IBM TotalStorage: FAStT600/900 and Storage Manager 8.4

New FAStT models

Storage Manager V8.4

VolumeCopy

Bertrand Dufrasne
Bernd Baeuml
Carlos De Nobrega
Ales Leskosek
Stephen Manthorpe
Jonathan Wright

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM TotalStorage: FAStT600/900 and Storage Manager 8.4

November 2003

SG24-7010-00

Note: Before using this information and the product it supports, read the information in "Notices" on page xix.

First Edition (November 2003)

This edition applies to FAStT Storage Manager Version 8.4 for use with the IBM TotalStorage FAStT600 (Turbo) FAStT700 and FAStT900 Storage Servers.

Contents

Figures

Tables

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:

AIX®
Approach®
Balance®
Enterprise Storage Server®
eserver™
FlashCopy®
ibm.com®
IBM®
Lotus®
Netfinity®
Notes®
pSeries™
Redbooks(logo) ™
Redbooks™
S/390®
SANergy™
Sequent®
Tivoli®
TotalStorage®
eserver ™
Wave®
xSeries®

The following terms are trademarks of International Business Machines Corporation and Rational Software Corporation, in the United States, other countries or both.

Rational Software Corporation®
Rational®

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 reviews, in detail, the hardware of the IBM TotalStorage FAStT600 and FASTtT900 Storage Servers and presents the new features introduced with the FAStT Storage Manager Version 8.4 (SM 8.4).

After explaining the concepts and functions used in planning and managing the FAStT Storage Server, the book offers a step-by-step guide to using the Storage Manager to create arrays, logical drives, and other basic (as well as advanced) management tasks.

This publication also contains practical information on diagnostics and troubleshooting, critical event descriptions, usage of scripts, and the Command Line Interface.

Finally, this publication includes chapters dedicated to the Copy Services premium features, and covers the VolumeCopy feature provided in the latest release of the Storage Manager software.

This publication 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 V8.4. It also targets those who have a FAStT storage subsystem and need detailed advice on how to configure it.

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.

Bertrand Dufrasne is a Certified Consulting I/T Specialist and Project Leader for Disk Storage Systems at the International Technical Support Organization, San Jose Center. He has worked at IBM for 21 years in many I/T areas. Before joining the ITSO he worked for IBM Global Services in the US as an I/T Architect. He holds a degree in Electrical Engineering.

Bernd Baeuml is a Senior Technical Support Specialist for IBM Storage Division Back Office based in Greenock, Scotland. He has seven years of experience with IBM hardware and software. His areas of expertise include FAStT Storage Server, Windows and Unix based operating systems. He is a co-author of two redbooks about Netfinitiy server management. He holds a Masters of Engineering degree from the Hochschule fuer Technik and Wirtschaft Dresden. He is PSE and CNE.

Carlos De Nobrega is a xSeries and FAStT Product Manager working in the Field Support Group in IBM Global Services in South Africa, performing level 2 support and product management for xSeries and FAStT Storage Products. He has five years of experience in the High Volume Intel platform/storage field. He has worked at IBM for eight years. He is qualified as an Electrical Engineer and his main focus of interest and expertise is in the Intel Server market and FAStT/Storage Area Networks.

Ales Leskosek is an I/T Specialist with IBM Slovenia and has nearly four years of experience as Field Technical Support Specialist for the CEMA region. His activities include pre-sales and post-sales support, from designing end-to-end storage solutions and competitive positioning to implementation and problem determination across the entire range of IBM TotalStorage products. Ales has taught storage classes and has spoken as a subject matter expert at industry events.

Stephen Manthorpe is a country Enterprise Systems Disk Specialist, based in Canberra, Australia. He joined IBM in 1988 and has 15 years of experience in S/390 systems, storage products, and laser printing systems. He has for three years provided technical support for ESS and SAN to Australia, New Zealand, and the ASEAN region.

Jonathan Wright is a technical specialist in New Zealand. He has 10 years of experience in the Intel server and storage field. His areas of expertise include xSeries hardware, Linux, clustering, and FAStT Storage

Figure 1 The team: Ales, Bernd, Bertrand, Jonathan, Stephen, Carlos

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

Barry Mellish
Deanna Polm
Emma Jacobs
International Technical Support Organization, San Jose Center

Julie Czubik
International Technical Support Organization, Poughkeepsie Center

Jay Smith
George Thomas
Shawn Bramblett
Chris Small
IBM SSG Tucson

Dolores Butcher
IBM Tucson

Michaell Quillen
IBM Beaverton

Arwed Tschoeke
IBM Germany

Bruce Allworth
Gene Cullum
James Goodwin
Todd Virnoche
IBM Advanced Technical Support - Americas

Mic Watkins
IBM Raleigh

Chuck Grimm
IBM Technical Support Marketing Lead

Harold Pike
IBM Raleigh -

John Murtagh
IBM SSG FAStT Product Manager

Tai Chang
IBM San Jose FAStT Program Manager

Ray Koehler
IBM US

Tim Simon
IBM US

Dave Worley
John Bish
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 FAStT

This chapter reviews the IBM TotalStorage FAStT family of storage products; introduces the most recent additions, the FAStT600, FAStT600 Turbo, and the FAStT900; and discusses where they fit in terms of storage solution.

The rest of this chapter summarizes the functions of the FAStT Storage Manager software, and how its features have evolved over time, up to the current 8.4 version.

1.1 FAStT models

IBM TotalStorage Fibre Array Storage Technology (FAStT) Storage Server is a redundant array of independent disks (RAID) storage subsystem that contains the Fibre Channel (FC) interfaces to connect both the host systems and the disk drive enclosures. The Storage Server provides high system availability through the use of hot-swappable and redundant components. This is crucial, because the Storage Server is placed in high-end customer environments such as server consolidation on Storage Area Networks (SANs).

Figure 1-1 shows the evolution of the IBM TotalStorage FAStT series.

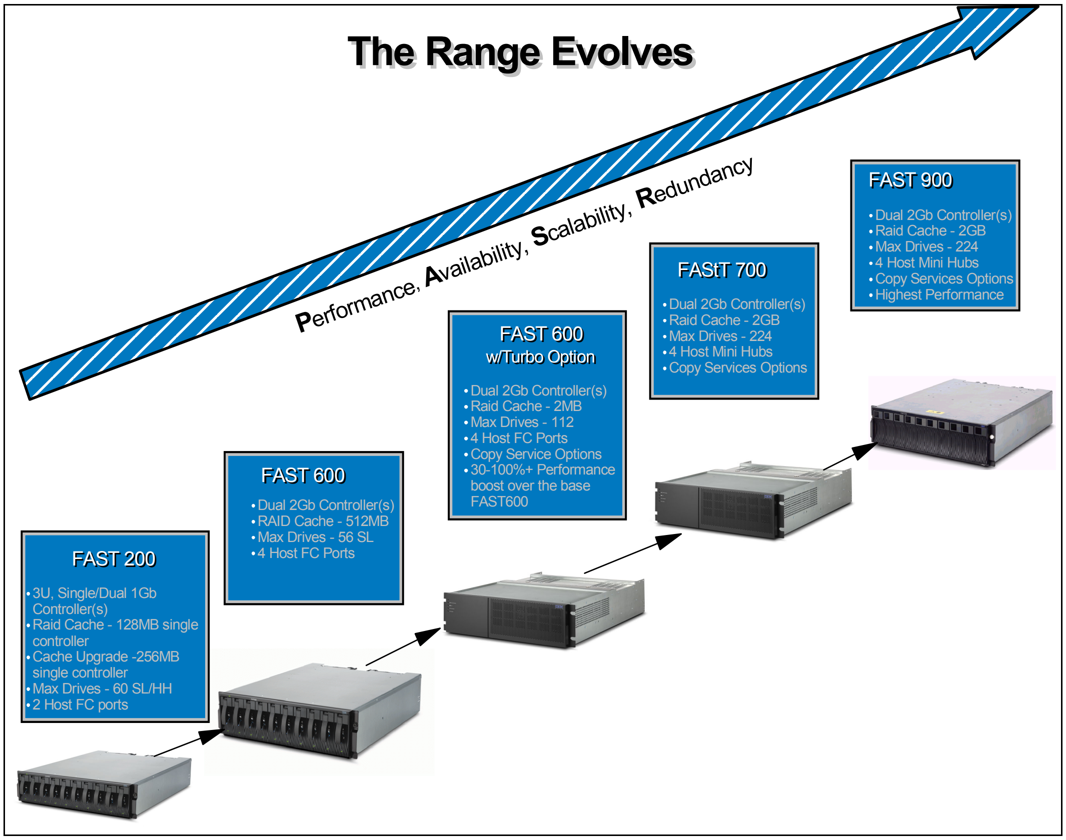

Figure 1-1 The FAStT evolution

The servers are:

- FAStT200 Storage Server

  The FAStT200 is designed for workgroup and departmental servers that require an external storage solution. The single controller model provides a cost-effective solution, while the FAStT200 High Availability (HA) model features a fully redundant configuration with dual-active controllers. As your storage requirements grow, you can easily expand storage capacity by adding IBM FAStT EXP500 (50 drives) or EXP700 (56 drives). Expansion units scale from 18Gb to 1.47TB in a compact 3U size with a maximum system capacity of 9.6TB. FlashCopy is supported and provides fast data duplication capability, reducing or eliminating the need for long shutdowns during backups and restores.

- FAStT600 and FAStT600 Turbo Storage Servers

  The FAStT600 and FAStT600 Turbo are among the latest additions to the FAStT family of products.

The FAStT600 is an entry level, highly scalable 2Gb Fibre Channel Storage Server.

It is designed to a be cost-effective, scalable storage server for consolidation and clustering applications. Its modular architecture can support on demand business models by enabling an entry configuration that can easily grow as storage demands increase (up to 8.2TB of capacity). Supporting up to 56 drives (using three EXP700 Expansion units), it is designed to deliver high performance of up to 400MB/sec. Dynamic capacity addition provides the ability to add an EXP700 enclosure to an existing FAStT600 without stopping operations. The FAStT600 can provide capacity on demand, allowing unused storage to be brought online for a new host group or an existing volume. The FASt600 can be upgraded to the Turbo model.

The FAStT600 Turbo is a mid-level storage server that can scale to over sixteen terabytes, facilitating storage consolidation for medium-sized customers. It uses the latest in storage networking technology to provide an end-to-end 2 Gbps Fibre Channel solution (the host interface on base FAStT600 is 2Gb, while Turbo auto senses to connect to 1Gb or 2Gb) and offers up to 70 percent performance improvement (with new Storage Manager V8.4 that ships with Turbo). It has higher scalability over the base FAStT600, up to 16.4TB for a total of 112 disks—using a maximum of 7 EXP700s. The FAStT600 Turbo supports up to 64 storage partitions. The cache has increased from 256MB per controller on base FAStT600 to 1GB per controller on Turbo; finally it offers autonomic functions such as Dynamic Volume Expansion and Dynamic Capacity Addition, allowing unused storage to be brought online without stopping operations, and FAStT Service Alert, which is capable of automatically alerting IBM if a problem occurs.

- FAStT500 Storage Server

  This Storage Server can support medium to high-end configurations with greater storage capability, as well as support heterogeneous host systems. This offers a higher level of availability, performance, and expandability than the FAStT200. The TotalStorage FAStT500 solution is designed to provide security against component failures. Dual hot swap RAID controllers help provide throughput and redundancy, and each controller supports up to 512MB of batter backed cache. Redundant fans, power supplies, and dynamic storage management further contribute to availability. Capacity is scalable from 18Gb to greater than 32TB supporting up 224 drives using either 22 EXP500 or 16 EXP700.

  Note: the FAStT500 has been withdrawn from marketing.

- FAStT700 Storage Server

  This offers 2Gbps technology for faster response time. Scales from 36GB to greater than 32TB of storage using 16 EXP700 expansion enclosures. Each expansion enclosure supports up to fourteen 2Gbps Fibre channel disk drives. Moreover, you can select the appropriate RAID level (from RAID 0,1, 3, 5, and 10) to match an application or suit particular needs. The FAStT700 also supports all premium features such as FlashCopy, VolumeCopy, Remote Volume Mirroring, and Storage Partitioning.

  This storage server supports high-end configurations with up to 64 heterogeneous host systems. High availability is critical in today's knowledge-based economy: The TotalStorage FASTT700 Storage Server is designed to support high availability, providing protection against component failures. Dual hot swap RAID controllers help provide high throughput and redundancy; each controller supports 1Gb of battery backed cache. Redundant fans, power supplies, and dynamic storage management further support high availability and help reduce the risk of costly downtime or the loss of valuable data.

- FAStT900 Storage Server

  IBM FAStT900 Storage Server delivers breakthrough disk performance and outstanding reliability for demanding applications in compute intensive environments. The FAStT900 is designed to offer investment protection with advanced functions and flexible features. Designed for today's on demand business needs, the FAStT900 easily scales from 36GB to over 32TB to support growing storage requirements. The FAStT900 offers advanced replication services to support business continuance and disaster recovery. The FAStT900 is an effective storage server for any enterprise seeking performance without borders.

  The FAStT900 uses 2Gb Fibre Channel connectivity to support high performance (772 MB/sec throughput from disk) for faster, more responsive access to data. It provides flexibility for multiplatform storage environments by supporting a wide variety of servers, operating systems and cluster technologies (certified for Microsoft Cluster Server, Novell clustering, HACMP, Veritas Cluster for Solaris). This storage server is well suited for high performance applications such as online transaction processing (OLTP), data mining and digital media.

Figure 1-2 shows how each TotalStorage FAStT model fits into its category of Open System Computing, from entry level to enterprise level.

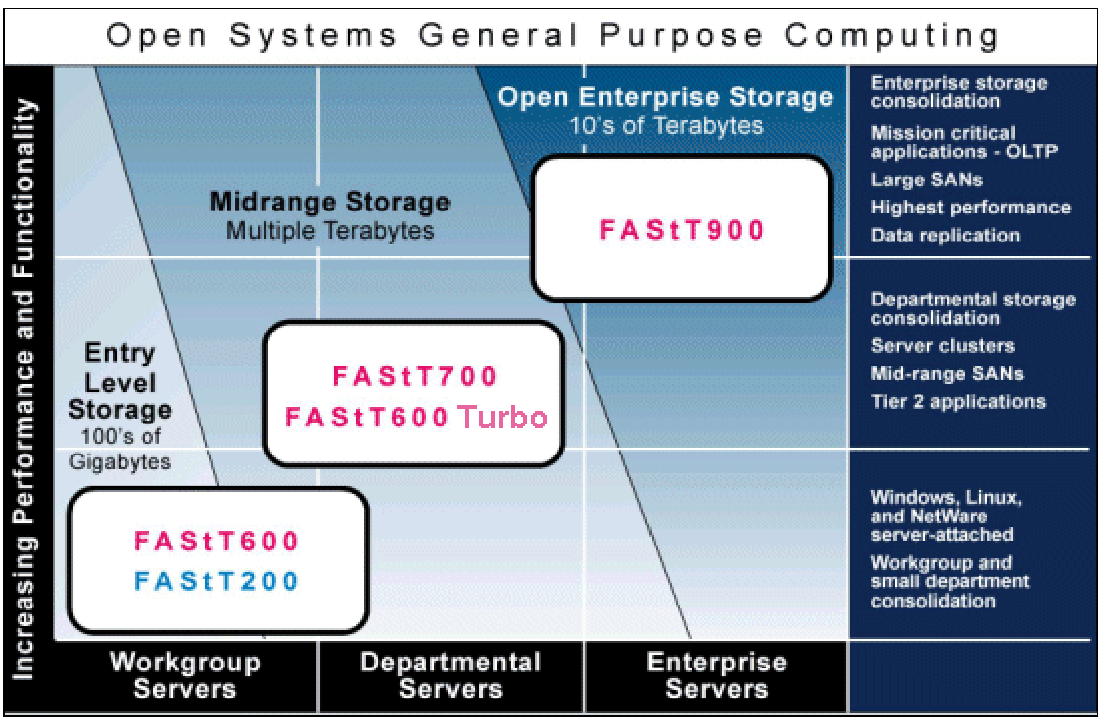

Figure 1-2 IBM TotalStorage FAStT Family

1.2 Introduction to Storage Manager

The IBM FAStT Storage Manager software is used to configure, manage and troubleshoot the FAStT Storage Server. Its front-end component, the Storage Manager client, is a Java-based management tool that is available for various operating systems.

It is used primarily to configure RAID arrays and logical drives, assign logical drives into storage partitions, replace and rebuild failed disk drives, expand the size of the arrays and logical drives, and convert from one RAID level to another. It allows troubleshooting and management tasks, like checking the status of the storage server components, updating the firmware of the RAID controllers, and managing the storage server. Finally, it offers advanced functions such as FlashCopy, Remote Volume Mirroring (RVM), and the new VolumeCopy (these are premium features that need to be purchased).

Chapter 4, “Managing the FAStT Storage Server” on page 67, and Chapter 5, “Step-by-step procedures using Storage Manager 8.4” on page 109, explain in detail the functions and features of Storage Manager, and include many illustrations of its usage.

Figure 1-3 shows the evolution of the FAStT Storage Manager.

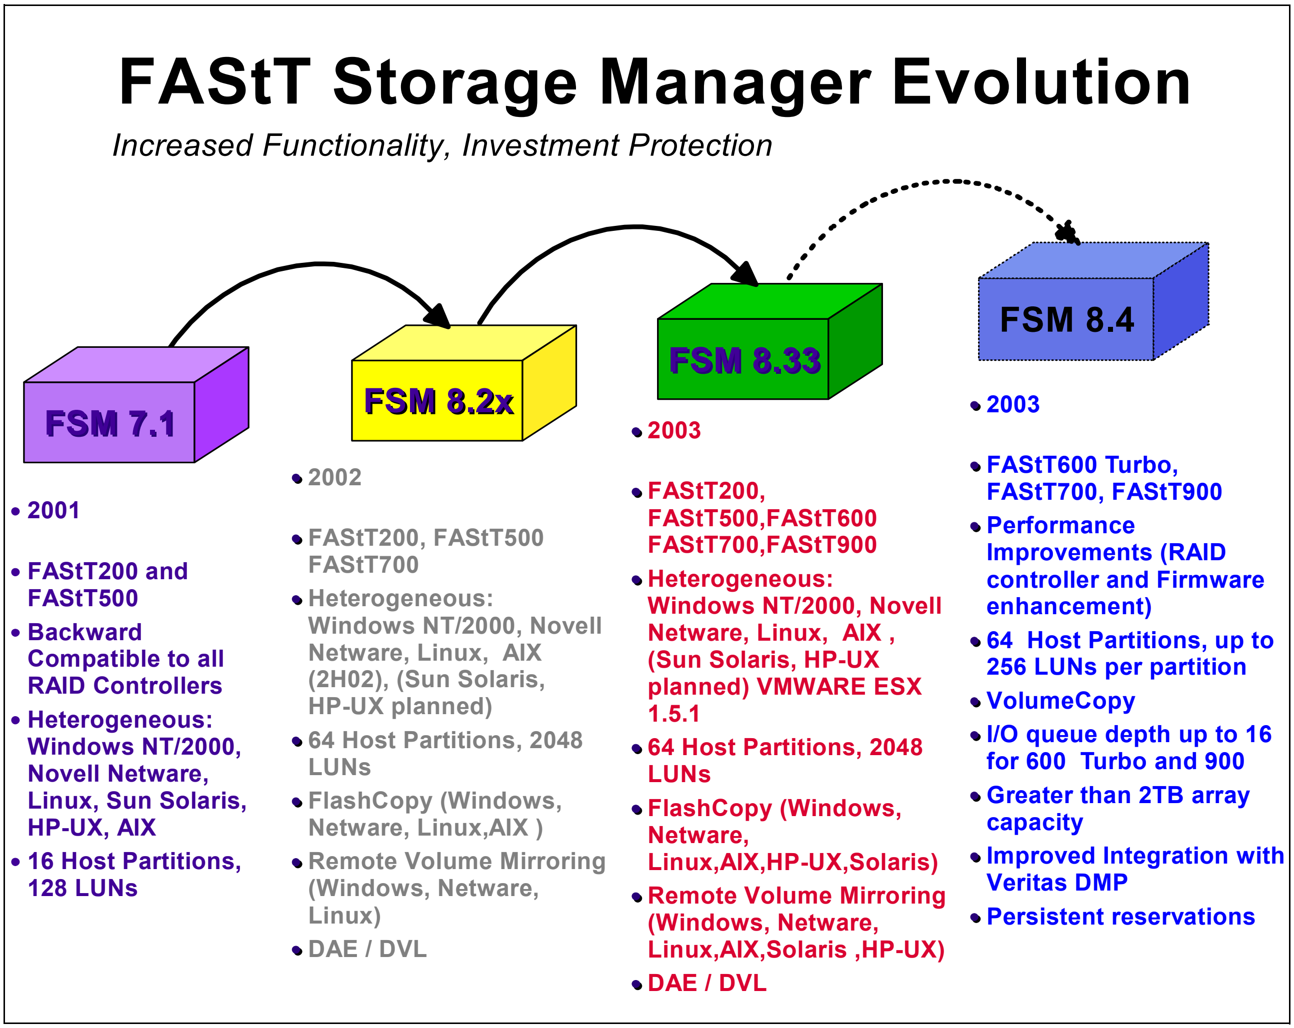

Figure 1-3 IBM FAStT Storage Manager evolution

Storage Manager 8.4 introduces several performance and capacity improvements as well as a new VolumeCopy premium feature. Refer to Chapter 2, “What is new in Storage Manager 8.4” on page 7, and Chapter 9, “VolumeCopy” on page 281.

1.2.1 FAStT Service Alert

FAStT Service Alert is available to all current and new FAStT Storage Server customers. FAStT Service Alert (hereafter called Service Alert) enables the IBM TotalStorage FAStT Storage Manager to monitor system health and automatically notify the IBM Support Center when problems occur. Service Alert sends an e-mail to an IBM call management center that identifies your system and location details such as your phone number. The IBM Support

Center analyzes the contents of the e-mail alert and contacts you to begin problem determination. The service is available worldwide.

FAStT Service Alert requires a services contract. Please contact your IBM representative for further information regarding the service offering.

For detailed information please refer to "IBM FAStT service alert" on page 158.

2

What is new in Storage Manager 8.4

For your convenience we have grouped in this chapter most of the new features introduced in Storage Manager 8.4.

It is assumed that before proceeding to this chapter the reader already has some familiarity with past versions of the Storage Manager software and a good overall understanding of its concepts and features. If this is not the case, you are encouraged to first read Chapter 4, “Managing the FAStT Storage Server” on page 67.

2.1 Storage Manager 8.4 new features

Storage Manager 8.4 and the corresponding firmware is only supported on the FAStT600 Turbo, FAStT700, and FAStT900.

2.1.1 VolumeCopy

VolumeCopy is an optional, firmware-based feature that is used to copy data from one logical drive (the source logical drive) to another logical drive (the target logical drive) in a single storage subsystem. VolumeCopy does a clone copy (track for track).

Note: VolumeCopy is the name of the feature; within the SM 8.4 code it is referred to as Logical Drive Copy.

VolumeCopy is a premium feature and must be enabled by purchasing a feature key file from IBM. The FlashCopy premium feature is also a prerequisite for the VolumeCopy feature.

This premium feature includes a Create Copy Wizard, to assist in creating a VolumeCopy; and a Copy Manager, to monitor logical drive copies after they have been created.

VoumeCopy can be used for:

- Copying data for greater access
- Backing up data
- Restoring FlashCopy logical drive data to the base logical drive

VolumeCopy is a major new feature available with SM8.4 and we review that feature in detail first in “VolumeCopy overview” on page 102 and then provide practical details on its usage in Chapter 9, “VolumeCopy” on page 281.

2.1.2 Persistent reservations

Persistent reservations is a SCSI-3 feature for restricting access to storage media, based on the concept of reservations that the host can establish and manipulate. Earlier versions of SCSI provide a simple reservation capability through the RESERVE and RELEASE commands; SCSI-3 persistent reservations provide a significant super-set of the earlier capability. Improvements that come with persistent reservations include:

- Well-defined model for reserving across multiple host and target ports
- Levels of access control, for example, shared reads, exclusive writes, exclusive reads and writes
- Ability to query the storage system about registered ports and reservations
- Provisions for persistence of reservations through power loss at the storage system

A logical drive reservation is a feature of the cluster software (the actual reservation of the logical drive is handled by the host application) that allows one or more host ports to reserve a logical drive, thus preventing other host ports to access the same logical drive. Persistent reservations, which are managed through cluster sever software, preserve logical drive reservations and prevent other hosts from accessing the logical drive.

Unlike other types of reservations, a persistent reservation reserves across multiple host ports, provides various levels of access control, offers the ability to query the storage subsystem about registered ports and reservations, and, optionally, provides for persistence of reservations in the event of a storage system power loss.

The benefits of the persistent reservations feature is that it allows the FAStTStorage Server to integrate with cluster solutions that use shared logical drives for increased availability, scalability and performance. For instance, this enables VERITAS Dynamic Multipathing (DMP) to manage FAStT when the VERITAS Cluster Server is in use (see “Supporting Veritas DMP for Solaris” on page 12).

To view persistent reservations

Select any desired node in the Subsystem Management window, and then select the **Advanced -> Persistent Reservations** pull-down menu option in the logical view.

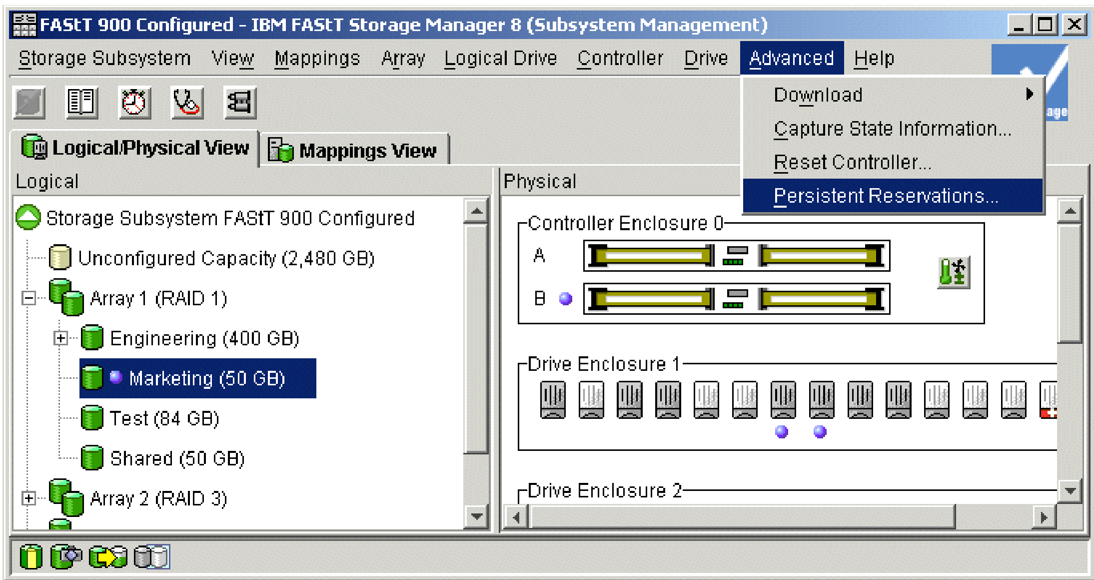

Figure 2-1 SM 8.4 -View persistent reservations

The Persistent Reservations window displays, showing any logical drives in the storage subsystem that have registrations, with the first logical drive in the list highlighted by default.

