switch failure 220
recreating a mirror relationship 221
role reversal 221
simple department with minimum redundancy 227
solution design 226
status 203
storage subsystem profile 214
supported configurations 201
supported hardware 201
supported operating systems 201
switch zoning 197
using 202
working with FlashCopy 200
RVM (Remote Volume Mirroring) 1, 7, 49

S

SAN fabric 220
SAN Switch F08 30
SAN Switch F16 30
Script Editor 58
comments 59
show statement 60
usage guidelines 59
Script Engine 7, 35
script view 59
secondary logical drive 51, 192
connection loss 199
secondary to primary logical drive 222
security, password setting 98
segment size 57, 125, 254
serial interface 5
serial port 15
SFP module 17, 30
SFP module port 15, 22
show statement 60
simple department with minimum redundancy 227
single equation single unknown 53
site failover 220
SM 7.10 xi, 358
SM 8 xi, 358
SMagent 6
small-form pluggable module 15
SMcli script 342
SMclient 6
SMdevices utility 7, 58, 141
SMflashcopyassist utility 7, 58
SMsnapassist 275
SMTP 111
configuring the mail server 112
SMutil 6
SNMP 111
trap destination 113
Solaris 83, 288, 290
creating a FlashCopy logical drive 295
installing host adapter driver 84
installing the FAStT Agent 84
installing the FAStT Client 84
installing the FAStT RDAC 84
Installing the FAStT Runtime Environment 84
reusing FlashCopy logical drives 298
Veritas Logical Drive Manager 290, 295
start percentage 54
starting cache flushing levels 252
static IP address 38
Status LED 19
stopping cache flushing levels 252
Storage Manager xi, 8, 33, 358
client 39
prerequisites 232
updating in a Microsoft Cluster environment 233
upgrading 231
Storage Manager 8.2 34, 48, 52
Storage Manager 8.2 Agent 6
Storage Manager 8.2 Client 6
Storage Manager 8.2 Utility 6
storage partition 4, 45–46
storage partitioning 7, 9, 35, 105
storage server 2
storage subsystem profile 214
storage subsystem unresponsive 265
Subsystem Management 6
Subsystem Management window 39–40, 50, 140
Sun Solaris 34, 47, 266
supported configurations for RVM 201
switch failure 220
switch zoning 197
synchronization 49
synchronization priority 198
synchronization priority rate 256
synchronous mirroring operation 49

T

terminal emulation utility 5
Total I/Os 250
Tray ID and Link Speed Module LEDs 25
tray ID switch 23
troubleshooting 23, 114
tuning 114, 250

U

Universal Xport Device 37
UNIX
creating a FlashCopy logical drive 287
Logical Drive Manager Logical 270

Logical Drive Manager Logical Drives 290
regular disks 286–287
reusing FlashCopy logical drives 289
unresponsive storage subsystem 265
updating Storage Manager in a Microsoft Cluster (MSCS) environment 233
updating the FAStT host software 232
upgrading from Windows NT 4.0 to Windows 2000 234
upgrading Windows in a cluster setup 234

V

Veritas Logical Drive Manager 290, 295

W

Windows 266
basic/regular disks 270–271
creating a FlashCopy logical drive 273
Dynamic Disks 280
dynamic disks 270
reusing FlashCopy logical drives 275, 278
Windows 2000 47
disable flashcopy script 330
diskpart script 330
initial FlashCopy script 328
Recreate FlashCopy Script 333
Windows 2000 basic disks 328
Windows 2000 dynamic disks 183, 335
disable flashcopy script 337
diskpart script 337, 341
initial flashcopy script 335
SMcli script 337, 342
Windows NT 47
Windows upgrade 234
Windows upgrade in a cluster setup 234
world-wide name (WWN) 45–46
write cache mirroring 253
Write caching 54
Write Caching parameter 118
Write Test 261
write-back caching 253
write-intensive 53
write-through caching 253
WWN (world-wide name) 45–46

X

XOR (exclusive-or operation) 53

Z

zoning in the fabric 264

IBM

Redbooks

IBM TotalStorage FAStT700 and Copy Services

(0.5" spine)
0.475"<->0.873"
250 <-> 459 pages

IBM TotalStorage FAStT700 and Copy Services

Explore the new features of the FAStT700 hardware

Learn about FlashCopy, its components, and how to use them

Find out how Remote Volume Mirroring can protect data

This IBM Redbook introduces the IBM TotalStorage FAStT700 Storage Server and the new Copy Services functionality provided in the latest release of the Storage Manager software. It explains FAStT700 and discusses the hardware in detail. It examines FAStT Storage Manager, with an emphasis on the release of version 8.2, which introduces Remote Volume Mirroring. It offers a step-by-step guide to using the Storage Manager to create arrays, logical drives, and other basic management tasks.

This redbook includes details on:

- FlashCopy: How it is implemented and how to use it
- Remote Volume Mirroring: What it is and how to set up and control it
- Migrating from Storage Manager 7.10 to Storage Manager 8
- Performance monitoring and error reporting

Plus you'll find practical information on:

- Host-specific instructions for FlashCopy
- Critical event descriptions
- Command Line Reference
- FlashCopy example scripts

This book targets IBM technical professionals, Business Partners, and Customers who want to learn more about the capabilities of the advanced functions introduced with FAStT Storage Manager Software version 8.2. It also targets those who have a FAStT storage subsystem and need detailed advice on how to configure it.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6808-00 ISBN 0738426717