As shown in Figure 2-2 on page 10, the following information is displayed:

- Logical Drive Name: Displays the user label of the logical drive with persistent reservations. Logical drives are listed in alphabetical order. If a logical drive user label is not available, then its logical drive identifier (WWN) is shown.
- LUN: Displays the assigned LUN number for the particular logical drive.
- Accessible By: Displays the associated host group or host name that has access to the logical drive. If the topology has not been defined, then the term `Default Group` is displayed instead.
- Registrants: Displays the number of registrations for the particular logical drive.
- Reservation type: Each addressable logical drive can have one reservation, and each reservation can grant access rights to one or more registrants, depending on the reservation type. A logical drive can be reserved for a specific access level by a group of registrants. All registrants within a group are restricted to the access level defined by the reservation type. Reservation types are as follows:
  - WE (Write Exclusive): Only the host port that reserved the logical drive may write to the logical drive. Any host port may read from the logical drive.
  - EA (Exclusive Access): Only the host port that reserved the logical drive may read or write to the logical drive.

– WE-RO (Write Exclusive - Registrants Only): All registered host ports may write to the logical drive. Any host port may read from the logical drive.
– EA-RO (Exclusive Access - Registrants Only): Only registered host ports may read or write to the logical drive.
– WE-AR (Write Exclusive - All Registrants): All registered host ports may write to the logical drive. Any host port may read from the logical drive.
– EA-AR (Exclusive Access - All Registrants): Only a host port may read or write to the logical drive.

To view the associated registrations (that is, the associations between the initiator host port, the controller port logical drive, and the logical drive reservation) double-click the desired logical drive or highlight the logical drive and select the **View Associated Registrations** check box in the upper left of the dialog (see Figure 2-2); note that only one logical drive's registration may be viewed at a time.

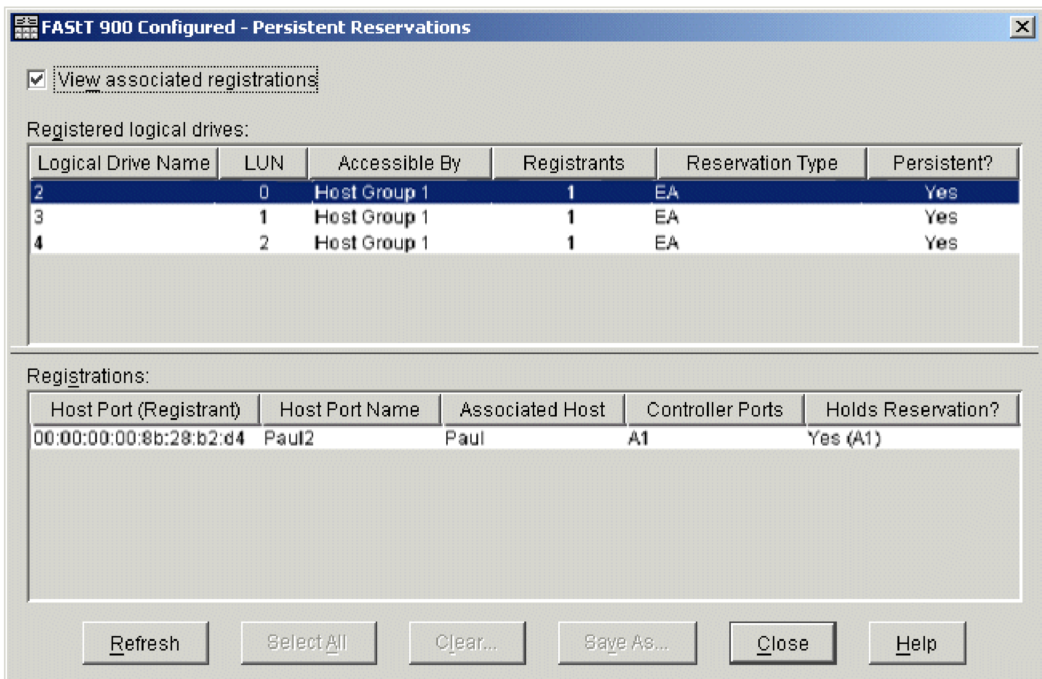

Figure 2-2 Display Persistent Reservation and Registrations

The registration information consists of:

- Host Port (Registrant) Identifier: Displays the WWN of the associated host port registered for the highlighted logical drive.
- Host Port Name: Displays the associated user label for the particular host port. If the host port name has not been provided, then `Not Available` is displayed.
- Associated Host: Displays the host associated with the particular logical drive.
- Controller Ports: Displays the controller slot (A or B) and the associated controller ports onto which the host port is logged. For example, A1, A2, B1, or B2 for controllers with four host-side interfaces.
- Holds Reservation?: Displays either Yes or No, depending on whether the particular host port is the reservation holder. If the registrant holds the reservation, the controller port that holds the reservation is indicated. For example, Yes (A1), Yes (B2). If the reservation type is such that all ports hold the reservation, then the designation is Yes (list of applicable

controller ports). For example, Yes (A1, A2, B1). For certain reservation types, all registrants for the specific logical drive are considered reservation holders and each registrant has Yes under this column. In most cases, only one registrant holds the reservation.

Persistent reservations are allowed on a primary logical drive in a Remote Mirror, but are not allowed on a secondary logical drive. If a logical drive has any type of reservation when designated as a secondary logical drive, the primary logical drive detects a reservation conflict at its first write request to the secondary logical drive and clears the reservation automatically. Subsequent requests to place a reservation on the secondary logical drive are rejected.

2.1.3 Support for 256 LUNs per storage partition

256 LUN support allows the storage subsystem to present up to 256 host-addressable LUNs (numbered 0–255) to a given host port, providing greater connectivity and storage capacity for SAN environments. For instance, this will allow support for larger databases and offer more ‘slots’ for FlashCopy and VolumeCopy logical drives.

This is an increase from a maximum of 32 in previous versions of Storage Manager. This capability is a fundamental attribute of the IBM TotalStorage FAStT product and will be present on any system executing a firmware revision level that supports this feature.

Important: The feature is operating system dependent and is enforced through NVSRAM settings.

Most hosts will be able to have 256 LUNs mapped per storage partition. Windows NT, Solaris with RDAC, NetWare 5.1, and HP-UX 11.0 are restricted to 32 LUNs. If you try to map a logical drive to a LUN that is greater than 32 on these operating systems, the host will be unable to access it. Solaris will require use of Veritas DMP for failover for 256 LUNs.

If storage partitioning is disabled, or is enabled, but host ports or volumes exist that have not been explicitly mapped to a partition, hosts make use of the default host group. The default host group partition also allows 256 LUNs to be mapped.

Figure 2-3 shows the dialog with the modified field that allows users to specify LUNs in the range 0 through 255.

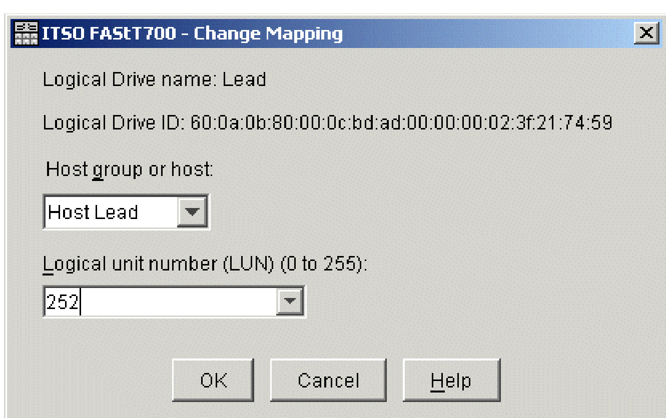

Figure 2-3 Mapping up to 256 LUNs

The access LUN (a mappable volume that allows "in-band" configuration and monitoring of the storage array by the host software client) can be set by the user, and is limited to the range of LUN numbers 0–31, inclusive, within a given partition.

2.1.4 Supporting Veritas DMP for Solaris

Veritas Logical Drive Manager with Dynamic Multi-Pathing (DMP) is a multi-path driver. In previous versions of Storage Manager, support for the Veritas DMP driver on the Solaris OS is achieved through firmware-based emulation of the Sun T3 RAID system. This method is somewhat limited in that it is based on a hard-coded inquiry string binding, which is internal to DMP.

Storage Manager 8.4 offers a more general method (rather than firmware based) of supporting Veritas DMP. The new method uses a software component, called an Array Support Library (ASL).

The ASL software provides information to the Veritas Logical Drive Manager for setting up the path associations for the driver. This gives the ability to have a multipathing solution on Solaris using Sun Fibre Channel host adapters and the ability to work with DMP without using Sun T3 emulation. It also gives the ability to direct the exclusion of access LUNs from the Veritas control.

To enable this support you must install ASL and also change the host type to Solaris with DMP, as illustrated in Figure 2-4. (Note that DMP does not provide load balancing.)

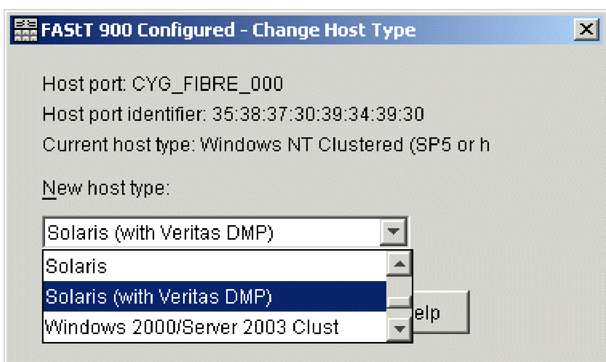

Figure 2-4 Support for Solaris with Veritas DMP

The ASL is bundled with the Storage Manager client code for Solaris (SMibmasl package). To install the SMibmasl package, using the following command:

```
#pkgadd -d SMibmasl_pkg
```

2.1.5 ADT alert notification and failover alert delay

Storage Manager 8.4 provides alert notification on ADT-induced logical drive ownership changes. The logical drive transfer alert notification is issued for any instance of a logical drive owned by a non-preferred controller, whether ADT is enabled or not, and is in addition to any informational or critical event already logged within the ADT or RDAC context.

As part of the enhancement, a failover alert delay can be specified that lets you delay the logging of a critical event if the multipath driver transfers logical drives to the non-preferred controller. If the multipath driver transfers the logical drives back to the preferred controller within the specified delay period, no critical event is logged. If the transfer exceeds this delay period, then a logical drive-not-on-preferred-path alert is issued as a critical event. This option also can be used to minimize multiple alerts when many logical drives fail over because of a system error, such as a failed host adapter.

The logical drive not-on-preferred-path alert is issued for any instance of a logical drive owned by a non-preferred controller and is in addition to any other informational or critical failover events.

Attention: Whenever a logical drive not-on-preferred-path condition occurs, only the alert notification is delayed; a needs attention condition is raised immediately.

To make the best use of this feature, set the failover alert delay period such that the host driver failback monitor runs at least once during the alert delay period. Note that a logical drive ownership change might persist through the alert delay period, but correct itself before you can inspect the situation. In such a case, a logical drive-not-on-preferred-path alert is issued as a critical event, but the array will no longer be in a needs-attention state.

To set the failover alert delay, select the storage subsystem from the Subsystem Management window, and then select the **Storage Subsystem -> Change -> Failover Alert Delay** pull-down menu option (see Figure 2-5).

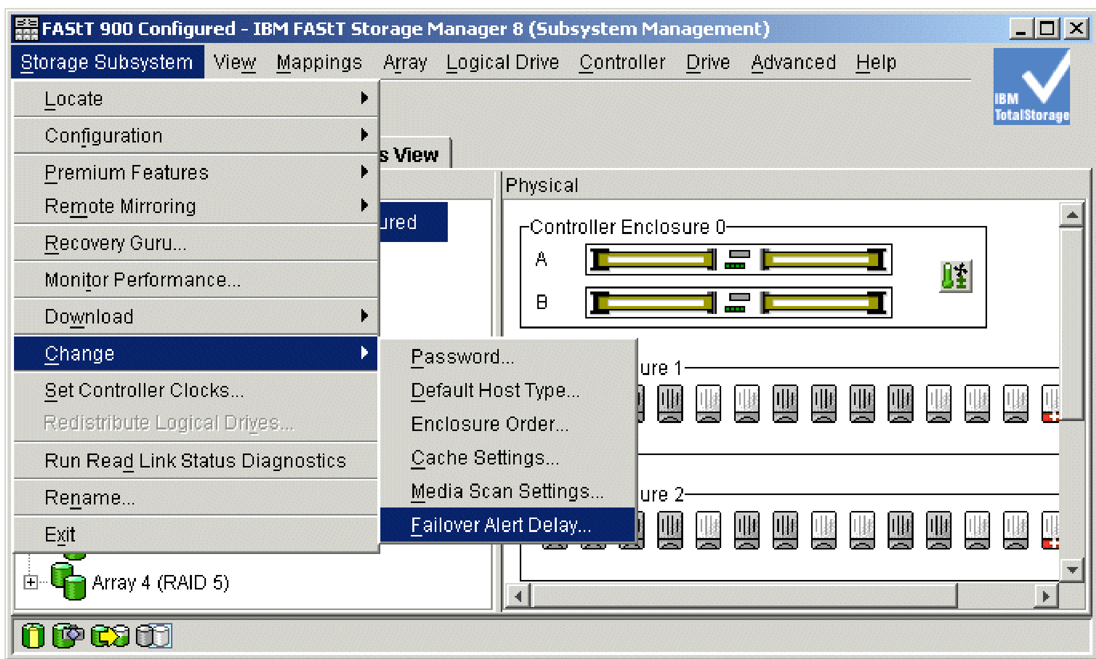

Figure 2-5 Select Failover Alert Delay

The Failover Alert Delay dialog is displayed. Enter the desired delay interval in minutes and select **OK** (see Figure 2-6 on page 14).

Important:

- The failover alert delay option operates at the storage subsystem level, so one setting applies to all logical drives.
- The failover alert delay option is reported in minutes in the storage subsystem profile as a storage subsystem property.
- The default failover alert delay interval is five minutes. The delay period can be set within a range of zero to 60 minutes. Setting the alert delay to a value of zero results in instant

notification of a logical drive not on the preferred path. A value of zero does not mean alert notification is disabled.

- The failover alert delay is activated after controller start-of-day completes to determine if all logical drives were restored during the start-of-day operation. Thus, the earliest that the not-on-preferred path alert will be generated is after boot up and the configured failover alert delay.

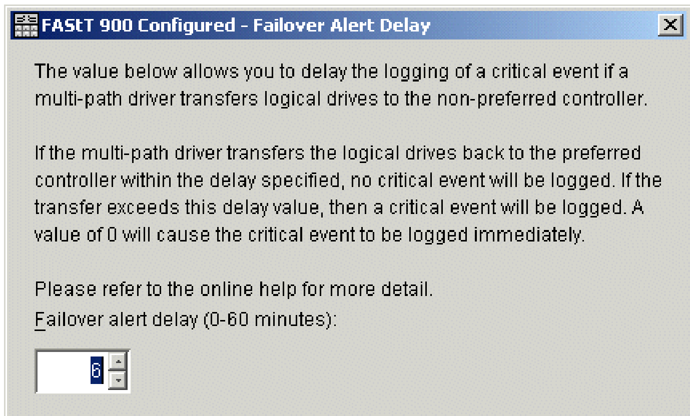

Figure 2-6 Set Failover Alert Delay

2.1.6 Recovery from intermittent drive path errors

This feature improves data availability by having the storage subsystem controller preemptively switch drive I/O operations from the preferred controller to the alternate controller when intermittent errors occur that prevent an I/O operation to a drive from being successfully completed.

Persistent intermittent errors are defined as errors that are typically recoverable, but which occur at a high enough rate, so as to be effectively unrecoverable. Previous versions of Storage Manager 8.4 were handling such errors by failing the I/O operation without resorting to a retry on the alternate path. This sometimes resulted in the failing of a good drive that was fully accessible on the alternate path.

With SM 8.4, the method of handling persistent intermittent errors is as follows: When the controller is trying to perform an I/O operation to a drive and encounters persistent intermittent errors to the extent that the retry limit is reached, the controller will mark the path degraded, but will not immediately fail the I/O. Instead, it will try to perform the I/O using the drive alternate path. If the I/O succeeds on the alternate path, then the drive is not failed, and the controller continues to use that path exclusively until failback occurs. If the I/O fails on the alternate path, then that path is also marked degraded and the I/O operation is failed. When both paths are degraded, the controller will make an arbitrary selection of which path to use for new I/Os.

The Storage Manager 8.4 GUI will show a needs-attention condition at the storage subsystem level (in both the Enterprise Management window and the Subsystem Management window).

There is a Recovery Guru procedure for the "path degraded due to intermittent errors" condition. The recommended recovery action of this procedure is to use the read-link-status data collection feature to identify the component in error.

Finally, the following is logged:

- Failover causes the storage array controller to post a critical Major Event Log (MEL) event. Because an I/O path problem can cause failover on multiple drives, the controller tries to consolidate the multiple drive failovers into one critical MEL event that lists all of the affected drives. In addition to the drive list, the MEL data includes (1) the identity of the path that failed, (2) the identity of the path that will be used in its place, or, if applicable, an indication that no alternate path is available.
- If the failover attempt resulted in the alternate path getting marked as degraded, a separate critical MEL event will be logged for that condition.
- As always, critical MEL events trigger alert notifications.
- Path failback due to loop initialization is recorded in the MEL as an informational event. Path failback due to controller reset or power cycle does not generate a MEL entry.

2.1.7 Delay in reporting loss of drive path redundancy

Previous versions of Storage Manager already indicated loss of drive path redundancy by raising a needs-attention condition.

In Storage Manager 8.4, the controller also reports a critical MEL event and a recovery guru action for this condition. However, because the conditions that cause the controller to detect loss of drive path redundancy can be transient in nature, the controller will not report the critical MEL event unless the condition persists for ten seconds. If, ten seconds after the first detection of the condition, the condition is no longer present, no event is logged.

The delay in reporting the loss of drive path redundancy condition also serves to minimize the number of critical events that are logged. If multiple occurrences of the loss of drive path redundancy are detected within ten seconds of the first one, only one event is logged.

When path redundancy is restored, the controller generates an informational MEL event indicating that fact. Again, due to the sometimes-transient nature of the condition, the controller does not log the event as soon as the event is detected, but delays for ten seconds to increase confidence that the state change is stable.

2.1.8 User control of network parameters

Using the SM8.4 client graphical interface users are now able to modify certain network parameter settings for each storage controller. In the past, some of these changes could only be made through the FAStT serial interface (see “Network setup of the controllers” on page 135).

The modifiable parameters are:

- Controller IP address
- Gateway IP address
- Network submask address
- bootp enabled or disabled
- Rlogin capability enabled or disabled

Changes made to these parameters go into effect immediately, without any need for a controller reboot or reset.

They are persistent; they are saved in both NVSRAM and DACstore, and will remain in effect across controller reboots and resets, until subsequently modified by the user.

They are also automatically propagated to a replacement controller.

Starting from the Subsystem Management window, select a controller in the Physical View, then select the **Controller -> Change -> Network Configuration** pull-down menu option (or **Change -> Network Configuration** from the right-mouse pop-up menu).

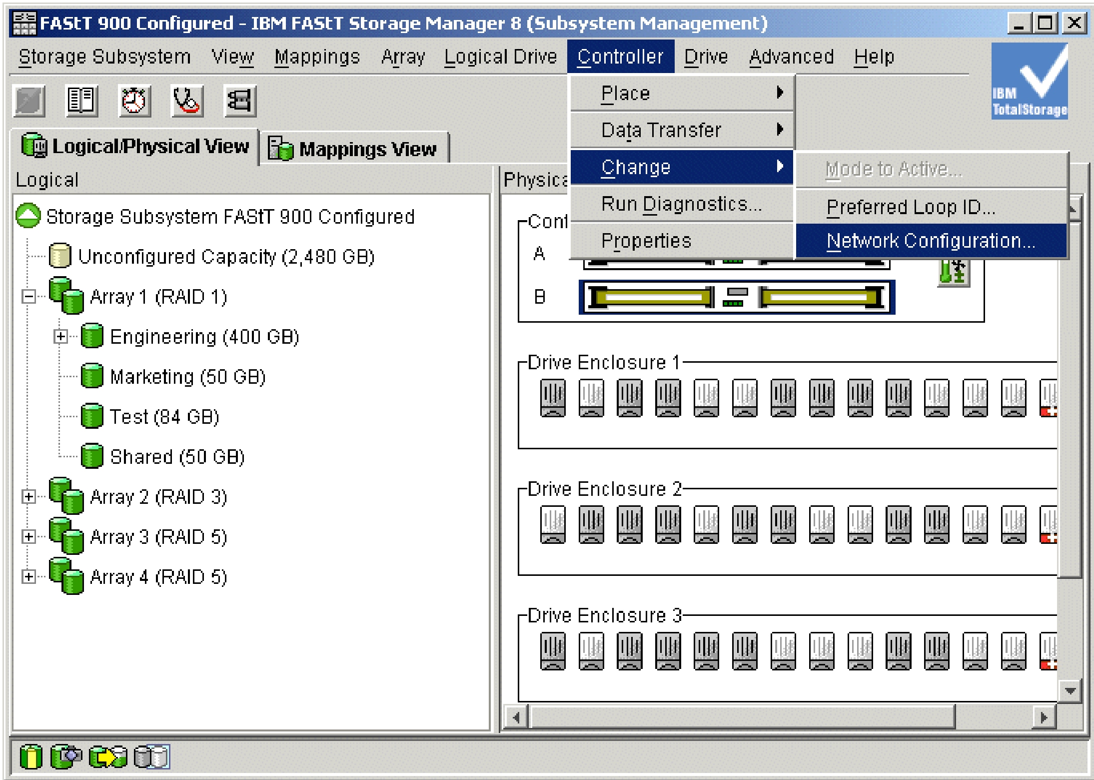

Figure 2-7 Change network parameters

The Change Network Configuration dialog is displayed (see Figure 2-8 on page 17). A separate tab is provided for each controller that displays the Ethernet port, the MAC address, and the host name. The dialog also displays the radio buttons for setting the IP address or enabling BOOTP.

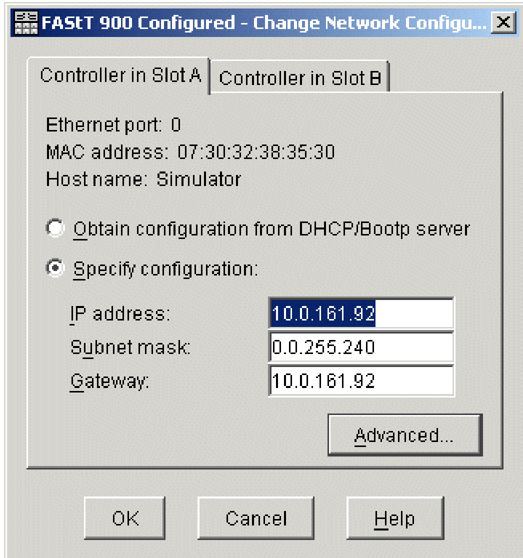

Figure 2-8 Change network parameters dialog

Select one of the two radio buttons. If you use a DHCP server, select the radio button **Obtain IP Address from DHCP/BOOTP Server**; otherwise select **Specify IP Address** and enter the controller's IP address, including the gateway address and network subnet mask.

Important: If you change your IP address, you will lose the management path to the device. You must remove the device from the Enterprise Management window and then add it back to the EMW by selecting **Edit -> Add device** and typing in the new IP address.

A confirmation dialog is displayed when you press the **OK** button.

When BOOTP is enabled, the controller will query the bootp server for the network settings to use. Any parameters received from the bootp server will override the saved parameters. Parameters received from a bootp server are used only while the controller remains operational and until the next bootp request is issued. These parameters are not saved in NVSRAM. When the controller goes through the next initialization sequence, because of a power cycle or reset, it will query the bootp server again to get its network parameters. Bootp is also invoked when a network link is reestablished, that is, when the network cable is connected to the controller. These received parameters may differ on each query, depending on the how the bootp server operates. Should the bootp request fail, the controller reverts back on the default fixed IP addresses.

Selecting **Advanced** allows you to enable remote login. Remote login should only be enabled to allow technical support access to the controller.

Attention: Leaving remote login enabled during normal operation poses a potential security risk. If you enable remote login for technical support to access your system, disable it when the session is completed to prevent unauthorized access to the system. Disabling remote login will terminate any current remote login sessions, and any new remote login requests will be rejected.

2.1.9 Capacity and performance improvements

Storage Manager 8.4 also introduces a series of capacity and performance improvements.

Increased drive queue depth

Each controller in the storage subsystem manages Input/Output (I/O) operations or requests for each drive. The term "queue depth" refers to how many outstanding I/O requests can be sent to a given drive. A higher queue depth increases the performance of applications with small-block, random I/O activity as it allows to boost the IOPS of the drive.

The queue depth is now up to a maximum of 16 for the FAStT600 Turbo and FAST900 (it remains at 8 for the FAStT700).

The queue depth limit for a given drive is reduced if the sum of the transfer lengths for outstanding I/O requests to that drive exceeds a pre-determined threshold. The threshold and reduced queue depth are governed by the upper limit on queue depth for that platform. If the upper limit for the drive queue depth is 8, then the transfer length threshold is 128KB, and the reduced limit is 2. If the upper limit for the drive queue depth if 16, then the transfer length threshold is 256KB, and the reduced limit is 4.

Increase command queue depth

Each controller in a storage subsystem can queue a minimum of 2048 host commands. The upper limit for a given host port is 2048 concurrent commands, subject to the overall controller limit. If the controller cannot queue any additional host commands for a given host port, incoming I/O operations for that port will be terminated with Queue Full status.

Increased number of Fibre Channel logins

The number of Fibre Channel logins per controller port is now equal to the maximum of host ports that can be defined (a maximum of 512 for the FASt700 and FAStT900).

This increase enables more host-to-storage I/O paths, which is critical for storage consolidation environments and to ensure there is always access to all defined hosts.

Increased large I/O size

If the transfer length specified for a host read or write operation exceeds a pre-determined size, the controller may break the I/O operation down into smaller, more manageable, steps. This pre-determined size is referred to as the large I/O size. In SM 8.4 the large I/O size has increased from 512 K to 2 MB for all logical drives and all controllers. This results in true bandwidth improvements by allowing full-stripe writes for larger segment sizes.

Arrays greater than two TB

Previous versions of Storage Manager supported up to 30 disk drives per array but with a two terabytes (TB) boundary for the overall array capacity. In other words, using disks of 146 GB authorized a maximum of only 14 disks per array or (28 disks of a 73 GB capacity each).

In 8.4, the limit of 30 disk drives per array remains but not the two TB boundary.

The benefit of this improvement is improved IOPS performance since you can have more spindles per logical drive (note that the maximum logical drive size remains at 2TB).

Cache-mirroring-enabled performance

A firmware enhancement that reduces cache traffic improves the performance when cache mirroring is enabled. Measurements have shown a 20 to 40 percent improvement on the FAStT700 and FAStT900 when compared to SM 8.3.

In addition, for the FAStT600 Turbo, an optimization of the RAID controllers can increase the overall performance by up to 70 percent when compared to the FAStT600 with SM8.33.

512K segment size

It is now possible to select a 512K segment size through the "createVolume" script engine command (CLI).

Note that using the SM GUI client you can only select the standard sizes (up to 256 K).

Disk drive capacity limitation

If a storage device is discovered that has a capacity greater than 2 TB, the drive will be "locked out" by the storage subsystem. Drives placed in this state will be uniquely represented to the user as failed, uncertified drives. The existing critical MEL event and "needs attention" condition associated with controller-failed drives will also be reported. Because the drive in question will be unusable by the storage subsystem, the only recovery action is to remove the drive from the system.

2.1.10 Other enhancements

We have grouped here other minor enhancements in Storage Manager 8.4.

Out-of-sync controller clocks

This feature provides an easy way for the user to know when the controller clocks are out-of-sync with respect to clock(s) on the host computer(s), and synchronize the storage server clocks with the computer that is running the host software.

Every time the user launches a Subsystem Management window, there is a time check on each controller. If either controller's time is out of sync by more than a set amount of time, a dialog is displayed that shows the time for each controller and the local computer's time. The dialog allows the user to synchronize the controller clocks with the local computer's time. The allowable variance between the local computer's time and the controller's time is defined in a properties file. The "out-of-box" setting is 5 minutes. A value of zero disables this feature.

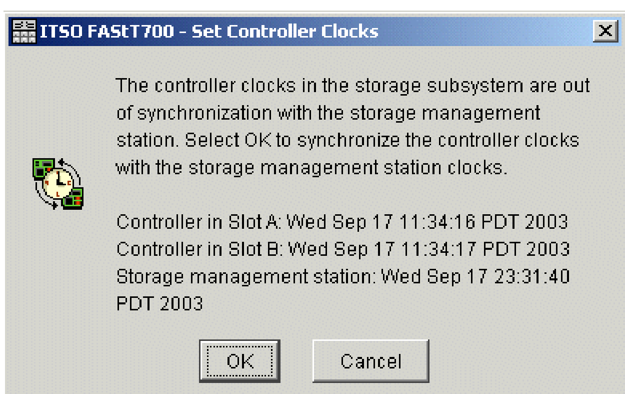

Figure 2-9 Set controller clocks

From the CLI, each time the user runs the **`show storageSubsystem healthStatus`** command on a particular storage subsystem, it will, in addition to its other responsibilities, check the time on each controller. If either controller's time is out of sync by more than a set amount of time, it will display a warning message that lists the time for each controller and the local computer's time.

Host port Information refresh

Up-to-date information about which host ports are known to the storage subsystem (that is, have logged in) can be obtained by clicking a refresh button. The refresh button is provided in two dialogs: The Define Host Port dialog and the Show All Host Port Information dialog, shown in Figure 2-10.

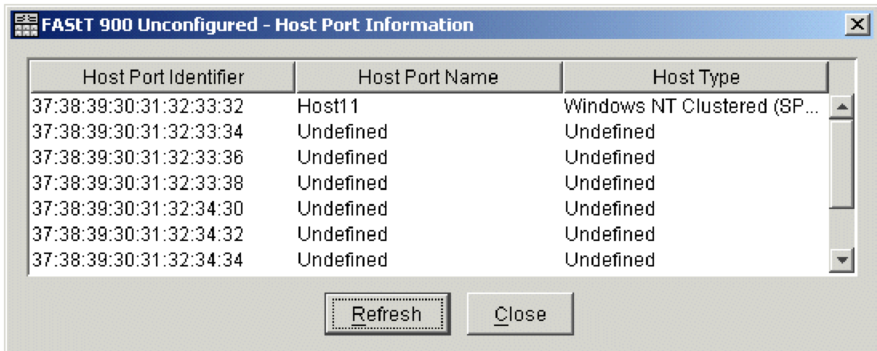

Figure 2-10 Host Port Information dialog

Additionally, up-to-date information about host ports can be obtained by executing the new **`show hostPorts`** script command.

Preferred owner information

The identity of the preferred owning controller is included with the volume information.

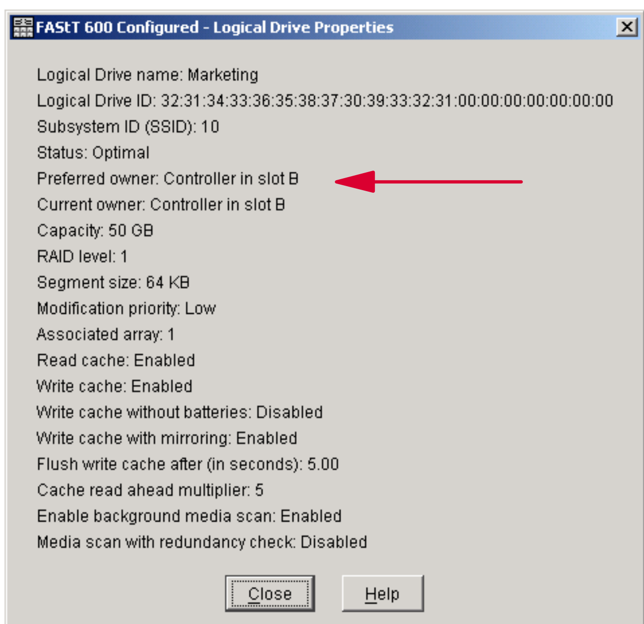

Figure 2-11 Preferred owner information

This same information continues to be available by following method: Right-click **Volume -> Change -> preferred-path**.

MEL viewer settings

The MEL viewer dialog settings are persistent after exiting the Subsystem Management Window, or after exiting the SM 8.4 altogether.

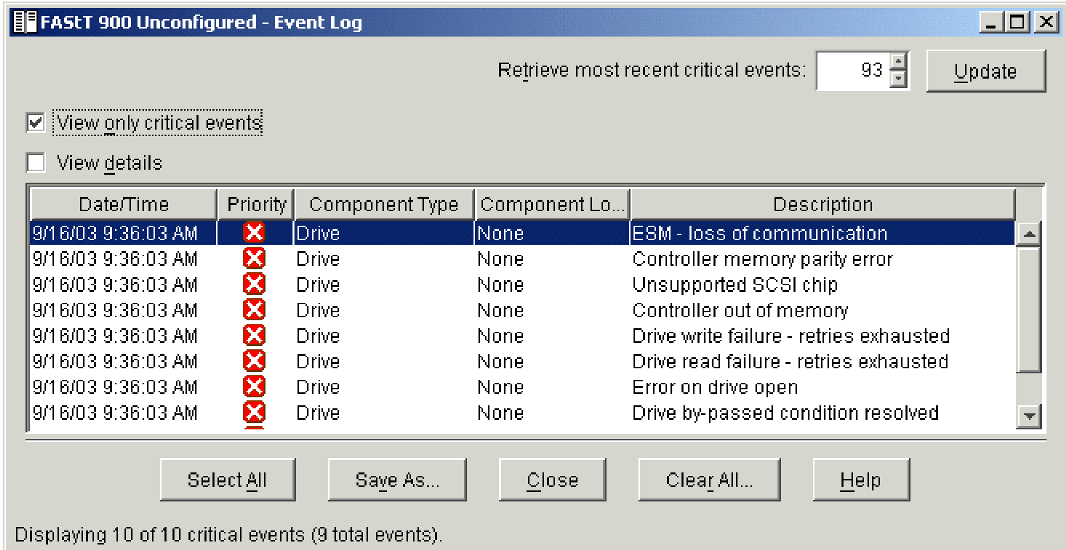

Figure 2-12 Persistent MEL viewer settings

The following affected settings are retained:

- Critical vs. non-critical checkbox
- View details checkbox
- Number of events to retrieve spinner box

Certain changes have also been introduced in the way the spinner box behaves:

- It no longer allows for different values, depending on whether all events or only critical events are to be displayed. Whatever value is entered into the spinner box applies to both critical-only and critical + informational viewing modes.
- It has a new default value of 100. (Previously the default was 50 for critical and 200 for critical + informational.)
- It increments by one for both critical and critical + informational events. (Previously it would increment by five for critical events and 25 for critical + informational events.)

Multiple logical drives deletion

When selecting a logical drive for deletion in the Subsystem Management window, the user is presented with a dialog, within which is a list box enumerating all delete-able logical drives on the storage subsystem from the time this interface was invoked.

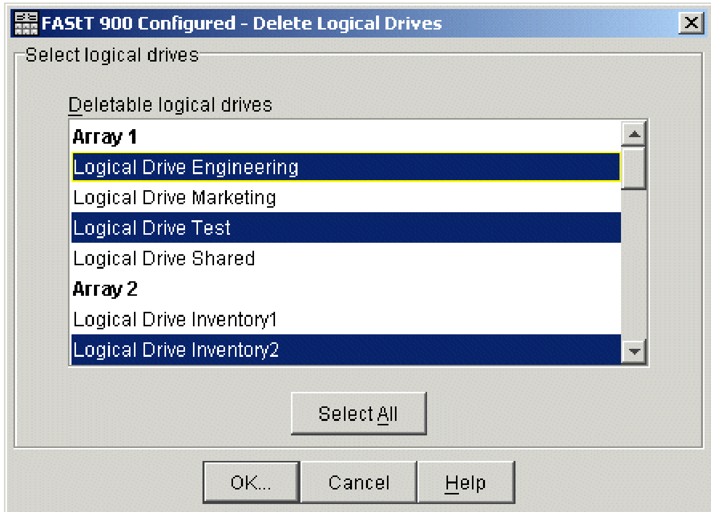

Figure 2-13 Multi logical drive deletion dialog

The user may select a single logical drive, several, or all logical drives in the list. The operation requires user confirmation before the delete will be performed.

Note that the list of delete-able logical drives does not include all logical drives in the array if some are volumes considered non-delete-able due to restricted states.

Recovery profile

This is an append-only file that can be used by IBM technical support for troubleshooting FAStT issues.

Optional controller batteries

The storage controller recognizes a new NVSRAM bit setting, "controller batteries not present." If set, the controllers will not consider missing batteries to be an invalid configuration state and will not report this condition; that is, no needs-attention indicators, no Recovery Guru information, and no MEL events.

In cases where there is a battery present in the storage subsystem but the NVSRAM is set to not include the batteries, the physical battery will take precedence over the NVSRAM setting. In this case, batteries will be reported by the controller firmware and the host software. The reason behind this behavior is that if a battery is physically installed in a storage array, the user intended to use a battery and the NVSRAM setting is not correct.

The ability to set this option is not available to end customers in the field and must be set by trained IBM technical support personnel.

Additional warnings/states

Some additional warnings/states are;

- Added new battery statuses (“charging” or “not present”)
- Warning - out of sync clocks (Mechanisms to synchronize, display date and time also provided)

Authentication failure reporting and lock-out

To reduce the risk of and to provide better management of unauthorized access to the storage subsystem, the following behavior is implemented in Storage Manager 8.4:

- Log an information MEL event for all authentication failures that were not a result of using a NULL password (authentication is required for some commands only).
- Log a critical MEL event if the frequency of authentication failures reaches ten within ten minutes. Authentication failures caused by a NULL password will not be counted towards this limit.
- Once the frequency threshold for authentication failures has been exceeded, all commands requiring authentication received by either controller over the next 10 minutes are failed. This statement applies for all SM clients, even ones that are not the source of failed attempts, or ones that were previously authenticated. This new state of the controllers where they are failing all authentication requests is called lock-out.

The behavior outlined above applies regardless of the type of program that originated the command—either the SM GUI or CLI.

RLS diagnostics data

Automatic save of the Read Link Status (RLS) diagnostic data.

Windows installation enhancements

In Storage Manager 8.4, the program version information is included in the text displayed for the Storage Manager Client package in the Add/Remove Programs dialog box on all supported Windows platforms.

The windows installation program now supports client automatic deletion and installation; this feature addresses some awkwardness in the way the Windows installation worked in past releases. The behavior was such that, if there was a pre-existing Storage Manager package at the time of the installation, the pre-existing package would be removed, and the install would terminate at that point, rather than proceed to install the new package. The user had to invoke the installation a second time in order to complete the installation of the new package.

In SM 8.4, the Windows installation always proceeds with the installation of the new package even when a pre-existing version is detected. It is no longer necessary to invoke the installation twice when a pre-existing package is detected. When the install completes, the new package is installed and the pre-existing version is no longer present.

3

Hardware details

This chapter covers the latest additions to the FAStT family: The IBM TotalStorage FAStT600 and FAStT900. We discuss additional hardware components that are essential for a complete storage solution. These include the EXP700 external expansion enclosure, cabling considerations, Host Bus Adapters and SAN switches.

3.1 IBM TotalStorage FAStT600 and FAStT600 Turbo

The FAStT600 Storage Server is a 3U rack-mountable Fibre Channel RAID controller and disk drive enclosure. It targets the entry and midrange segment of the Fibre Channel storage market. A typical usage of the FAStT600 is a two-node cluster directly attached to it. It can support up to 56 Fibre Channel disk drives. The model number is1722. This is shown in Figure 3-1.

3.1.1 Features of the FAStT600 (base)

- Compact 3U rack-mountable enclosure.
- Dual 2 Gbps hot swappable RAID controllers with 512 MB of battery backed cache (256 MB per controller).
- Redundant, hot swappable power supplies and cooling.
- Support for RAID-0 (implemented as a RAID-10), 1, 3, and 5.
- Supports global hot spares.
- One expansion port per controller.

Figure 3-1 FAStT600 storage server

- Supports three IBM TotalStorage FAStT EXP700 Expansion Units (a license key is required); this allows for up to 56 disk drives, offering a potential capacity of 8.2TB.

> **Note:** FAStT600 only supports the EXP700. The FAStT600 does not support connection to the EXP500 as the FAStT600 is 2 Gb system only, and does not support 1 Gb devices. (FAStT600 Turbo auto senses to connect to 1Gb or 2Gb.)

- There are four host ports, two for each controller; this enables a cluster solution without the use of a switch.
- Supports both short and long wave Fibre Channel host attachment, allowing distances up to 10 km (6.2 miles).
- Supports one storage partition in standard configuration. There is an option to expand up to 4, 8 or 16 storage partitions.
- Standard configuration only supports Windows connection; Linux, Netware, and UNIX support needs an orderable option.

The FAStT600 can be upgraded to the FAStT600 Turbo.

3.1.2 FAStT600 Turbo Highlights

The FAStT600 Turbo is an upgraded version of the bas FAStT600 model and presents the following characteristics and features in comparison:

- Up to 70 percent performance improvement with new Storage Manager v8.4 that ships with Turbo
- Higher scalability over base FAStT600, scalable to 16.4TB by attaching up to seven EXP700s
- Increased cache from 256MB per controller on base FAStT600 to 1GB per controller on Turbo
- Autonomic functions such as Dynamic Volume Expansion and Dynamic Capacity Addition, allowing unused storage to be brought online without stopping operations, and FAStT Service Alert, which is capable of automatically alerting IBM if a problem occurs
- Host interface on base FAStT600 is 2Gb, Turbo auto senses to connect to 1Gb or 2Gb
- Up to 64 storage partitions
- VolumeCopy (premium feature)

3.1.3 FAStT600 supported operating systems and infrastructures

The following is included for illustration purposes. The range of supported platforms is continually increasing. For the latest support information please refer to the following Web site:

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

Supported operating systems

The FAStT600 supports the following operating systems:

- Windows 2000 Server (with Service Pack 3)
- Windows 2000 Advanced Server (with Service Pack 3)
- Windows Server 2003 (32bit - 64bit support later)
- Windows NT 4.0 (with service pack 6a)
- VMware ESX 1.5.2
- Novell NetWare 5.1 and 6.0, 6.5 and Cluster
- SuSE Linux Enterprise Server 8
- Red Hat Advanced Server 2.1
- IBM AIX 4.3.3, 5.1, and 5.2
- HP UX 11,11i
- Sun Solaris 2.6, 7 and 8

Clustering support

For clustering support:

- Microsoft Cluster Service
- Novell Cluster Service
- IBM HACMP
- PSS, GPFS/RVSD support
- Solaris 8 with Veritas Cluster Server 1.3, 2.0 and 3.5
- Veritas Volume manager 3.1, 3.2 and 3.5 (Solaris)
- HP MC/Service Guard

- GPFS for Linux is *not* supported on the FAStT600

Supported server platforms

Supported server platforms are:

- IBM xSeries and pSeries
- HP (Intel)
- Dell (selected models)
- HP and Sun UNIX based
- Compaq (Intel selected model)

Supported SAN switches

Supported SAN switches are:

- IBM TotalStorage SAN Switch:F08/F16/F32/M12
- Most McData, Cisco and Inrange switches and directors

3.1.4 Front view

From the front panel (Figure 3-1 on page 26) you can access up to fourteen hot swappable Fibre Channel hard disk drives and view the two status LEDs on top:

- Green Power LED

  This LED indicates the DC power status is OK.

- Amber General-System-Fault LED

  When a storage server component fails (such as a disk drive, fan, or power supply), this LED will be on.

Each disk drive also has its own status LEDs:

- Green Activity LED

  When this LED is flashing as quickly as about three times per second, it indicates some disk drive activity. When it is steadily on, it means the disk drive is OK and properly installed. When it is flashing as slowly as once per second, it means that the RAID controller does not detect the HDD. If the drive is off, it means the FAStT Storage Server is not powered up, or drive has not spun up.

- Amber Disk-Drive-Fault LED

  If a disk drive fails, this LED will be steadily *on*. While the drive is being identified, rebuilding or rebuilt, it flashes. In normal operation it is off.

The number of HDDs in the enclosure can vary from 1–14. To ensure proper cooling and avoid Electro Magnetic Interference (EMI), all 14 bays should be occupied either by a HDD or by a filler module.

The Blue LED situated right to the General System error LED, is used as a locate option from Storage Manager to locate and identify a drive enclosure in the Storage Subsystem.

3.1.5 Rear view

You access the RAID controllers, power supplies, and fans from the rear side of the FAStT600 unit (see Figure 3-2 on page 29). Two hot swappable power supplies and fan units provide redundancy and, therefore, offer a high availability level. If a fan unit fails you should not remove it from the storage server until the replacement is available. This is because the cooling airflow will be less than optimal when a fan unit is missing.

Figure 3-2 FAStT600 storage server rear view

The same is true for power supplies: Do not remove a failed power supply unit before you have a replacement available, or else the cooling will not be efficient.

The RAID controllers themselves are also redundant and hot swappable.

3.1.6 RAID controller

The FAStT600 has two RAID controllers; each RAID controller contains an Intel 80321 processor, running at 600 MHz and using 256 MB of Error Checking and Correcting (ECC) cache (1GB of ECC cache on the FAStT600 Turbo).

RAID controller connectors

The RAID controller unit (shown in Figure 3-3) has three ports for SFP modules, which are used to connect Fibre Channel cables. Two SFPs can connect to host systems; the third SFP is used to connect to EXP700 disk expansion units.

Each RAID controller reserves 128 MB of cache memory. A backup battery protects the cache memory. If the power fails, the battery will maintain the data in cache for at least 72 hours. The life cycle of the battery is three years. For more information, see “Cache backup battery” on page 31.

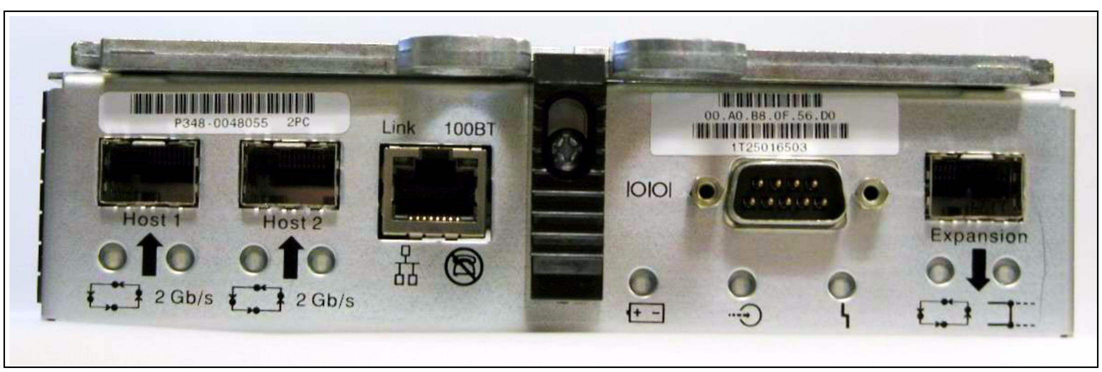

Figure 3-3 FAStT600 storage server controller

The following connectors are on each controller as shown in Figure 3-3.

- Host port SFP connection

  There are two host ports available in each controller (see Figure 3-4 on page 30). They are on two independent Fibre Channel loops. Each offers a fabric connection to the host. You can install either a short (maximum 500m) or long (maximum 10 km) wave SFP into the slot. This design makes a cluster without the use of a switch or hub possible. Please refer to “Host port connection” on page 32 for connection topology.

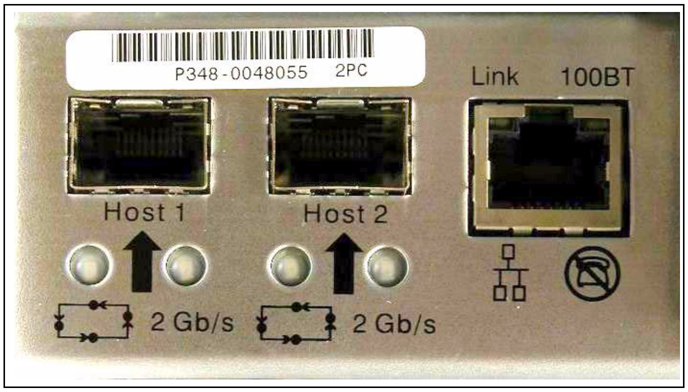

Figure 3-4 FAStT600 controller host side

- Expansion port SFP connection (see Figure 3-5)

  This connector provides storage expandability. In the standard configuration, the expansion port is disabled. You will need a license key to enable it to connect three EXP700s (and up to 7 EXP700 can be attached to the FAStT600 Turbo). The connection is supported only with short wave SFP.

Figure 3-5 FAStT600 expansion port

- RJ-45 Ethernet connector

  This connector is for an RJ-45 10BASE-T or 100BASE-T Ethernet connection. The intention of using this connection is for inband management using the Storage Manager client.

  Integrated Ethernet Chip’s MAC address is printed on the label shown in Figure 3-5.

- Serial port

  This serial port is used for management and diagnostics purpose. You can use a PC with a terminal emulation utility such as Hyper Terminal to access the command set.

> **Attention:** Managing the FAStT Storage Server through the serial interface has potential risks. With certain commands you can initialize the RAID controller, and therefore lose all your data. You should only use this interface when instructed to do so by IBM Level 3 Technical Support. In this publication, we provide only limited instructions on how to use the serial interface. If it is required to access the serial interface you will receive all instructions from IBM Technical Support.

RAID controller LEDs

There are nine indicator LEDs on FAStT600's RAID controller, which are shown in Figure 3-6.

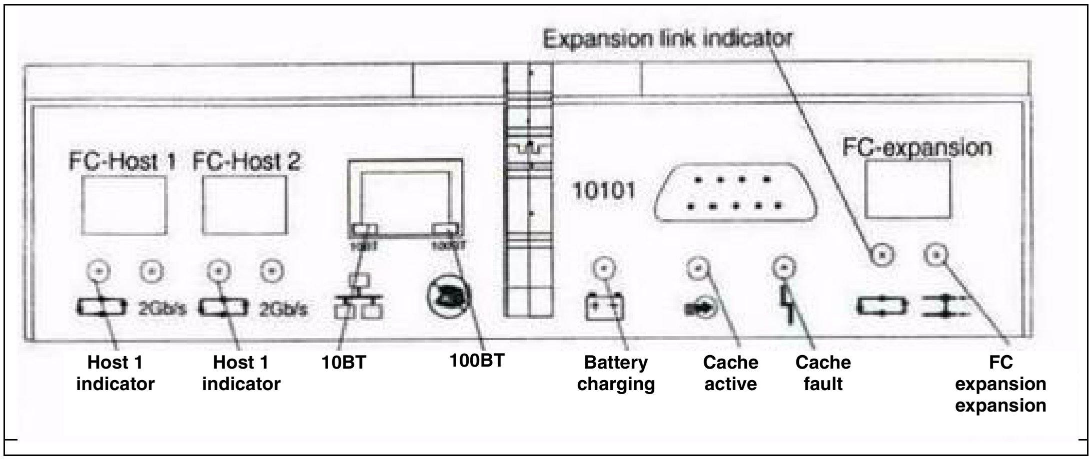

Figure 3-6 FAStT600 RAID controller indication LED

Several diagnostic LEDs are present on the RAID controller unit:

- Host loop LED (green)

  This LED should be on, which means that the host connection loop is good. If it is off, the following problems may have occurred:

  – The host loop is down, not turned on, or not connected.
  – A SFP has failed, or the host port is not occupied.
  – The RAID controller circuitry has failed, or the RAID controller has no power.

- Cache activity LED (green)

  This LED is on when the data is in cache. If it is off, one of the following situations has occurred:

  – There is no data in cache.
  – The cache option is not selected for the array.
  – The cache memory has failed, or the battery has failed.

- Battery charged LED (green)

  Normally, this LED should be on. If it is off, it indicates a battery fault. The LED blinks while the battery is charging or performing a self-test.

- Expansion port bypass LED (amber)

  The LED will be on if nothing is plugged into the expansion port, or the expansion is powering off.

- Expansion loop link LED (green)

  Normally on when the drive side Fibre Channel loop is operating normally.

- Controller fault LED (amber)

  During normal operation, this LED should be off. It is on during the power up and power down sequence, or if the controller fails.

3.1.7 Cache backup battery

Each RAID controller unit contains its own backup battery. When the battery expiration date is near, you should replace it. The battery is installed inside the RAID controller, so in order to replace it, you have to remove the RAID controller unit from the FAStT600. When you are using the FAStT600 in a redundant configuration, you can avoid downtime by keeping the FAStT600 up and running during the battery replacement. Be aware that you have to complete the replacement procedure within 30 minutes. This time should always be sufficient. When the procedure takes more than 30 minutes, overheating might occur.

Note: The FAStT600 is different from the FAStT700 and FAStT900 in this aspect. FAStT700 and 900 use a separate battery module; this can be replaced while both controllers remain operational.

3.1.8 FAStT600 indicator lights

The FAStT600 Storage Server indicator lights display the status of the FAStT600 Storage Server and its components. The green indicator light means normal operation status; amber indicator lights mean a possible failure and an event logged in Storage Manager.

It is important to check all the indicator lights on the front and back of the controller. The indicator lights might blink intermittently when power is just turned on. Wait until the FAStT600 Storage Server completes its power up sequence before checking for faults. It can take up to 15 minutes for the battery to complete its self-test, and up to 24 hours to be fully charged, particularly after an unexpected power loss of more than a few minutes.

The indicator lights for the components of the FAStT600 Storage Server are displayed in Figure 3-7. Each power supply contains two LEDs:

- DC power good LED (green), which indicates the DC power is OK and is normally on.
- Fault LED (amber), which indicates a failure if it is on.
- There is a fan fault LED (amber) on the fan module. This LED is normally off. If it is on, it indicates a failure of the fan in the module.

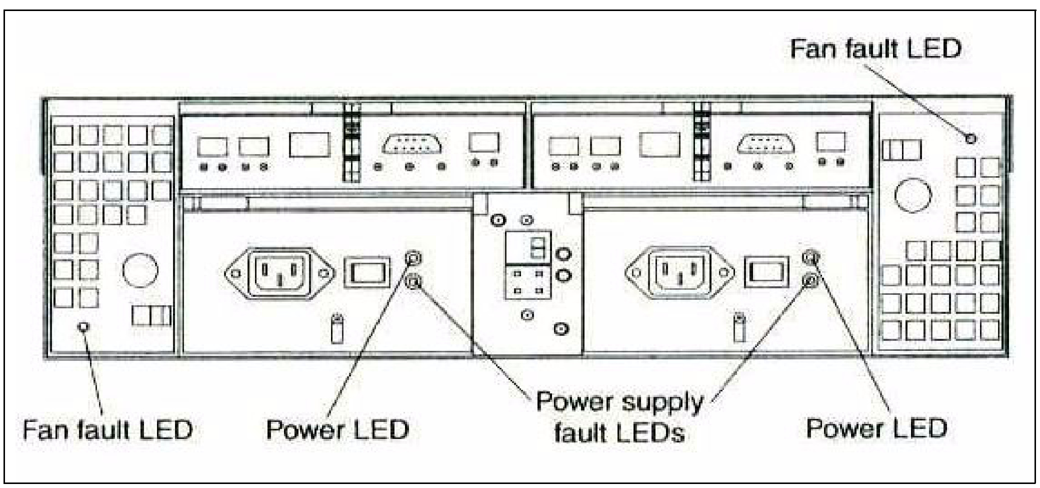

Figure 3-7 FAStT600 storage server light indicator

3.1.9 Host port connection

You can direct attach the FAStT600 to a single host with two HBAs. It offers fault tolerance on both HBAs and FAStT controllers. If one HBA were to go faulty, the host still has access to the storage server via the other HBA in the host, making a direct connection to the second host side connection on the FAStT Controller A or B (as there are dual host connections per controller). At the same time, you can get higher performance, since the dual controller allows for load balancing. See the left side of the chart in Figure 3-8.

FAStT600 supports a dual node cluster without using a switch. This is shown in Figure 3-8.

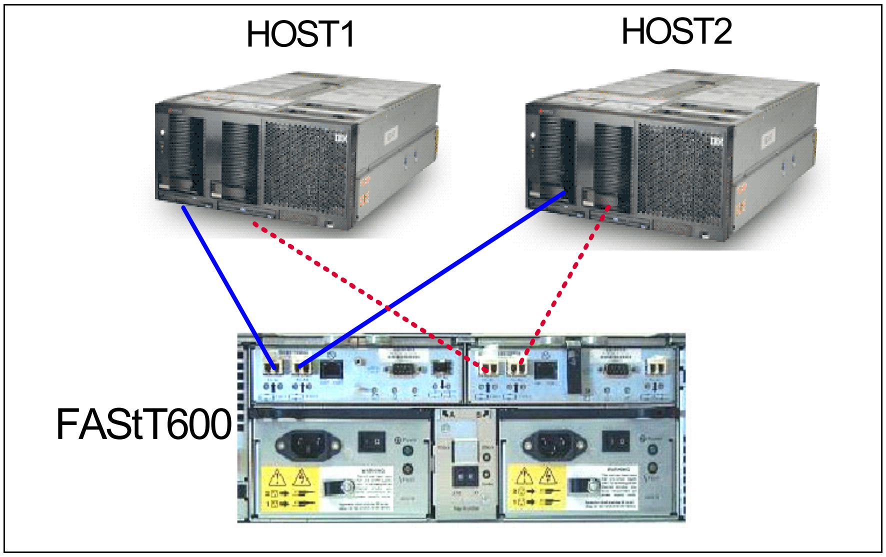

Figure 3-8 FAStT600 connection

You also can connect a switch (recommended) to a FAStT600, so that it will support additional hosts (see Figure 3-9).

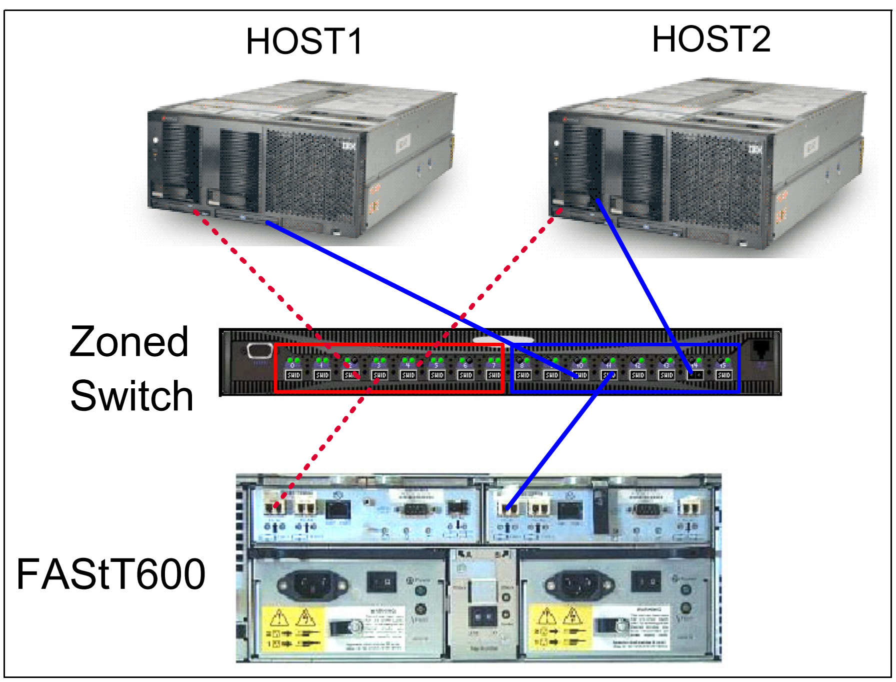

Figure 3-9 FAStT600 connection through a switch

Additionally, you can expand the capacity by attaching EXP700 expansion units (see Figure 3-10 on page 34).

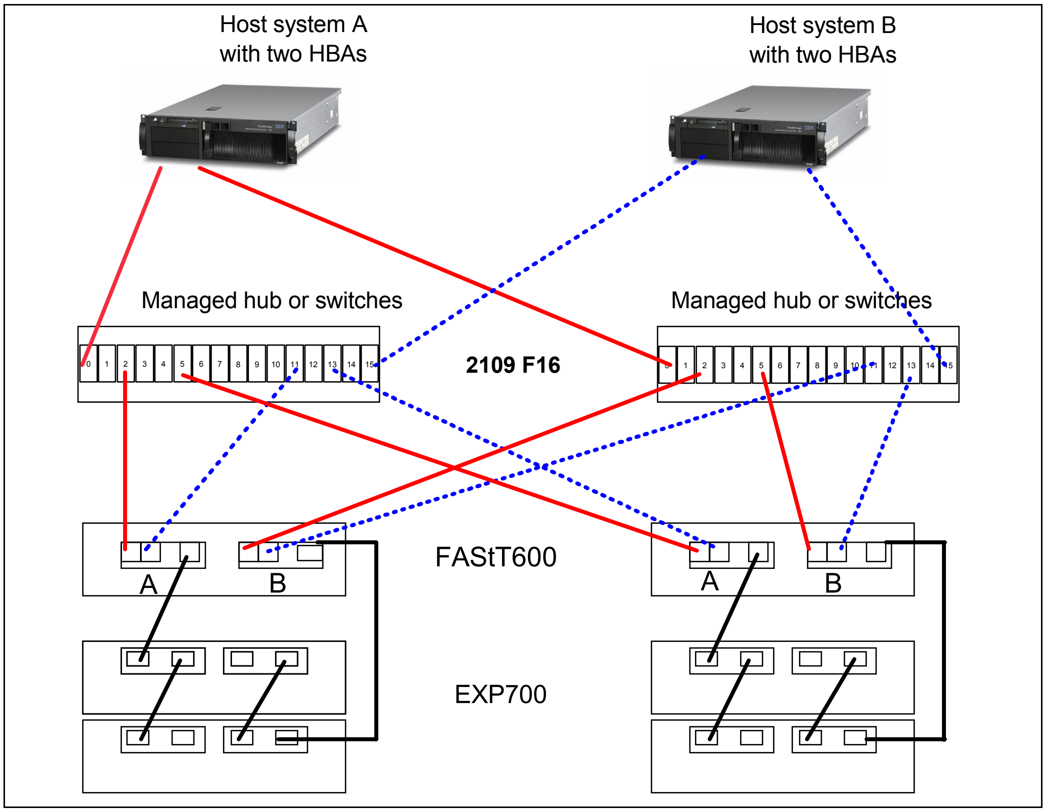

Figure 3-10 FAStT600 fault tolerance connection with switches

Note the importance of zoning as well (zones are differentiated by the line type: Continuous versus dotted line), such that any given HBA in a server can only see a single FAStT controller.

3.2 IBM TotalStorage FAStT900 server

The IBM TotalStorage FAStT900 Storage Server expands the FAStT family's highly scalable offerings with improved performance capabilities. Built upon a new fourth generation Fibre Channel (FC) RAID controller, the FAStT900 brings high availability, advanced functionality, scalable capacity, and connectivity to a wide range of storage area network (SAN) applications in mission-critical enterprise networks. The FAStT900 is designed to support applications running on IBM @server pSeries, and xSeries servers, as well as servers from Sun, Hewlett Packard, and other Intel-based providers.

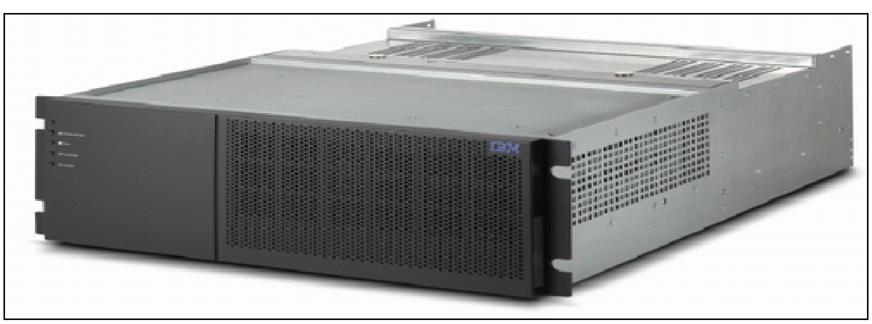

Figure 3-11 FAStT900 storage server

Features of the FAStT900

The FAStT900 is a high-performance, highly available, Fibre Channel based storage product for demanding applications in the UNIX and Intel SAN environments.

Machine type 1742, models: 90U, 90X:

- Compact 4U rack-mountable enclosure.
- Uses an Intel Pentium III processor in each RAID controller unit.
- New shared memory bus controller/PCI bridge/RAID assist engine - 1.6 GB/s bandwidth.

> **Note:** Throughput rates are heavily dependent on the internal controller bandwidth. FAStT900's design just addressed this point.

- Dual, redundant 2 Gb RAID controllers with 2 GB of Rambus cache memory (1 GB per RAID controller). The data in the cache is protected by battery backup for at least seven days.
- Cache backup battery is hot-swappable.
- Redundant, hot-swappable power supplies and fans.
- Supports RAID 0, 1 (implemented as RAID 10), 3, and 5.
- Supports global hot spares.
- Supports four host-side Mini-Hubs for two controllers. Each controller is responsible for two Mini-Hubs. Each Mini-Hub has two ports. Four Host Ports (two on each controller) provides a cluster solution without using a switch. Two host-side Mini-Hubs are standard.
- Supports four host Fibre Channel loops, eight host connections with selectable 1 Gb or 2 Gb speed.
- Supports four drive-side Mini-Hubs for two controllers. Supports four loops in Expansion connection. Can be connected in a fully redundant way to avoid any single point of failure. Standard configuration has two drive-side Mini-Hubs.
- Supports connecting to 16 EXP700, up to 224 disk drives, offering a capacity of over 32.8TB using flexible combinations of 18.2, 36.4, 73.4 and 146.8 GB drives.
- Supports both short (maximum of 500m) and long (maximum of 10 Km) wave Fibre Channel host attachment, allowing distances up to 10 km (6.2 miles).
- Supports multiple heterogeneous servers and operating systems.
- Includes IBM FAStT Storage Manager V8.3 with up to 16 storage partitions and an option to expand up to 64 storage partitions. When upgraded to Storage Manager 8.4, the FAStT900 sees a performance improvement between 20 percent and 40 percent when write Cache Mirroring is on.
- Supports FlashCopy, Remote Volume Mirroring (RVM) and VolumeCopy (requires Storage Manager 8.4).
- Integrated ESM firmware, hard drive microcode download function; ESM firmware can be downloaded when storage is receiving I/O.
- New service alert function. Controllers send an alert e-mail to the IBM Help Center when an error occurs (see Appendix B, “IBM FAStT Service Alert” on page 391).

Supported operating systems and infrastructures

The following is included for illustration purposes. The range of supported platforms is continually increasing. For the latest support information please refer to the following Web site:

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

Supported Operating Systems

Supported Operating Systems are:

- Windows NT 4.0 (with Service Pack 6a)
- Windows 2000 Server (with Service Pack 3)
- Windows 2000 Advanced Server (with Service Pack 3)
- Windows Server2003
- Novell NetWare 5.1 and 6.0 (Novell Netware .1 support planned availability 3Q 2003)
- SuSE Linux Enterprise Server 7, 8
- Red Hat Linux Advanced Server 2.1
- SuSE Linux Enterprise Server 7,8
- VMWARE ESX 1.5.2
- IBM AIX 4.3.3,5.1 and 5.2
- Sun Solaris 2.6,7,8 and HP-UX 11,11i

Clustering support

Clustering support is for:

- Microsoft Cluster Service
- NetWare Cluster Service
- IBM HACMP
- PSS, GPFS/ VSD support
- Solaris 8 with Veritas Cluster Server 1.3, 2.0 and 3.5
- Veritas Volume manager 3.1, 3.2 and 3.5 (Solaris)
- HP MC/Service Guard
- GPFS for Linux is *not* supported on the FAStT900

Supported server platforms

Supported server platforms are:

- IBM xSeries and pSeries
- HP (Intel)
- Dell (selected models)
- HP and Sun UNIX based
- Compaq (Intel selected model)

Supported SAN switches

Supported SAN switches are:

- IBM TotalStorage SAN Switch:F08/F16/S08/S16
- McData and Inrange switches and Directors

3.2.1 Front view

After removing the front bezel, we can access the following components (see Figure 3-13 on page 37).

Two RAID controllers

Each RAID controller blade contains diagnostic LEDs and a reset switch.

Controller fan module

Two fans are integrated in this module. Their role is to provide cooling to both RAID controller blades. FAStT900 diagnostic LEDs are present on the module, and there is an alarm switch that can turn the sound on and off.

Battery module

It is possible to hot-replace the battery module. During the replacement, the cache will temporarily lack battery protection. It is best to disable writing cache to avoid any possibility of data loss when changing the battery.

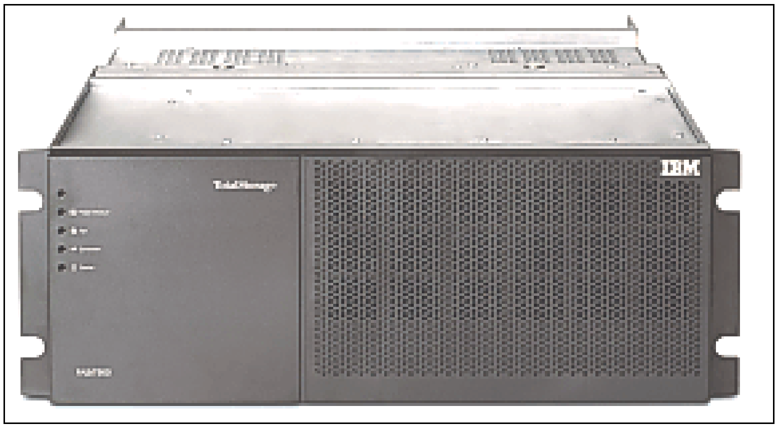

Figure 3-12 Front view FAStT900

Figure 3-13 shows the front view of the FAStT900 with the bezel detached.

Figure 3-13 Front view of the FAStT900 with bezel detached

3.2.2 Rear view

The rear view is shown in Figure 3-14 on page 38. Components accessible on this side of the FAStT900 include:

- Host connection: Mini-Hub ports with SFP module for host connectivity.
- Drive connection: Mini-Hub ports with SFP module for Fibre Channel expansion enclosure connectivity.
- Communication module has management ports and fan.

  There are two serial and two Ethernet ports on the communication module. The upper side serial port communicates with RAID Controller A, the other one is for Controller B. We can see letters A, B beside the appropriate port. This applies to the Ethernet ports. They are used for managing and troubleshooting purposes.

  Besides management ports, this module has two fans that provide redundant cooling to both power supplies.

► Two hot-pluggable and redundant power supplies

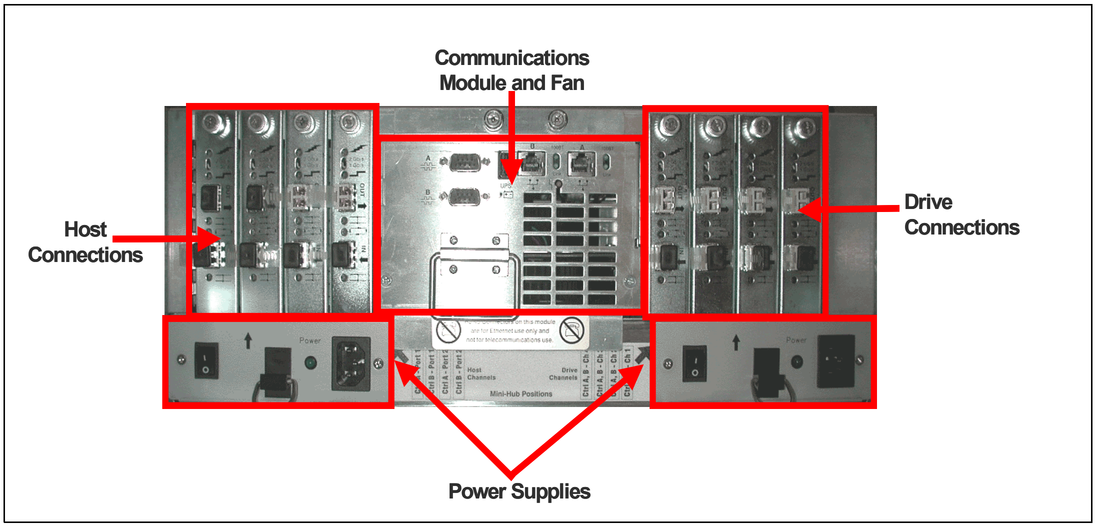

Figure 3-14 FAStT900 rear view

3.2.3 RAID controller

Each RAID controller contains an Intel Pentium III 850MHZ processor. Each Controller supports 1 GB of cache. A backup battery protects the cache memory. If the power fails, the battery will maintain the data in cache for at least seven days.

RAID controller LEDs and status LED advance

RAID controller LEDs' layout and meaning is the same as the FAStT700. Each RAID controller has ten indicator lights: One power, one fault, and eight status lights, as shown in Figure 3-15 on page 39.

Note: To check the RAID controller indicator lights, you should remove the FAStT900 Storage Server's front bezel. There are eight status lights that glow in various patterns, depending on the controller status. Always use the storage management software to identify the failure. If the controller is offline, the indicator will be lit. This does not indicate failure.

The LEDs are:

► Controller power LED (green)

  This LED is normally on to indicate that the controller is powered up.

► Fault controller LED (amber)

  This LED is normally off. If it is on, some failure happened on the controller.

► Heartbeat LED (green)

  This LED indicates the controller is working properly; it normally flashes for one second on and one second off.

► Status LEDs (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.

.

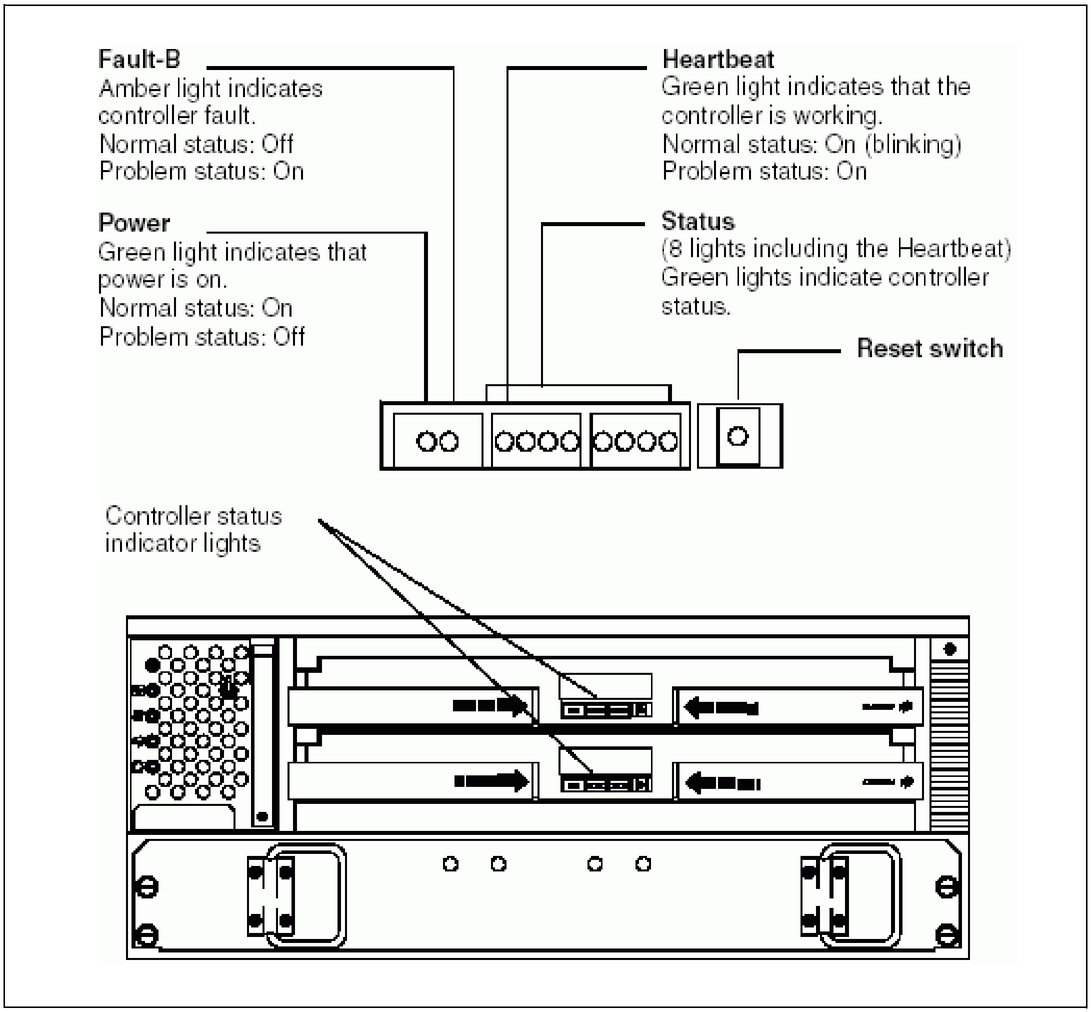

Figure 3-15 FAStT900 RAID controller indicator light

Status LED advance

The fault LED on indicates that the code or hardware module on the RAID controller is failing. The eight status LEDs are very useful for diagnostic purposes. The eight LEDs on and off stand for 1 and 0. It is a two-digit hex number, which corresponds to the different status. From Figure 3-16 on page 40, we can find out if only the Heartbeat LED is on, the binary form is 1000 0000, and in hex will be 0x80. If all the LEDs are off, the hex will be 0x00. If all the 8 LEDs are on, the hex will be 0xFF.

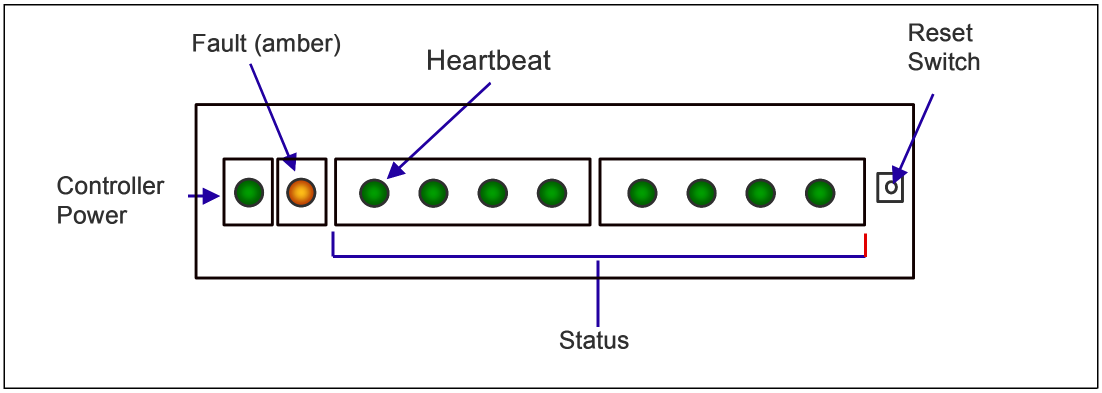

Figure 3-16 FAStT900 RAID controller status LED

0x00/0x80: Heartbeat LED is flashing (on and off) and the other LEDs are off. This means the RAID controller is active and working fine (see Figure 3-17).

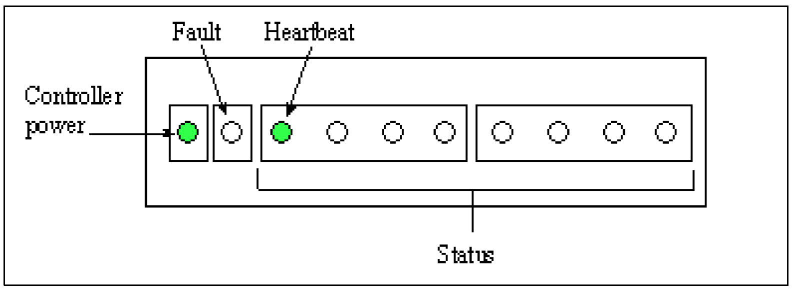

Figure 3-17 FAStT900 active indicator lights

0x6E/0xEE: RAID Controller is in passive status and there is no fault (see Figure 3-18).

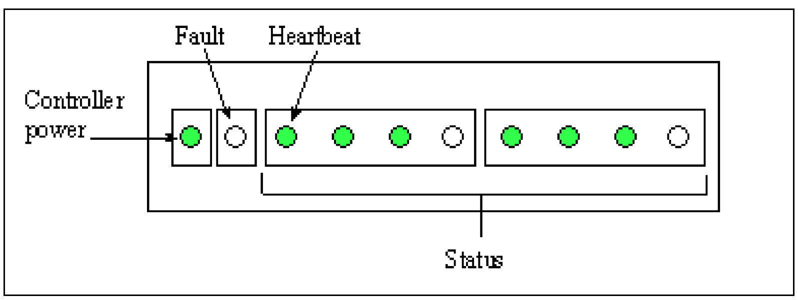

Figure 3-18 FAStT900 passive indicator lights

0xFF and Fault LED is on indicating RAID Controller is held in reset (see Figure 3-19 on page 41).

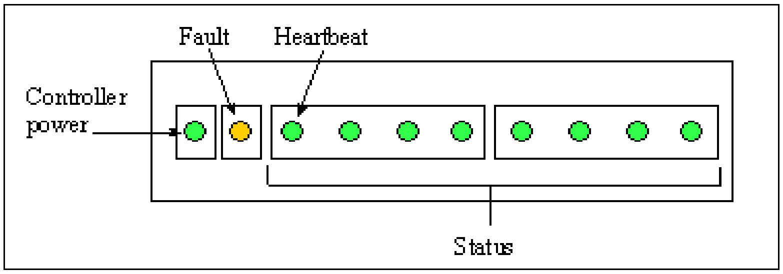

Figure 3-19 Hold in reset indicator light

0x01 and fault LED on indicate that RAID controller is currently receiving a firmware update. This may take up to 30 minutes, depending on the firmware package size. If more than 30 minutes have elapsed, there maybe a problem with the controller (see Figure 3-20).

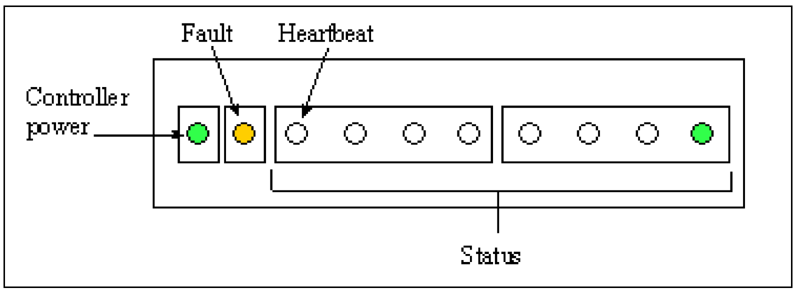

Figure 3-20 Firmware download indicator light

When all the eight status LEDs cycle (one-by-one), this indicates the RAID controller is held in boot menu.

When all the eight status LEDs cycle (add one), this indicates RAID controller firmware is being upgraded.

0x36/0x00: RAID controller processor SIMM failure.

0x37/0x00: RAID controller data cache SIMM failure.

Status LEDs in power up:

1. Cycle flashes one by one for 5–10 seconds as the controllers go through boot and load self-tests.
2. Heartbeat LED should start flashing within 15 seconds of power-on.
3. Heartbeat LED flashes one-second on and one-second off.

All other combinations of LEDs indicate a RAID Controller failure.

3.2.4 FAStT 900 indicator lights

All the indicator lights are just the same as FAStT700. The FAStT900 Storage Server indicator lights display the status of the FAStT900 Storage Server and its components. Green indicator lights mean normal operating status; amber indicator lights mean a possible failure.

It is important to check all the indicator lights on the front and back of the controller. The indicator lights might blink intermittently when power is just turned on. Wait until the FAStT900 Storage Server completes its power up before checking for faults. It can take up to 15 minutes for the battery to complete its self-test, and up to 24 hours to be fully charged, particularly after an unexpected power loss of more than a few minutes.

The indicator lights for the components of the FAStT900 Storage Server are described in the following sections.

Storage server indicator lights

The storage server has five indicator lights (as shown in Figure 3-21 with an explanation). To view the storage server indicator lights, we do not have to remove the FAStT900 Storage Server bezel.

Fan and communications module indicator lights

The fan and communications module contains five LEDs shown in Figure 3-22 on page 43.

- Ethernet link speed LEDs (green)

  There are two LEDs on each Ethernet port; they are dedicated to controller A and B. The upper LED indicates a proper link with 100Mb/sec; the lower one indicates a link with 10 Mb/sec. If the FAStT900 is attached to the Ethernet, one of the LEDs per Ethernet connection is normally on.

- Fan fault LED (amber)

  This LED is normally off. If it is on, it indicates the failure of a fan in the module.

Battery indicator lights

The battery has four indicator lights as shown in Figure 3-23 on page 43. To view the battery indicator lights, you should remove the FAStT900 Storage Server front bezel.

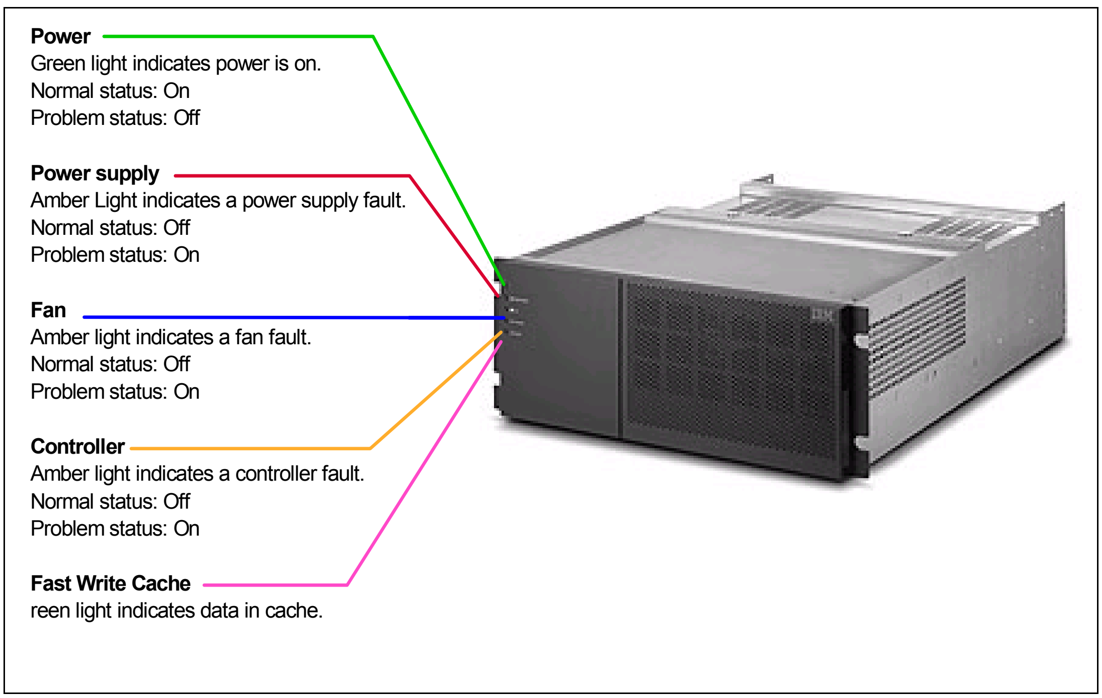

Figure 3-21 FAStT900 storage server indicator lights

Figure 3-22 on page 43 shows the FAStT900 fan and communication module indicator lights.

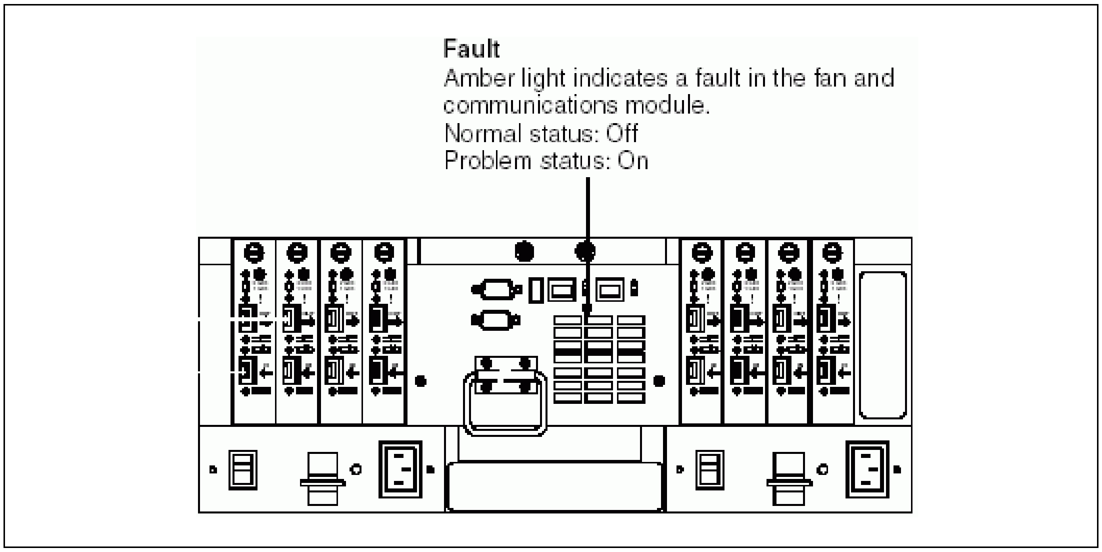

Figure 3-22 FAStT900 fan and communication module indicator lights

Figure 3-23 shows the FAStT900 battery indicator lights.

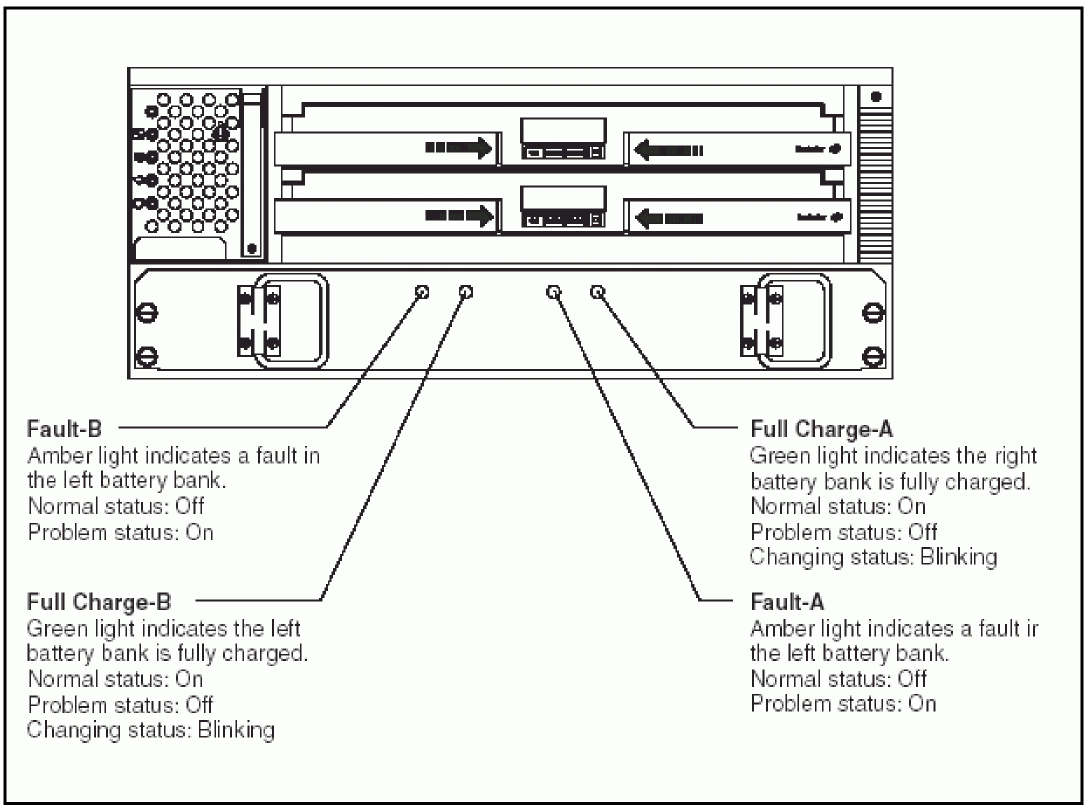

Figure 3-23 FAStT900 battery indicator light

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.

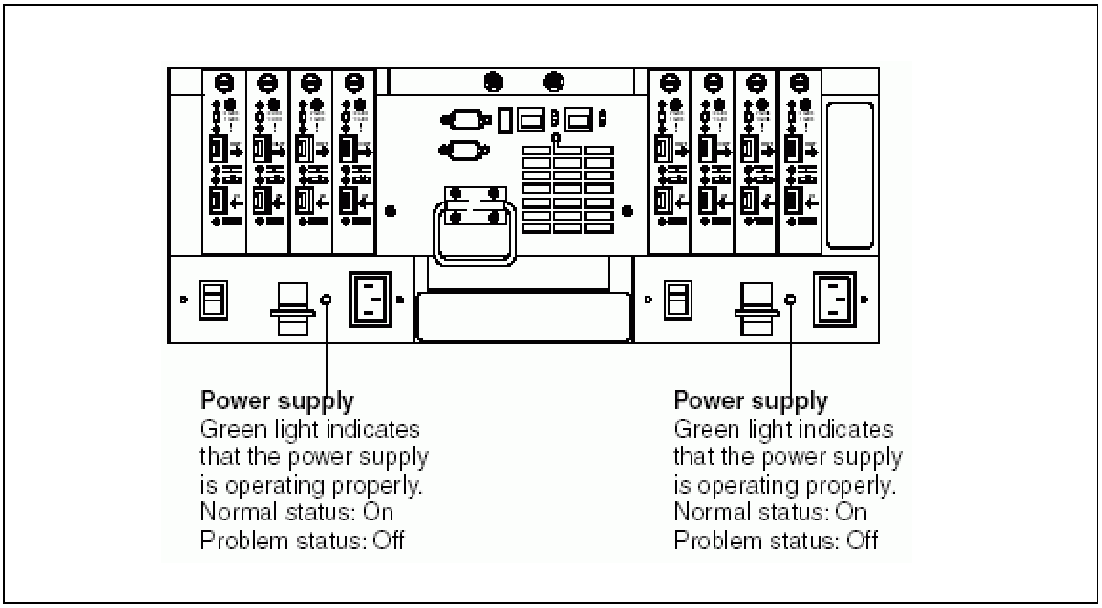

Figure 3-24 FAStT900 power indicator lights

Mini-Hub indicator lights

There are five host-side Mini-Hub indicator lights and five drive-side Mini-Hub indicator lights. Figure 3-25 shows the host-side indicator lights. The drive-side indicator lights are the same; however, the possible conditions indicated by the LEDs might be different.

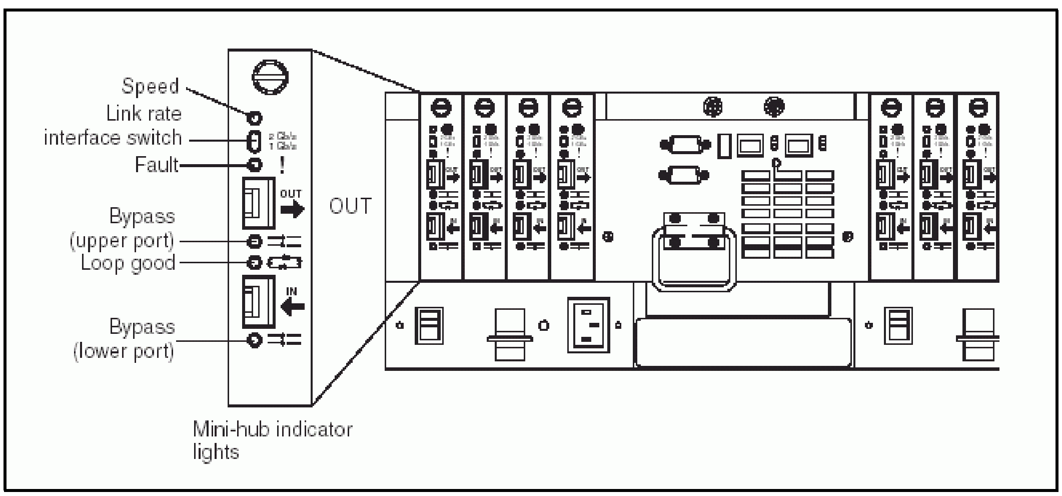

Figure 3-25 FAStT900 Mini-Hub indicator light

The LEDs present on the Mini-Hub are:

- Speed LED (green)

  This LED is on when the selected link speed is 2 Gb/s and a link is up. This LED is off when the FAStT900 RAID Controller works on 1 Gb/s.

- Fault LED (amber)

  This LED should normally be off; if on it indicates a fault of the Mini-Hub or one of the SFP modules.

- Two Bypass LEDs (amber)

  There is one bypass LED for each SFP module. 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 (for example, no cable or faulty cable, or host not powered on) it will go on.

- Loop good LED (green)

  This LED should be normally on; it may be off if we have link errors.

> **Important:** The Fibre Channel link speed must be set manually through the DIP switch; the Mini-Hub does not auto detect this setting. Figure 3-26 shows the DIP switch that allows the link speed selection for the SFP module between 2 Gb/s and 1 Gb/s.

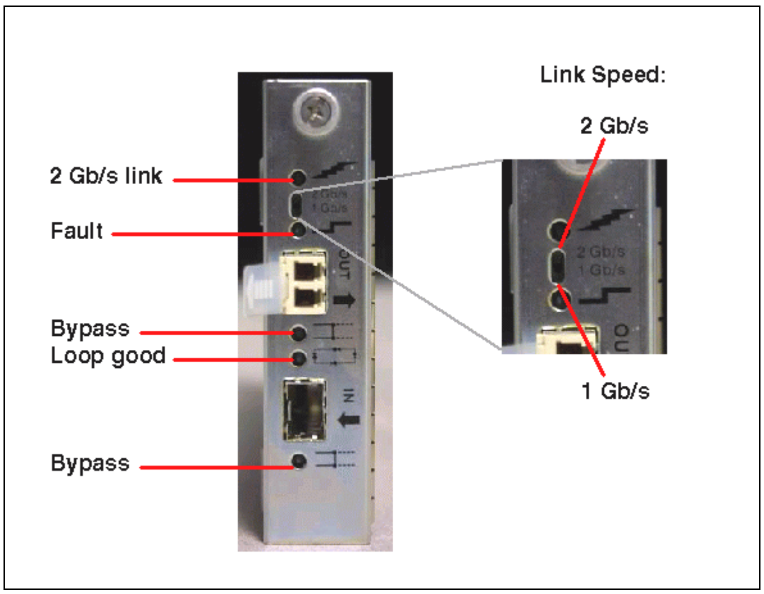

Figure 3-26 FAStT900 Mini-Hub DIP switch

3.2.5 Host side connection

There are up to four host Mini-Hubs: Two are standard. Cards 1 and 3 correspond to the top controller (controller A), and cards 2 and 4 correspond to the bottom controller (controller B). Each pair of Mini-Hubs accommodates one host channel. Each Mini-Hub provides host loop capability and self-diagnostic features. To ensure redundancy, you must connect each host to both RAID controllers, which means a host can connect to 1 and 2; 1 and 4; 2 and 3 by two HBAs or through the use of switches.

FAStT900 host-side Fibre Channel cabling is shown in Figure 3-27 on page 46:

1. Connect a host adapter to the Host 1 (top) port on host-side Mini-Hub 1. For redundancy, connect a second host adapter to the Host 1 (top) port on host-side Mini-Hub 2.
2. For a second redundant host, connect two host adapters to the Host 2 (bottom) ports on host-side Mini-Hubs 1 and 2.
3. For a third redundant host, connect two host adapters to the Host 3 (bottom) ports on host-side Mini-Hubs 3 and 4.
4. For a fourth redundant host, connect two host adapters to the Host 4 (top) ports on host-side Mini-Hubs 3 and 4.

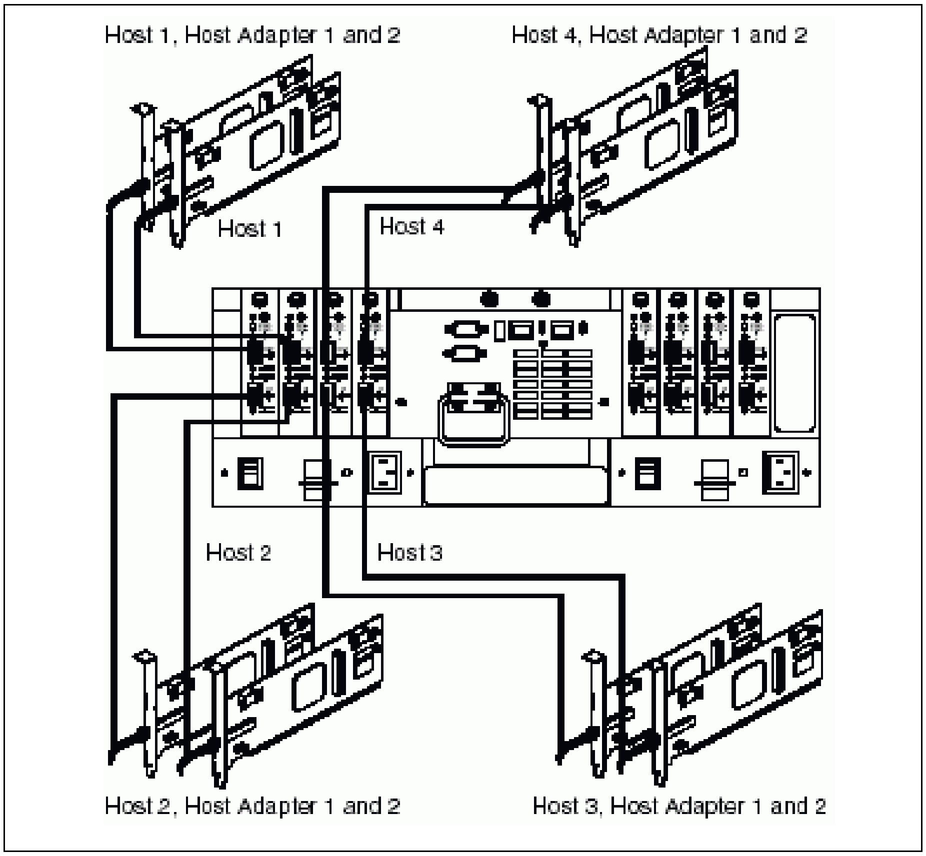

Figure 3-27 FAStT900 host-side connection

Tip: When a host side Mini-Hub on the FAStT900 is disconnected, and it has only one SFP installed, it will cause a delay on the SAN when the controller attempts to bring it up. It is recommended to either remove the unused Mini-Hub, or install two SFP modules in the unused Mini-Hub.

3.2.6 Drive-side connection

Four Mini-Hubs (interface cards) are fitted, two as standard configuration. Each Mini-Hub is attached to both controllers. Devices can be dynamically added to the Mini-Hubs. The Fibre Channel loop supports 127 addresses, so can support eight EXP700 expansion enclosures, or eleven EXP500 per drive loop, with a total of 112 or 110 drives being addressed.

Note: Addresses are assigned 1 per hard disk drive as well as one per enclosure, therefore limiting to a maximum of eleven EXP500s or eight EXP700s.

Two such fully redundant loops can be used; therefore, we can connect up to 16 EXP700 expansion enclosures or 22 EXP500 expansion enclosures. This means we can attach up to 224 disk drives by using the EXP700 or 220 disk drives using the EXP500 to the FAStT700 without a single point of failure.

Important: We recommend using EXP700 with the FAStT900: The EXP700 complements the FAStT900 for a full 2Gb Fibre Channel solution; if any EXP500 is connected into the loop, you must manually set the 1 Gb speed switch to force all the devices and hosts connected to this FAStT900 to work at 1 Gb speed.

On the drive side Mini-Hub, one SFP module port is marked as IN, the other one OUT. In fact, both ports function identically and it is thus 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 the FAStT900 to incoming ports on EXP700, you will ensure clarity and consistency of your cabling implementation, and this will allow for much easier, quicker, and more efficient troubleshooting.

FAStT900 drive side Fibre Channel cabling is illustrated below (see Figure 3-28 on page 48):

1. Start with the first expansion unit of drive enclosures group 1, and connect the In port on the left ESM board to the Out port on the left ESM board of the second (next) expansion unit.
2. Start with the first expansion unit of drive enclosures group 1 and connect the In port on the right ESM board to the Out port on the right ESM board of the second (next) expansion unit.
3. If you are cabling more expansion units to this group, repeat steps 1 and 2, starting with the second expansion unit.
4. If you are cabling a second group, repeat step 1 to step 3 and reverse the cabling order; connect from the Out ports on the ESM boards to the In ports on successive expansion units according to the illustration on the left; see Figure 3-28 on page 48.
5. Connect the Out port of the drive-side Mini-Hub 4 (leftmost drive side) to the In port on the left ESM board of the last expansion unit in the drive enclosures group 1.
6. Connect the In port of drive-side Mini-Hub 3 to the Out port on the right ESM board of the first expansion unit in the drive enclosures group 1.
7. If you are cabling a second group, connect the Out port of the drive-side Mini-Hub 2 to the In port on the left ESM board of the first expansion unit in drive enclosures group 2; then, connect the In port of the drive-side Mini-Hub 1 (rightmost drive side) to the Out port on the right ESM board of the last expansion unit in drive enclosures group 2.
8. Ensure that each expansion unit has a unique ID (switch setting) and that the left and right ESM board switch settings on each expansion unit are identical.

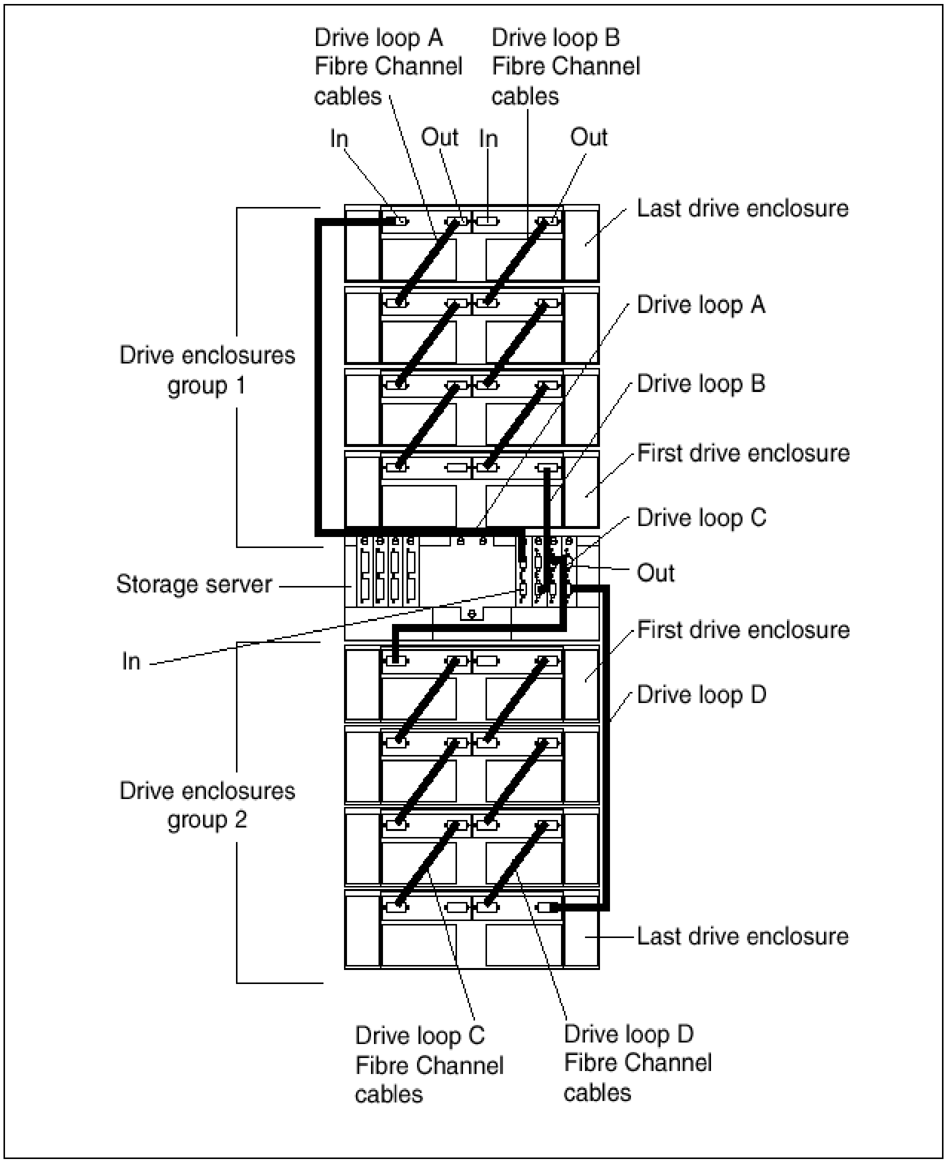

Figure 3-28 FAStT900 - drive side Fibre Channel cabling

Note: Only one port on each Mini-Hub of FAStT 900 is ever used. We recommend removing all the SFP modules on the Mini-Hub port, which does not connect to any device.

3.2.7 FAStT700 and FAStT900 comparison

FAStT700 and 900 are all 2 Gbps products; their outlook is different when you detach the front cover (see Figure 3-29 on page 49).

In comparison to the FAStT700, the performance capabilities of the FAStT900 exceed that of the FAStT700 and with a 772 MB/sec disk throughput and in conjunction using advanced caching algorithms to optimize disk performance. This makes the FAStT900 the leading Enterprise Storage Server over the FAStT700

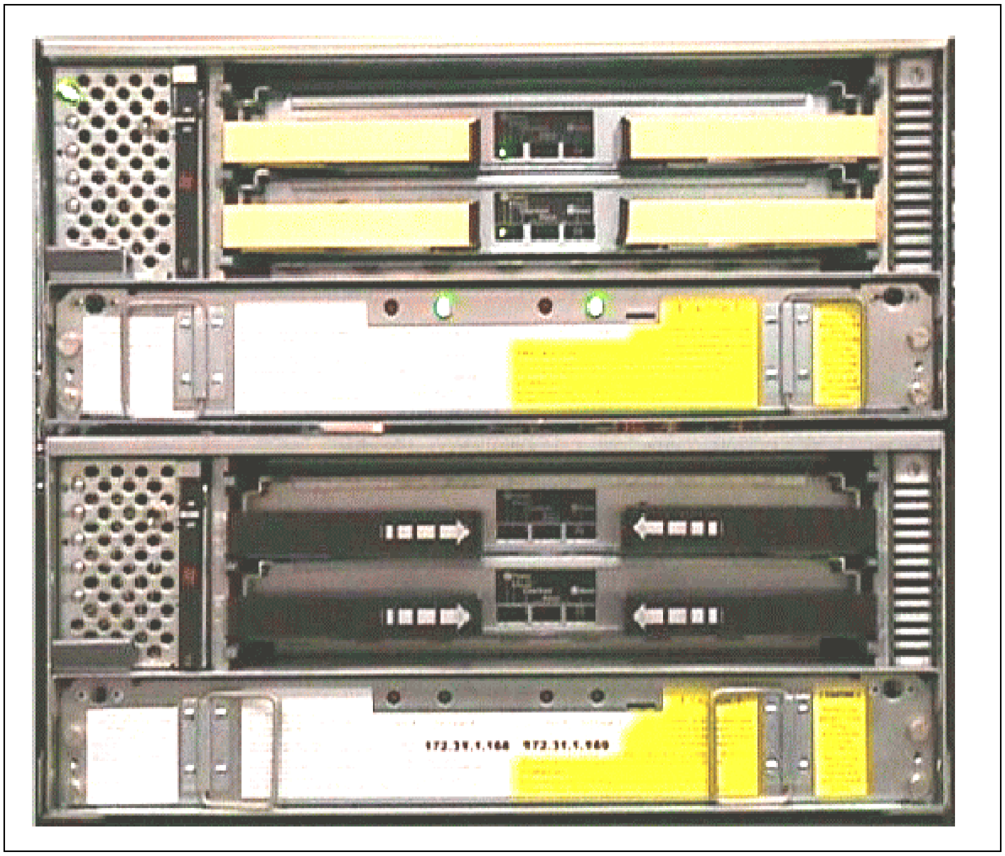

Figure 3-29 FAStT900 (upper), FAStT700 (lower)

The design of the controller has differences with the FAStT700, which gives the FAStT900 double the throughput performance of the FAStT700 (see Table 3-1).

Table 3-1 FAStT 700 and 900 comparison

FAStT700	**FAStT900**
566 MHZ Celeron processor	850 MHZ Pentium III processor
2GB Battery Backed cache	2GB Battery Backed cache
32 KB NVRAM	128 KB NVRAM
32 bit, 33 MHZ PCI bus, Figure 3-30	64 bit, 66 MHZ PCI bus, Figure 3-31
533 MB/s memory bus, Figure 3-30	1.6 GB/s memory bus, Figure 3-31

Figure 3-30 on page 50 shows the internal bus structure of the FAStT700.

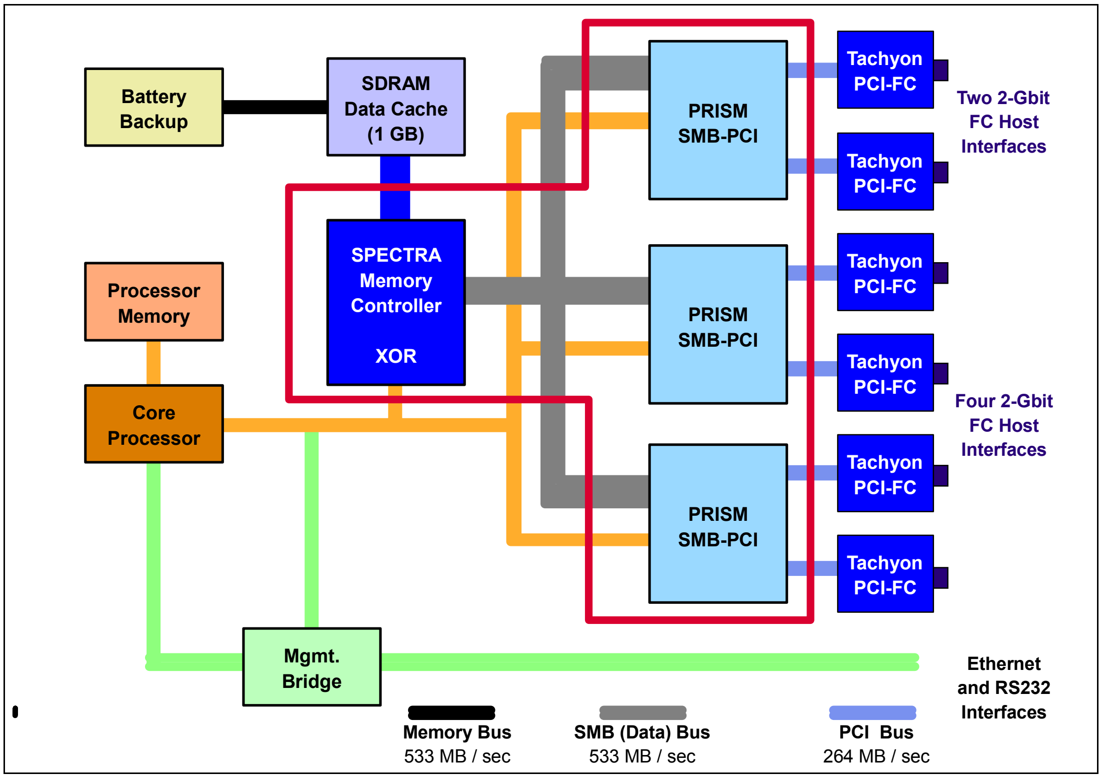

Figure 3-30 FAStT700 internal topology

Figure 3-31 shows the internal bus structure of the FAStT900.

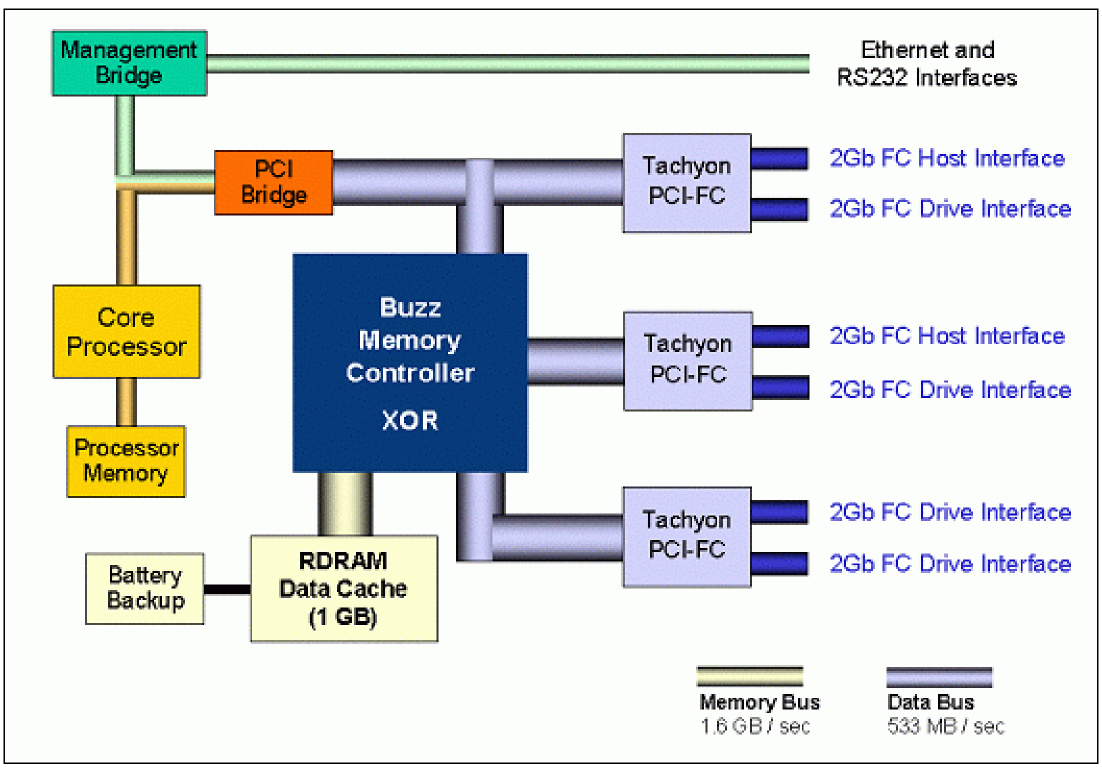

Figure 3-31 FAStT900 internal topology

3.3 IBM TotalStorage FAStT EXP700 expansion enclosure

In this section we discuss how storage can be added to the FAStT using the EXP700 expansion unit.

3.3.1 Overview of the EXP700

The EXP700 is a rack-mountable storage expansion enclosure that contains fourteen bays for slim-line hot-swappable Fibre Channel disk drives. It occupies 3U inside a rack, and features hot-pluggable and redundant power supplies and fans. Additionally, it contains two Enclosure Service Monitor (ESM) boards that can connect to the FAStT Storage Server in a fully redundant Fibre Channel loop. See Figure 3-32 for the front view of the EXP700. The machine type is 1740.

Figure 3-32 EXP700 front view

3.3.2 Front view

From the front, we can access up to14 hot-swappable Fibre Channel hard disk drives and view the two status LEDs on top (Figure 3-33 on page 52):

- Green Power LED

  This LED indicates the DC power status is OK.

- Amber General-System-Fault LED

  When a storage server component fails (such as a disk drive, fan or power supply), this LED will be on.

Each disk drive also has its own status LEDs:

- Green Activity LED

  When this LED is flashing as quickly as about three times per second, it indicates disk drive activity. When it is steadily on, it means the disk drive is OK and properly installed. When it is flashing as slowly as one time per second, it seems that the RAID controller does not recognize the HDD.

- Amber Disk-Drive-Fault LED

  If a disk drive fails, this LED will be steadily on. While the drive is being identified or rebuilt, it flashes.

- The blue LED situated right to the General System error LED, is used as a locate option from Storage Manager to locate and identify a drive enclosure in the Storage Subsystem.

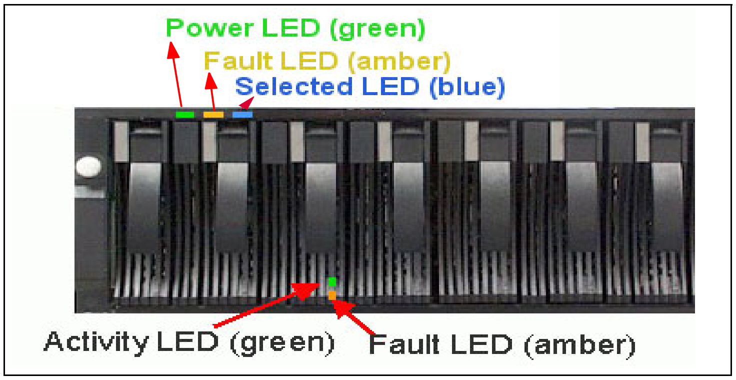

Figure 3-33 EXP700 front view

The HDDs in the enclosure can vary from 1 to 14. To ensure proper cooling and avoid EMI, all 14 bays should be occupied either by a HDD or by a filler module.

3.3.3 Rear view

You can access the ESM boards, power supplies, and fans from the rear side of the EXP700 (Figure 3-34). Two hot-swappable power supplies and fan units provide redundancy, and therefore offer a higher availability level. If a fan unit fails, do not remove it from the storage server until the replacement is available. The cooling airflow will not be 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, else, the cooling will not be efficient anymore.

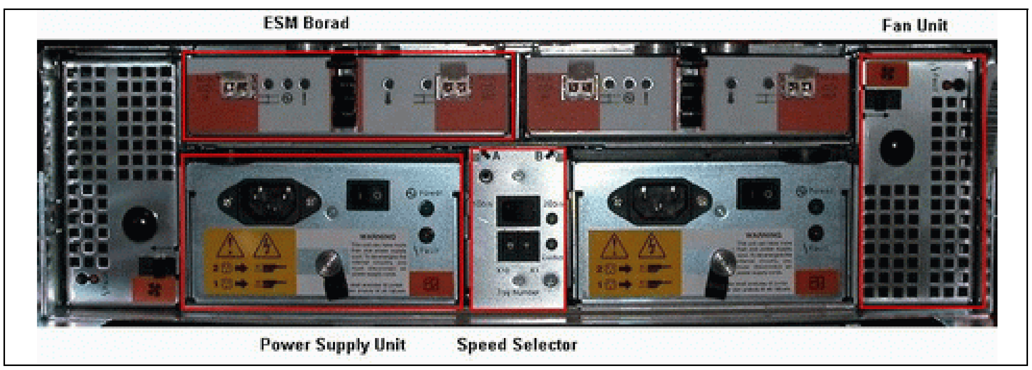

Figure 3-34 EXP700 rear view

Each ESM board contains five LEDs (see Figure 3-35).

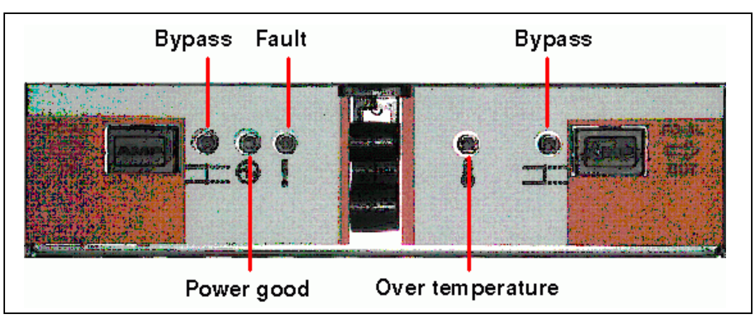

Figure 3-35 EXP700 ESM LED

- Bypass LED (green)

  This LED should normally be off when:

  – No SFP module is installed.
  – A SFP module is installed and a proper link is detected.

  The Bypass LED will turn on when a SFP module is present and a link error is detected. For example, no cable or faulty cable. There is one Bypass LED for each SFP module.

- Power Good LED (green)

  This indicates the EXP700 is powered up; it is normally on.

- Fault LED (amber)

  If a failure in the ESM board occurs, this LED will light; normally it is off.

- Over Temperature LED (amber)

  This LED indicates a over temperature condition; normally it is OFF.

3.3.4 ESM board

The ESM boards provide the connectivity of the drives inside the EXP700 to the FAStT Storage Server. As you can see in Figure 3-36, both ESM boards are connected to all fourteen disk drive bays in the EXP700. This provides redundancy in the case of one ESM board or one loop connection failed. All disk drives remain connected through the other ESM board or FC loop. Each ESM board has an incoming and outgoing SFP module port.

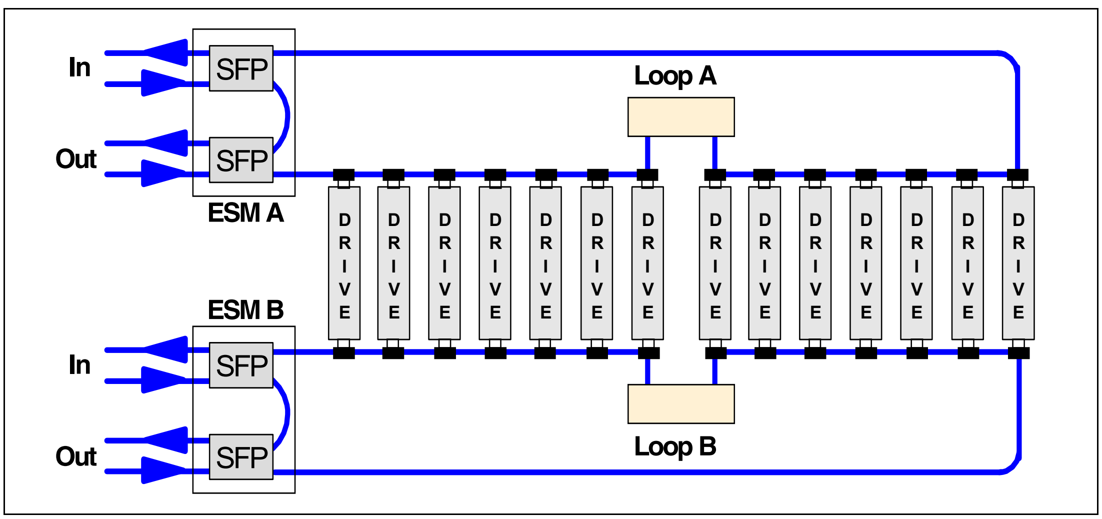

Figure 3-36 EXP700 ESM board diagram

One SFP module port is marked as IN, the other one OUT. A close look at Figure 3-36 reveals 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 our cabling implementation, and this will allow for much easier, faster, and more efficient troubleshooting.

It is very important to correctly set the tray ID switches on ESM boards. They are used to differentiate multiple EXP700 enclosures that are connected to the same FAStT Storage Server. Each EXP700 must use a unique value. The FAStT Storage Manager uses the tray

IDs to identify each EXP700 enclosure. Additionally, the Fibre Channel 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:

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

> **Important:** Every EXP attached to the same FAStT must have a different tray ID. If the FAStT Storage Server has internal HDDs, the EXP attached to it should be set to a different tray ID, otherwise, the front panel of both FAStT Storage Server and EXP will have an Amber LED.

We can therefore set any number between 0 and 99 (see Figure 3-37).

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 prohibit changing it accidentally. If you are attaching the EXP700 to a FAStT 200 and 500 Storage Server, we should change the switch to a 1 Gb setting. Figure 3-37 shows the two types of speed selection modules.

> **Note:** Newly produced models of the EXP700 changed the metal cover to a power-switch type of selection button (right side of the Figure 3-37). It is easier to operate.

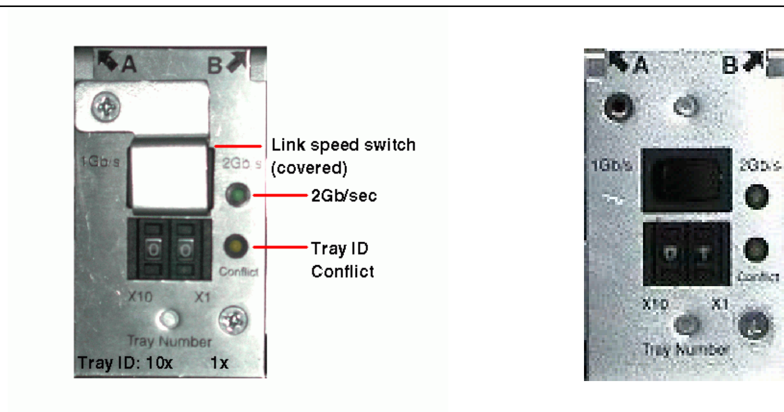

Figure 3-37 EXP700 tray ID and link speed

> **Important:** Be sure that 2 Gb/s is selected in a pure 2 Gb environment. If there are mixed connections with 1 GB products, 1 Gb/s should be selected, or else the EXP700 will not function properly.

There are two diagnostic LEDs (see Figure 3-37):

- Tray ID conflict (amber)

  The LED is normally off. This LED indicates that another EXP with the same tray ID is attached to the FAStT.

- 2 Gb/s link speed (green)

This LED indicates a link speed of 2 Gb/sec and will be lit when such a link is detected.

3.3.5 EXP700 Indicator light

Each fan module provides one LED as shown in Figure 3-38.

- Fan fault LED (amber)

  This LED is normally off. If it is on, then there is a fan failure on the module.

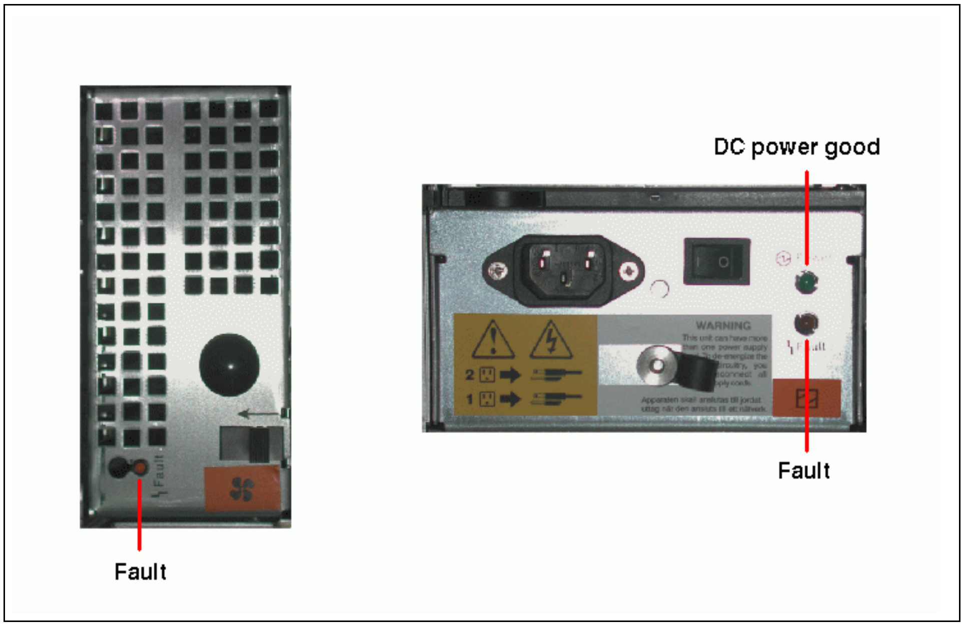

Figure 3-38 EXP700 fan LED

Each power supply has one diagnostic LED (see Figure 3-38):

- DC power good LED (green)

  This LED is normally on and indicates the DC power is OK.

- Fault LED (amber)

  This LED is normally off. If it is on, then the power supply failed.

HDD

FAStT EXP700 Supports up to 14 slim 2 Gb hot-swappable Fibre Channel drives:

- 36.4 GB (Option Number:5205, P/N:06P5761) 10 K RPM drives
  73.4 GB (Option Number:5206, P/N:06P5762) 10 K RPM drives
  146 GB (Option Number:5207, P/N:32P0765) 10 K RPM drive
- 18.2 GB (Option Number:5211, P/N:06P5771) 15 K RPM drive
  36.4 GB (Option Number:5212, P/N:06P5772) 15 K RPM drive
  73.4 GB (Option Number:5213, P/N:32P0768) 15 K RPM drive

EXP500

The IBM FAStT EXP500 enclosure is compatible with the following IBM products:

- Type 3552 FAStT500 RAID controller
- Type 3542 FAStT200 controller
- Type 1742 FAStT700 controller
- Type 1742 FAStT900 controller

Important: It is strongly recommended that you connect *only* EXP700 to a 2 Gb FAStT Storage Server. Connecting the EXP500 to the 2 Gb environment will reduce the speed of the whole system to 1 Gb.

HDD

FAStT EXP500 supports up to 10 half high 1 Gb hot-swap Fibre Channel drives:

- 36.4 GB (Option Number:5005, P/N:19K0653)10 K RPM 1 Gb drive
  73.4 GB (Option Number:5006, P/N:19K0654) 10 K RPM 1 Gb drive
- 18.2 GB 1(Option Number:5011, P/N:06P5707) 15 K RPM 1 Gb drive
- 36.4 GB (Option Number:5205)10 K RPM 2 Gb drive
  73.4 GB (Option Number:5206) 10 K RPM 2 Gb drive
  146.8 GB (Option Number:5207) 10 K RPM 2 Gb drive
- 18.2 GB (Option Number:5211) 15 K RPM 2 Gb drive
  36.4 GB (Option Number:5212)15 K RPM 2 Gb drive
  73.4 GB (Option Number:5213) 15 K RPM 2 Gb drive

3.4 FAStT physical installation and cabling considerations

In this section we review some topics of importance when planning for the physical installation of a FAStT system.

3.4.1 Rack considerations

The FAStT and possible expansions are mounted in rack enclosures.

Preparing the physical site

Before you install a rack enclosure, be sure you:

- Understand the rack specifications and requirements.
- Prepare a layout for the racks.
- Prepare the physical site.

General planning

For general planning:

1. Determine:
   - The size of the floor area required by the equipment (see “Product Specifications”):
     - Floor-load capacity
     - Space needed for expansion
     - Location of columns
   - The power and environmental requirements.
2. Create a floor plan to check for clearance problems.
3. Make a full-scale template (if necessary) of the rack and carry it along the access route to check for potential clearance problems through doorways and passage ways, around corners, and in elevators.
4. Provide space for storage cabinets, card files, desks, communication facilities, daily storage of tapes, and other supplies.
5. Store all spare materials that can burn in properly designed and protected areas.

Rack layout

To be sure you have enough space for the racks, create a floor plan before installing the racks. You might need to prepare and analyze several layouts before choosing the final plan.

If you are installing the racks in two or more stages, prepare a separate layout for each stage.

Consider the following when you make a layout:

- The flow of work and personnel within the area.
- Operator access to units, as required.
- If the rack is on a raised floor, position it over a cooling register. The bottom of the rack is open to facilitate cooling.
- If the rack is not on a raised floor, determine:
  - The maximum cable lengths
  - The need for such things as cable guards and ramps to protect equipment and personnel
- Location of any planned safety equipment.
- Expansion.

Begin with an accurate drawing of the installation area (blueprints and floor plans are appropriate).

Be sure to include the following on the layout:

- Service clearances required for each rack or suite of racks.
- If the equipment is on a raised floor:
  - Things that might obstruct cable routing
  - The height of the raised floor
- If the equipment is not on a raised floor:
  - The placement of cables to minimize obstruction
  - If the cable routing is indirectly between racks (such as along walls or suspended), the amount of additional cable
- Location of:
  - Power receptacles
  - Air conditioning equipment and controls
  - File cabinets, desks, and other office equipment
  - Room emergency power-off controls
  - All entrances, exits, windows, columns, and pillars

Review the final layout to ensure that cable lengths are not too long and that the racks have enough clearance.

You need at least 152 cm (60 in) of space between 42-U rack suites. This space is necessary for opening the front and rear doors, and for installing and servicing the rack. It also allows air circulation for cooling the equipment in the rack. All vertical rack measurements are given in rack units (U). One U is equal to 4.45 cm (1.75 in). The U levels are marked on labels on one front mounting rail and one rear mounting rail. Figure 2-4 shows an example of the required service clearances for a 9306-900 42U rack.

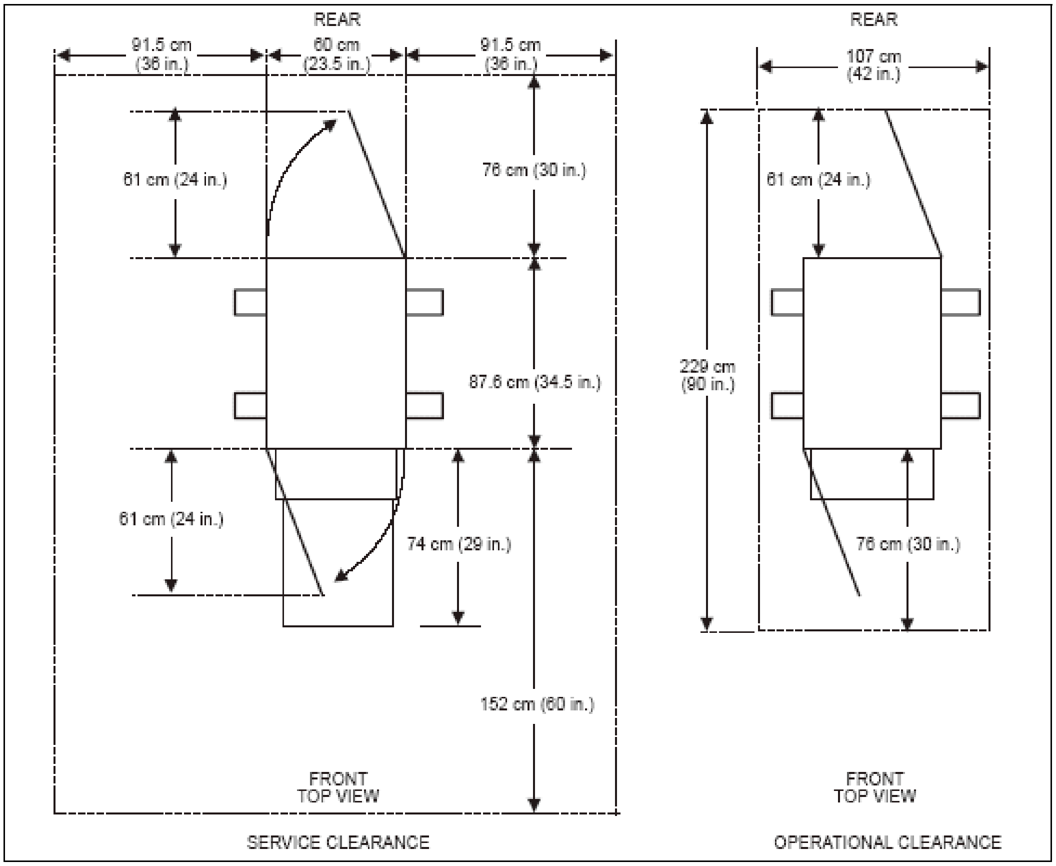

Figure 3-39 Example of service clearances for 9306-900 Rack

3.4.2 Cable management and labeling

Cable management and labeling for solutions utilizing racks, n-node clustering and Fibre Channel is increasingly important in Intel-processor solutions. Cable management and labeling needs have expanded from the traditional labeling of network connections to management and labeling of most cable connections between your servers, disk subsystems, multiple network connections, power and video subsystems. Examples of solutions include Fibre Channel configurations, n-node cluster solutions, multiple unique solutions located in the same rack or across multiple racks, and solutions where components may not be physically located in the same room, building or site.

Why more detailed cable management is required

The necessity for detailed cable management and labeling is due to the complexity of today's configurations, potential distances between solution components, and the increased number of cable connections required to attach additional value-add computer components. Benefits from more detailed cable management and labeling include: Ease of installation, ongoing solutions/systems management, and increased serviceability.

Solutions installation and ongoing management is easier to achieve when your solution is correctly and consistently labeled. Labeling helps make it possible to know what system you are installing or managing (for example, when it is necessary to access the CD-ROM of a particular system and you are working from a centralized management console). It is also helpful to be able to visualize where each server is when completing custom configuration tasks such as node naming and assigning IP addresses.

Cable management and labeling improve service and support by reducing problem determination time, ensuring the correct cable is disconnected when necessary. Labels will assist in quickly identifying which cable needs to be removed when connected to a device such as a hub that may have multiple connections of the same cable type. Labels also help identify which cable to remove from a component. This is especially important when a cable connects two components that are not in the same rack, room, or even the same site.

Cable planning

Successful cable management planning includes three basic activities: Site planning (before your solution is installed), cable routing, and cable labeling.

Site planning

Adequate site planning completed before your solution is installed will result in a reduced chance of installation problems. Significant attributes covered by site planning are location specifications, electrical considerations, raised/non-raised floor determinations, and determination of cable lengths. Consult the documentation of your solution for special site planning considerations. IBM Netfinity Racks document site planning information in the *IBM Netfinity Rack Planning & Installation Guide* (part number 24L8055).

Cable routing

Effective cable routing will allow you to keep your solution's cables organized, reduce the risk of damaging cables, and allow for effective service and support. To assist with cable routing, IBM recommends the following:

- When installing cables to devices mounted on sliding rails:
  - Run the cables neatly along equipment cable-management arms and tie the cables to the arms. (Obtain the cable ties locally.)

    Note: Do not use cable-management arms for Fibre cables.

  - Take particular care when attaching fiber optic cables to the rack. Refer to the instructions included with your fiber optic cables for guidance on minimum radius, handling, and care of fiber optic cables.
  - Run the cables neatly along the rack rear corner posts.
  - Use cable ties to secure the cables to the corner posts.
  - Make sure the cables cannot be pinched or cut by the rack rear door.
  - Run internal cables that connect devices in adjoining racks through the open rack sides.
  - Run external cables through the open rack bottom.
  - Leave enough slack so the device can be fully extended without putting a strain on the cables.
  - Tie the cables so the device can be retracted without pinching or cutting the cables.
- To avoid damage to your fiber-optic cables, follow these guidelines:
  - Do not route the cable along a folding cable-management arm.
  - When attaching to a device on slides, leave enough slack in the cable so that it does not bend to a radius smaller than 76 mm (3 in) when extended or become pinched when retracted.
  - Route the cable away from places where it can be snagged by other devices in the rack.

– Do not overtighten the cable straps or bend the cables to a radius smaller than 76 mm (3 in.).

– Do not put excess weight on the cable at the connection point and be sure that it is well supported.

Additional information for routing cables with IBM Netfinity Rack products can be found in the *IBM Netfinity Rack Planning & Installation Guide* (part number 24L8055). This publication includes pictures providing more details on recommended cable routing.

Cable labeling

When labeling your solution, follow the below tips:

- As you install cables in the rack, label each cable with appropriate identification.
- Remember to attach labels to any cables you replace.
- Document deviations from the label scheme you use. Keep a copy with your Change Control Log book.

Whether using a simple or complex scheme, the label should always implement a format including these attributes:

- Function (optional).
- Location information should be broad to specific (for example, site/building to specific port on a server or hub).
- The optional Function row helps identify the purpose of the cable (that is, Ethernet vs. Token Ring or between multiple networks).

Common cabling mistakes

Some of the most common mistakes include:

- Leaving cables hanging from connections with no support
- Not using dust caps
- Not keeping connectors clean
- Leaving them dragging on the floor for people to kick and trip over

3.4.3 Cabling and connectivity

In this section, we review one essential physical characteristic of fibre cables and connectors.

Cable types

Fibre cables are basically available in multi-mode fiber (MMF) or single mode fiber (SMF). Both types can support shortwave as well as longwave.

Multi-mode fiber (MMF) allows light to disperse in the fiber so that it takes many different paths, bouncing of the edge of the fiber repeatedly to finally get to the other end (multi-mode means multiple paths for the light). The light taking these different paths gets to the other end of the cable at slightly different times (different path - different distance - different time). The receiver has to figure which signals go together as they all come flowing in: The maximum distance is limited by how “blurry” the original signal has become. The thinner the glass the less the signals “spread out” and the further you can go and still figure out what is what on the receiving end. This dispersion (called modal dispersion) is the critical factor in determining the maximum distance a high speed signal can go; it is more relevant than the attenuation of the signal (from an engineering standpoint it is easy enough to increase the power level of the transmitter and/or the sensitivity of your receiver, but too much dispersion cannot be decoded no matter how strong the incoming signals are).

Single mode fiber (SMF) is so thin (9 microns) that the light "can barely squeeze through" and basically tunnels through the center of the fiber using only one path (or mode). This behavior can be explained (although not simply) through the laws of optics and physics. The result is that since there is only one path that the light takes to the receiver, there is no "dispersion confusion" at the receiver. However, the concern with single mode fiber is attenuation of the signal

Table 3-2 .Cable type overview

Fiber type	Speed	Maximum distance
9 micron SMF (longwave)	1Gb	10km
9 micron SMF (longwave)	2Gb	2km
50 micron MMF (shortwave)	1Gb	500m
50 micron MMF (shortwave)	2Gb	300m
62.5 micron MMF (shortwave)	1Gb	175m/300m
62.5 micron MMF (shortwave)	2Gb	90m/150m

Interoperability of 1Gb and 2Gb devices

The Fibre Channel standard specifies a procedure for speed auto-detection. Thus, if a 2Gb port on a switch or device is connected to a 1Gb port it will negotiate down and run the link at 1Gb. If there are two 2Gb ports on either end of a link then the negotiation will end up running the link at 2Gb if the link is up to spec. A link that is too long, or "dirty" could end up running at 1Gb even with 2Gb ports at either end so watch your distances and make sure your fiber is good.

Connectors

The FAStT600, FAStT700, FAStT900 and EXP700 use LC connectors. The Mini-Hubs in FAStT500 and EXP500 use SC connectors. The differences in the two types of connectors are shown in Figure 3-40 on page 62.

- To Install an SFP module and Fibre Channel cable: Remove the protective caps from the SFP module and from the Fibre channel cable. *Do not touch* the exposed fibre-optic cable (see Figure 3-41 on page 62). Insert the SFP module into an SFP module port on a FAStT600 FC host or expansion port and then connect the fibre cable to SFP module.

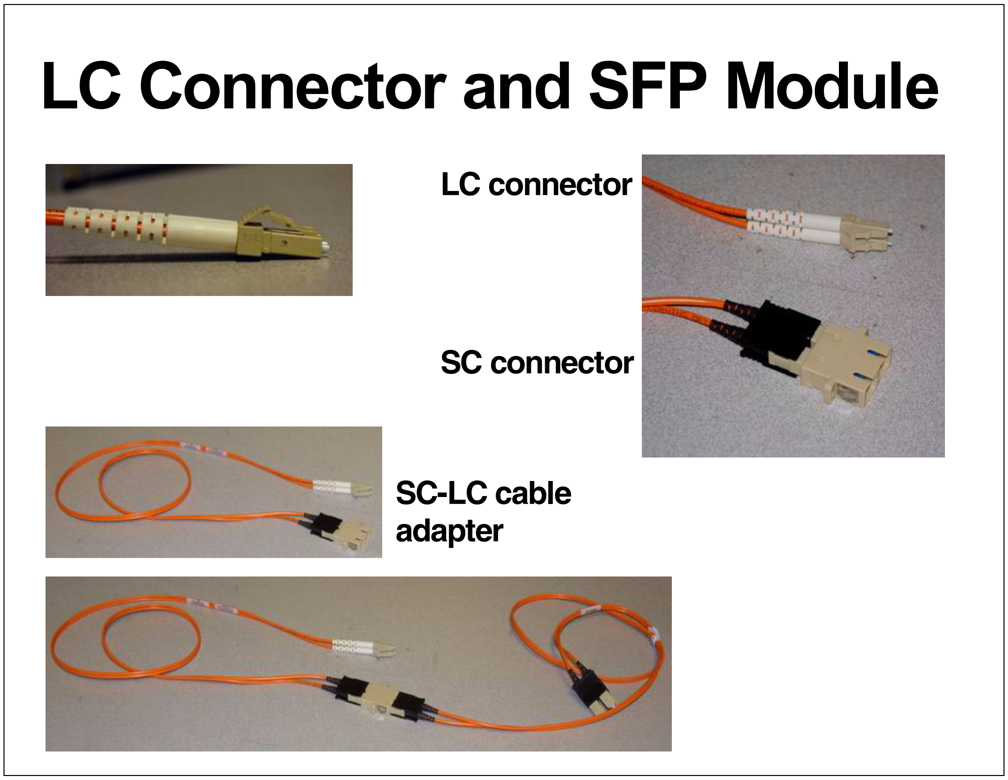

Figure 3-40 LC and SC connectors

- To Remove a Fibre Channel cable and SFP module: For SFPs with plastic tabs, pull the SFP module plastic tab outward at a 10 degree angle; for the wire tab pull the wire tab outward at 90 degrees angle.

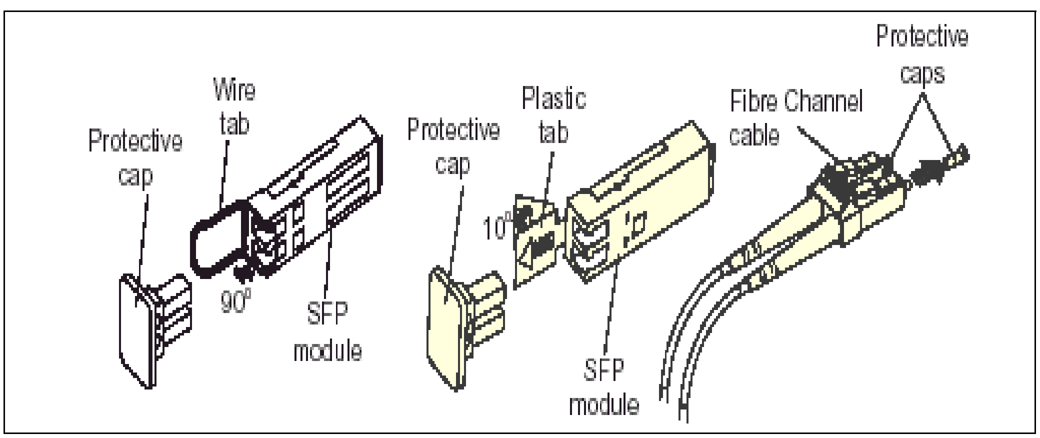

Figure 3-41 SFP Module and LC Fibre cable connector

- When planning the cabling for a FAStT700, use care to ensure that the correct cables are ordered with the correct connectors, SC or LC. Also, you need converters to connect FAStT700 and EXP500, as well as FAStT700 and older HBAs and switches.

3.4.4 Host Bus Adapters

We provide hereafter a list per hardware platform, of the most popular Host Bus Adapters (HBA) used to connect to the FAStT 600/900.

- For Intel based systems running Windows 2000 or Windows 2003, Linux, or Netware 5.1 and 6.0, use either:
  - IBM FAStT FC-2 133 Host BUS adapter (FC 2104) BIOS 1.34
  - IBM FAStT FC-2 133 Host BUS dual port adapter BIOS 1.34
- For pSeries systems running AIX 5.1 or 5.2, use either a:
  - IBM pSeries FC 6227 Fibre Channel Adapter for PCI Bus
  - IBM FC 6239 2Gb PCI-X FCP Adapter
- For Sun Solaris 2.6, 7,8:
  - JN I FCE-6460 PCI 64 bit
  - JNI FCE -1473 SBUS 64 bit

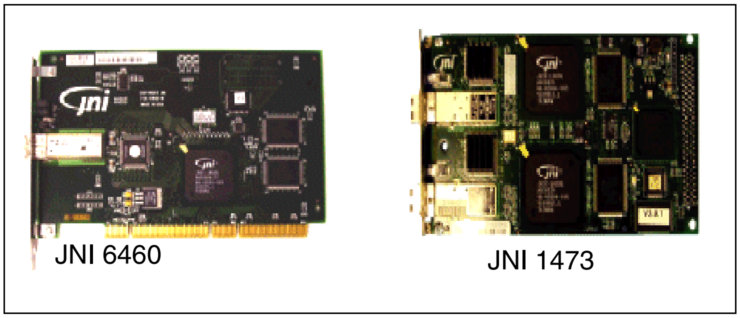

Figure 3-42 JNI adapters

- For HP-UX:

  HP model A 5158A (1GB) and A6795A (1 0r 2 GB)

For the latest list of supported adapter models and their driver please check:

`http://www.storage.ibm.com/support`

There is a search tool at:

`http://knowledge.storage.ibm.com/HBA/HBASearchTool`

Host system bus

As you can see, there is a choice of high-speed adapters for connecting disk drives; obviously, fast adapters can provide better performance, but you must be careful not to put all the high-speed adapters on a single system bus; otherwise the computer bus becomes the performance bottleneck.

We recommend distributing high-speed adapters across several busses. When you use PCI adapters, make sure you first review your system specifications: Some systems include a PCI adapter placement guide.

The number of adapters you can install obviously depends on the number of PCI slots available on your server, but also depends on what traffic volume to expect on your SAN.

Do you want only failover capabilities on the storage side (one HBA, two paths) or do you want to share the workload and have fully redundant path failover with multiple adapters and over multiple paths? In general all operating systems support two paths to the FAStT Storage Server.

Windows 2000 and Windows 2003 support up to four paths to the storage controller. As illustrated in Figure 3-2, AIX can also support four paths to the controller, provided there are two partitions accessed within the FAStT subsystem.

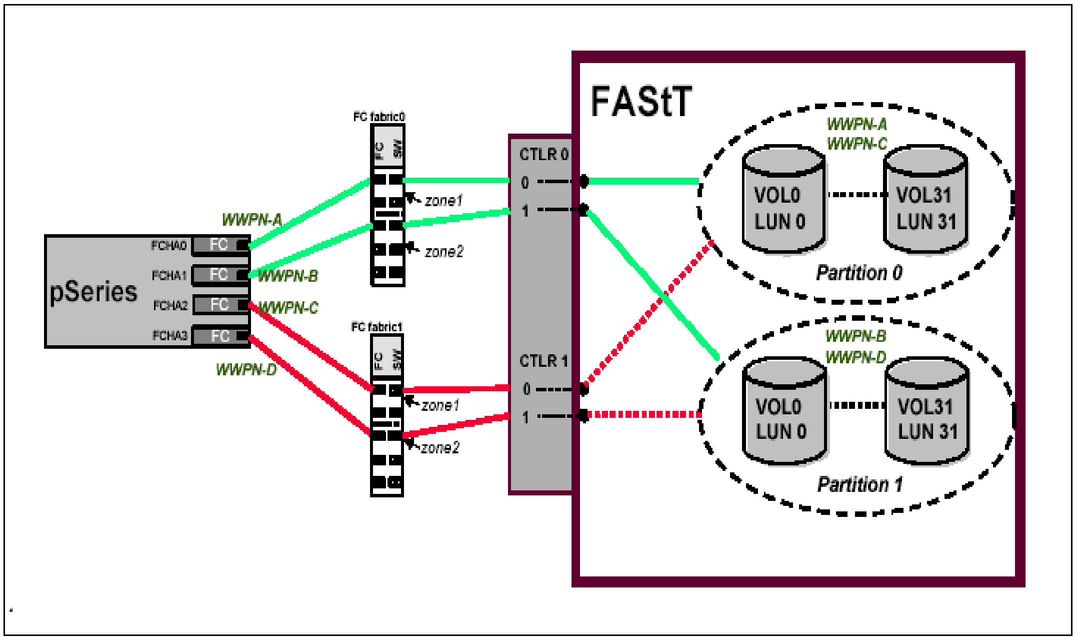

Figure 3-43 AIX - four paths to FAStT

3.4.5 SAN switches

IBM offers a wide variety of SAN switches that can be used when connecting FAStT Storage servers in a SAN infrastructure. All switches offer 1 or 2 GB connectivity.

We list below the most popular SAN switches.

- IBM TotalStorage SAN switches 2109-F08, 2109-F16,2109 -F32,3534-F08, 2109-M12
- Inrange FC 9000 Director
- McData ED 5000

For a full list of products, please consult:

http://www.storage.ibm.com/ibmsan

For the latest Firmware level see:

http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/san

Figure 3-44 on page 65 illustrates connectivity options for the FAStT600 and FAStT900.

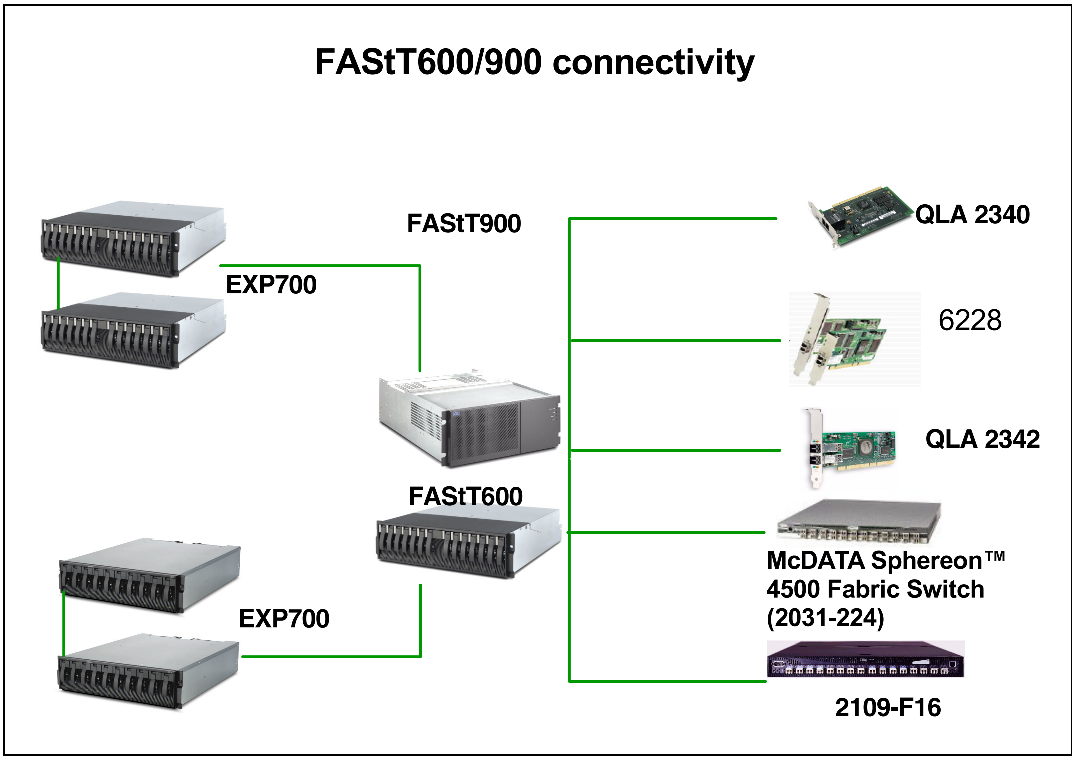

Figure 3-44   Connectivity overview

4

Managing the FAStT Storage Server

This chapter presents the main features and components of the FAStT Storage Manager software, which is used to configure, manage, and troubleshoot the FAStT Storage Servers.

The chapter starts by introducing the components of Storage Manager and the management methods. It continues with an overview of the concepts and tasks used for basic management of the storage server; finally we review premium features and advanced management functions offered by the Storage Manager software.

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
- Tune the FAStT Storage Server
- Perform diagnostic and troubleshooting tasks
- Execute advanced functions such as FlashCopy, VolumeCopy and Remote Mirroring

Attention: We used Storage Manager Version 8.4 in preparation of this publication. SM 8.4 and some of the new features it introduces are only supported on FASt600 Turbo, FAStT700 and FastT900 models.

In addition, some of the functions discussed in this section were not formally supported on all operating systems at the time this book was written. Please 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 FAStT supported servers Web site at:

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

4.1 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. Storage Manager 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. Depending on the version of Storage Manager and availability (at your installation) of optional premium features, you can configure and manage FlashCopy volumes, and perform VolumeCopy and Remote Volume Mirroring (RVM).

The FAStT Storage Manager software consists of the following:

- Microsoft Virtual Machine (MSVM)

  Microsoft Virtual Machine is required to support the Event Monitor option of the SMclient package on a host computer running Windows NT4.0 or Windows Server 2003 (32-bit editions). You must manually install the Microsoft Virtual Machine on your management station, host computer, or server as applicable to your configuration.

- Storage Manager client and Event Monitor

  The Storage Manager client (SMclient) component provides the graphical user interface (GUI) for managing storage subsystems through the Ethernet network or from the host computer.

  The Storage Manager client is called thin because it only provides an interface for storage management, based on information supplied by the storage subsystem controllers. When you install the SMclient on a management station, you send commands to the storage-subsystem controllers. The controller firmware contains the necessary logic to carry out the storage management commands. The controller is responsible for validating and executing the commands and providing the status and configuration information that is sent back to the SMclient.

  The Event Monitor is a separate program that is bundled with the SMclient. It runs in the background and can send alert notifications in the event of a critical problem

  Please refer to “The FAStT Storage Manager client” on page 72 for additional details; the functions of the SMclient are reviewed throughout this book

- Redundant Dual Active Controller (RDAC) multi-path driver software

  RDAC is a Fibre Channel I/O path failover driver that is installed on host computers. Usually, a pair of active controllers is located in a storage subsystem enclosure. Each logical drive in the storage subsystem is assigned to one of the controllers, which control the Fibre Channel I/O path between the logical drive and the host computer through the Fibre Channel network. When a component in the Fibre Channel I/O path fails, such as a cable or the controller itself, the RDAC multi-path driver transfers ownership of the logical drives assigned to that controller to the other controller in the pair. The hot-add part of RDAC allows you to register new logical drives to the operating system dynamically.

  Note: Some operating systems have built-in Fibre Channel I/O path failover drivers and do not require RDAC multi-path driver.

- Storage Manager agent

  The Storage Manager agent (SMagent) package contains the host agent software. You can use the host agent software to manage storage subsystems through the host computer Fibre Channel I/O path. The host agent software takes requests from a management station that is connected to the host computer through a network connection and passes the requests to the storage subsystem controllers through the Fibre Channel I/O path. The host agent, along with the network connection on the host computer, provides an in-band host agent type network management connection to the storage subsystem instead of the out-of-band direct network management connection through the individual Ethernet connections on each controller. The management station can communicate with a storage subsystem through the host computer that has host agent management software installed. The host agent receives requests from the management station through the network connection to the host computer and sends them to the controllers in the storage subsystem through the Fibre Channel I/O path. The SMagent package contains the SMdevices utility that associates logical drives you create using the storage-management software with their operating-system device names.

Host computers with the host-agent software installed are automatically discovered by the storage-management software and appear in the device tree in the Enterprise Management window along with their attached storage subsystems. A storage subsystem might be duplicated in the device tree if you are managing it through its Ethernet connections and it is attached to a host computer with the host-agent software installed. In this case, the duplicate storage subsystem icon can be removed from the device tree using the Remove Device option in the Enterprise Management window.

- Storage Manager Utilities (SMutil)

  The Storage Manager Utilities package contains command line tools hot_add, SMdevices, SMflashcopyassist and SMrepassist.

  Please refer to “Storage Manager utilities” on page 74 for details.

4.1.1 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. This requires installation of the Storage Manager agent package. The management information can either be processed in the host or passed to the management station through the network connection, as shown in Figure 4-1 on page 70. For more information how to install it refer to “Installing the HBA device drivers and host software” on page 110.

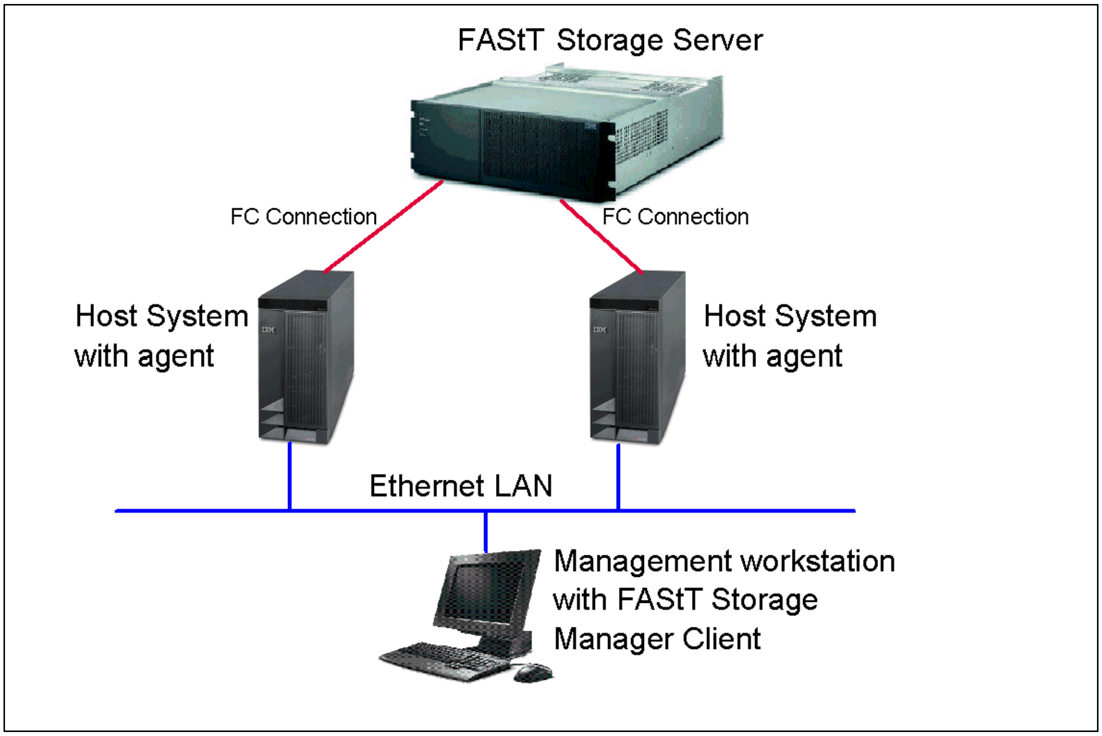

Figure 4-1 In-band management

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 can be 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.

Notes: Take note of the following:

- 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 or you must set up the FAStT Storage Server to use DHCP/BOOTP settings. The controllers must be attached, and a cable must be attached to the Ethernet connectors on each of the storage subsystem controllers, as shown in Figure 4-2.

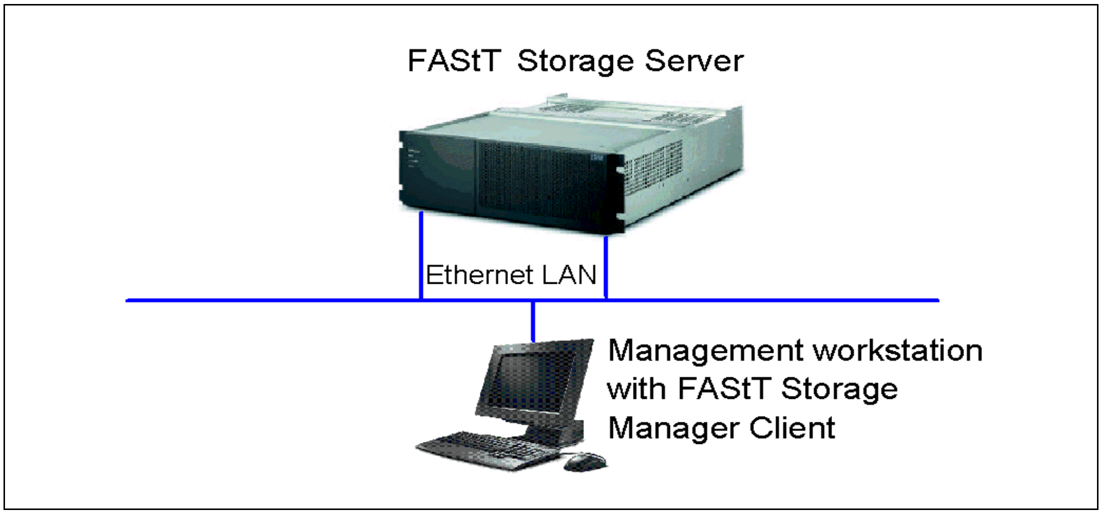

Figure 4-2 Out-of-band management

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 the current level of Storage Manager.
- 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. The user can manage and troubleshoot the storage subsystem even when there are problems with the Fibre Channel links.
- Security is enhanced because you can create management LANs/VLANs and use more advanced solutions, such as VPN, to manage the system remotely.
- More FAStT systems in the network can be managed through one Storage Manager interface.

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, refer to "Network setup of the controllers" on page 135. With Storage Manager 8.4 it is also possible to set the network parameters using the client GUI (see "User control of network parameters" on page 15).

If the storage subsystem controllers have firmware version 05.00.xx or later, then FAStT will have the default settings shown in Table 4-1, only if no DHCP/BOOTP server is found.

Table 4-1 Fixed IP addresses

Controller	IP address	Subnet mask
A	192.168.128.101	255.255.255.0
B	192.168.128.102	255.255.255.0

4.1.2 The FAStT Storage Manager client

We continue in this section our overview of the Storage Manager client.

Supported host systems

The FAStT Storage Manager client is a Java-based GUI utility that is available for various operating systems. The list of the operating systems supported by Storage Manager 8.4 was not finalized at the time this redbook was written. For information about support for specific FAStT models you must check the FAStT supported server Web site at:

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

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

The Enterprise Management window (Figure 4-3 on page 73) is the first window that opens when you start the storage-management software. 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 on the same subnet 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.

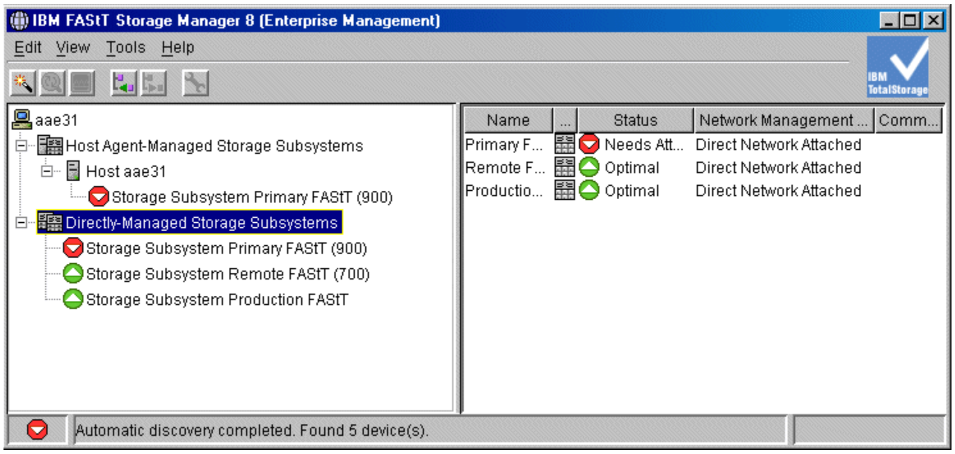

Figure 4-3 Enterprise Management window

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 4-4, 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. The storage-subsystem capacity is organized into 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. The Physical view displays the physical devices in the storage subsystem, such as controllers and drives. When you click a logical drive or other item in the Logical view, the associated physical components are displayed in the Physical view, as shown on Figure 4-4. The selection on the controllers and disk drives enables you to perform the tasks on them.

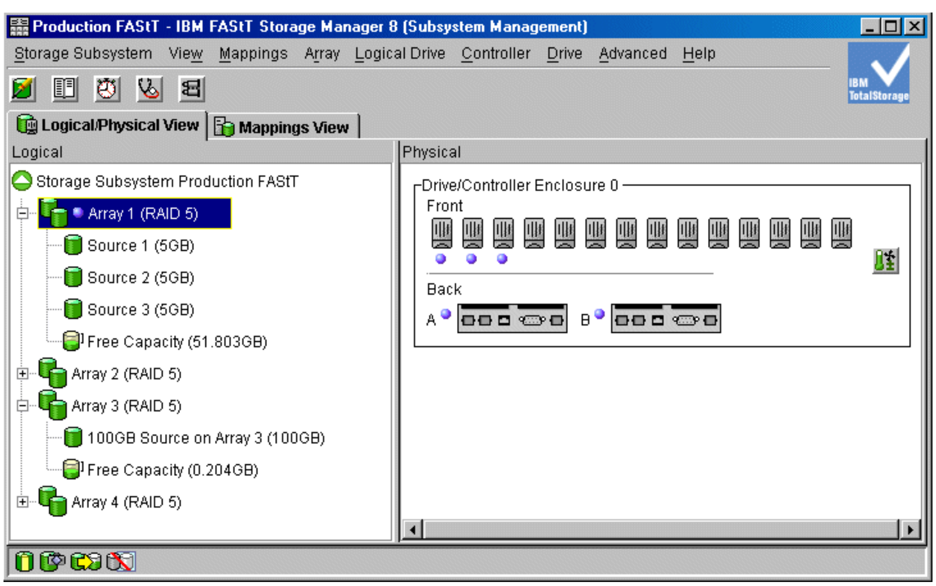

Figure 4-4 The Logical/Physical view

Mappings view

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, configure, or reconfigure which server is allowed to see which logical drive. Figure 4-5 shows the Mappings view of the logical drives that have been created in the FAStT Storage Server.

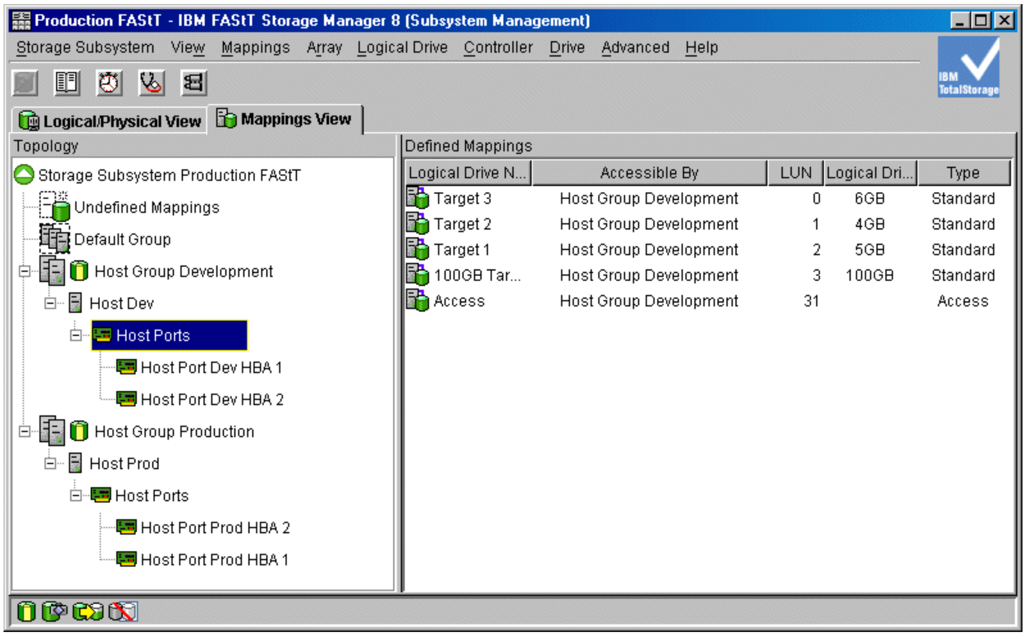

Figure 4-5 The Mappings view

Event Monitor

The Event Monitor handles notification functions (e-mail and SNMP traps) and monitors storage subsystems whenever the Enterprise Management window is not open. Monitoring is performed either in the Enterprise Management window or, when the Enterprise Management window does not run, by the Event Monitor.

The Event Monitor is a separate program bundled with the Storage Manager client software. (The Event Monitor cannot be installed without the client.) The Event Monitor can only be installed on a management station or host computer connected to the storage subsystems. For continuous monitoring, you are required to install the Event Monitor on a host computer that runs 24 hours a day. Once installed, the Event Monitor 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.

For additional information about how to set up the Event Monitor refer to “Monitoring and alerting” on page 157.

4.1.3 Storage Manager utilities

FAStT Storage Manager V8.4 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 (SMrepassist): This utility is a host-based utility for Windows platforms that performs certain functions needed to make the subsystem hardware FlashCopy work smoothly. It 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.

    This utility is extended in Storage Manager 8.4 to provide similar support for the VolumeCopy feature. Since the utility supports more than just FlashCopy, it is being distributed under a new name, *SMrepassist*, to connote replication assistance. The utility will also continue to be distributed under the name SMflashcopyassist.

For information on how to install Storage Manager utilities refer to “Installing the HBA device drivers and host software” on page 110.

4.2 FAStT management concepts and basics

This section describes common administration tasks you perform to manage a FAstT Storage Server. In addition, this section discusses and provides guidance on several storage concepts like RAID levels, controller ownership, and caching.

For specific details on how to set up and tune the system, refer to Chapter 5, “Step-by-step procedures using Storage Manager 8.4” on page 109.

4.2.1 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 physical drives.

Creating arrays and logical drives

A logical drive (or volume) is a logical structure you create on a storage subsystem for data storage. Creating arrays and logical drives is one of the most basic steps and is required before you can start using the physical disk space. That is, you divide your disk drives into arrays and create one or more logical drives inside each array.

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.

> **Tip:** Unless you have unique requirements, we highly recommend letting the system automatically create arrays. This usually ensures the most optimal balance between capacity, performance and redundancy. A manual configuration will typically not have the most optimal settings.

For more information refer to “Creating arrays and logical drives” on page 144.

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 FlashCopy, VolumeCopy or 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 specific limitations regarding specific model please refer to “Advanced functions” on page 90.

In simple configurations, you can use all of your drive capacity with just one array and create all of your logical drives in that unique array. However, this presents the following drawbacks:

- If you experience a drive failure, the rebuild process will affect all logical drives and the overall system performance will go down.
- Read/write operations of different logical drive still being made to the same physical hard drives.

With the advent of larger capacity drives and the ability to distribute logical drives across controllers, creating more than one logical drives are per array is a good way to use your storage capacity and protect your data.

Number of drives

The more physical drives you have per array the shorter the access time for read and write I/O operations.

You can determine how many physical drives should be associated with a RAID controller by looking at disk transfer rates (rather than at the megabytes per second). For example, if a hard disk drive is capable of 75 nonsequential (random) I/Os per second, then about 26 hard disk drives working together could—theoretically—produce 2,000 nonsequential I/Os per second, or enough to hit the maximum I/O handling capacity of a single RAID controller. If the hard disk drive can sustain 150 sequential I/Os per second, it then takes only about 13 hard disk drives working together to produce the same 2,000 sequential I/Os per second and keep the RAID controller running at maximum throughput.

Tip: More physical disks for the same overall capacity gives you:

- Performance: By doubling the number of the physical drives, you can expect up to a 50 percent increase in throughput performance.
- Flexibility: Using more physical drives gives you more flexibility to build arrays and logical drives according to your needs.
- Data capacity: When using RAID-5 logical drives more data space is available with smaller physical drives because less space (capacity of a drive) is used for parity.

Refer also to “Expanding arrays and logical drives” on page 90.

Channel protection

You achieve channel protection by spreading the array across multiple enclosures. There are two reasons why you would do this:

- Higher method of redundancy
- Higher performance

By implementing channel protection you reduce the chance of an array going offline by spreading it over many enclosures: If you have an enclosure failure, only one drive from the array is affected, leaving your data intact. For a RAID 1 array, you would need two enclosures; for a RAID 5 array, you would need a minimum of three enclosures. You also increase performance because the I/O request is processed by multiple ESM boards. Figure 4-6 shows an array that is created with channel protection.

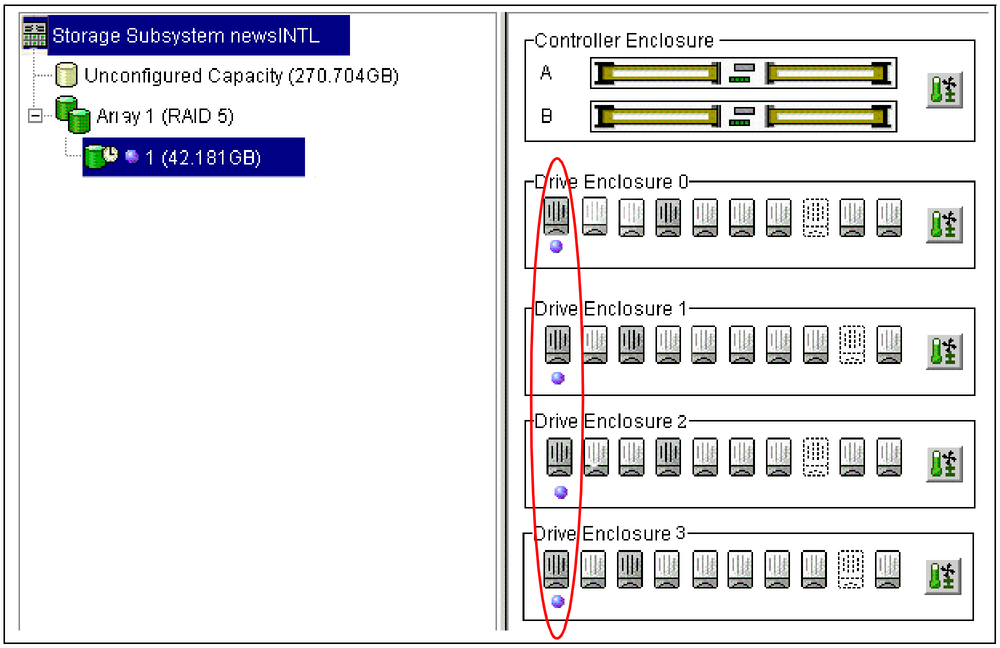

Figure 4-6 ARRAY created with channel protection

The automatic configuration feature will always choose channel protection across all available enclosures. See Figure 4-7.

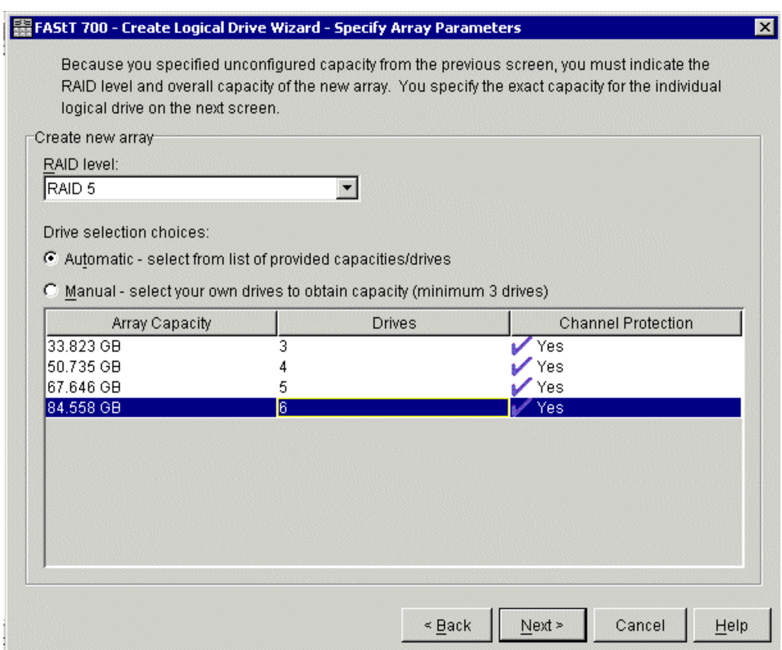

Figure 4-7 Automatic configuration feature

RAID levels

When creating an array, you must specify the RAID level.

We review here the different RAID levels and explain why we would choose this particular setting in this particular situation and readers can hopefully draw their own conclusions.

RAID 0 - For performance: RAID 0 is also known as *data striping*. It is well-suited for program libraries requiring rapid loading of large tables or, more generally, applications requiring fast access to read-only data, or fast writing. RAID 0 is only designed to increase performance; there is no redundancy, so any disk failures require reloading from backups. Select RAID Level 0 for applications that would benefit from the increased performance capabilities of this RAID Level. Never use this level for critical applications that require high availability.

RAID 1 - For availability/good read response time: RAID 1 is also known as *disk mirroring*. It is most suited to applications that require high data availability, good read response times, and where cost is a secondary issue. The response time for writes can be somewhat slower than for a single disk, depending on the write policy; the writes can either be executed in parallel for speed or serially for safety. Select RAID Level 1 for applications with a high percentage of read operations and where the cost is not the major concern.

Because the data is mirrored, the capacity of the logical drive when assigned RAID level-1 is 50 percent of the array capacity.

Some recommendations for using RAID-1 are:

- Use RAID-1 for the disks that contain our operating system. It is a good choice because the operating system can usually fit on one disk.
- Use RAID-1 for transaction logs. 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. Thus, we 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 will not complete until both writes have been done (two disks), performance of writes can be improved through the use of a write cache. When using a write cache, be sure it is battery-backed up.

RAID 3 - Sequential access to large files: RAID 3 is a parallel process array mechanism, where all drives in the array operate in unison. Similar to data striping, information to be written to disk is split into chunks (a fixed amount of data), and each chunk is written out to the same physical position on separate disks (in parallel). This architecture requires parity information to be written for each stripe of data.

Performance is very good for large amounts of data but poor for small requests because every drive is always involved, and there can be no overlapped or independent operation. It is well-suited for large data objects such as CAD/CAM or image files, or applications requiring sequential access to large data files. Select RAID 3 for applications that process large blocks of data. It provides redundancy without the high overhead incurred by mirroring in RAID 1.

RAID 5 - High availability and fewer writes than reads: RAID level-5 stripes data and parity across all drives in the array. RAID level-5 offers both data protection and increased throughput. When you assign RAID level-5 to an array, the capacity of the array is reduced by the capacity of one drive (for data-parity storage). RAID level-5 gives you higher capacity than RAID level-1, but RAID level-1 offers better performance.

RAID 5 is best used in environments requiring high availability and fewer writes than reads.

RAID-5 is good for multi-user environments such as database or file system storage where typical I/O size is small and there is a high proportion of read activity. 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.

Use write caching on RAID-5 arrays, because RAID-5 writes will not be completed until at least two reads and two writes have occurred; the response time of writes will be improved through the use of write cache (be sure it is battery backed up). RAID-5 arrays with caching can give as good a performance as any other RAID level, and with some workloads the striping effect gives better performance than RAID-1.

RAID 10 - Higher performance than RAID 1: RAID 10, also known in as RAID 0+1, implements block interleave data striping and mirroring. In RAID 10, data is striped across multiple disk drives, and then those drives are mirrored to another set of drives.

The performance of RAID 10 is approximately the same as RAID 0 for sequential I/Os. RAID 10 provides an enhanced feature for disk mirroring that stripes data and copies the data across all the drives of the array. The first stripe is the data stripe; the second stripe is the mirror (copy) of the first data stripe, but it is shifted over one drive. Because the data is mirrored, the capacity of the logical drive is 50 percent of the physical capacity of the hard disk drives in the array.

The recommendations for using RAID 10 are:

- Use RAID 10 whenever the array experiences more than 10 percent writes. RAID 5 does not perform as well as RAID 10 with a large number of writes.
- Use RAID 10 when performance is critical. Use write caching on RAID 10. Because RAID-10 write will not be completed until both writes have been done, performance of the writes can be improved through the use of a write cache (be sure it is battery backed-up).

When comparing to RAID 5:

- RAID-10 writes a single block through two writes. RAID-5 requires two reads (read original data and parity) and two writes. Random writes are significantly faster on RAID-10.
- RAID-10 rebuild takes less time than RAID-5 rebuild. If a real disk fails, RAID-10 rebuilds it by copying all the data on the mirrored disk to a spare. RAID-5 rebuilds a failed disk by merging the contents of the surviving disks in an array and writing the result to a spare.

RAID 10 is the best fault-tolerant solution in terms of protection and performance, but it comes at a cost: You must purchase twice the number of disks that are necessary with RAID 0.

Note: 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.
- RAID 10 offers higher performance than RAID 1.

Table 4-2 RAID levels comparison

RAID	Description	APP	Advantage	Disadvantage
0	Stripes data across multiple drives	IOPS MB/s	Performance due to parallel operation of the access	No redundancy; one drive fails, data is lost
1	Disk's data is mirrored to another drive	IOPS	Performance as multiple requests can be fulfilled simultaneously	Storage costs are doubled
10	Data is striped across multiple drives and mirrored to same number of disks	IOPS	Performance as multiple requests can be fulfilled simultaneously	Storage costs are doubled
3	Drives operated independently with data and parity blocks distributed across all drives in the group	MB/s	High performance for large, sequentially accessed files (image, video, graphical)	Degraded performance with 8-9 I/O threads, random IOPS, smaller more numerous IOPS
5	Drives operated independently with data and parity blocks distributed across all drives in the group	IOPS MB/s	Good for reads, small IOPS, many concurrent IOPS and random I/Os	Writes are particularly demanding

See “Changing RAID levels” on page 92 for additional information.

4.2.2 Segment size

The choice of a segment size can have a major influence on performance in both IOPS and throughput. Large segment sizes increase the request rate (IOPS) by allowing multiple disk drives to respond to multiple requests. A small segment size increases the data transfer rate (MB/s) by allowing multiple disk drives to participate in one I/O request. Use a small segment size relative to the I/O size to increase sequential performance.

You can use the performance monitor (see “Performance monitor data” on page 85) to evaluate how a given segment size affects the workload. Use the following guidelines:

- If the typical I/O size is larger than the segment size, increase the segment size in order to minimize the number of drives needed to satisfy an I/O request. This is especially true in a multi-user, database, or file system storage environment. Using a single drive for a single request leaves other drives available to simultaneously service other requests.
- If we are using the logical drive in a single-user, large I/O environment (such as for multimedia application storage), performance is optimized when a single I/O request can be serviced with a single data stripe (the segment size multiplied by the number of drives in the array that are used for I/O). In this case, multiple disks are used for the same request, but each disk is only accessed once.
- Normally, a small segment size is used for databases, normal sizes for a file server, and large segment sizes for multimedia applications.
- If we increase the segment size we gain more throughput.

Tips: The possible segment sizes available are 8k, 16k, 32k, 64k, 128k, 256k, and 512K.

- Storage Manager sets a default block size of 64 KB.
- For database applications block sizes between 4–16 Kb have shown to be more effective.
- In a large file environment like on media streaming or CAD, 128k and above are recommended.
- For a Web server, file and print server the range should be between 16–64 kb.
- A segment size of 512K can only be set through the Command Line interface, not the client GUI.

Note: You should do performance testing in you environment before you go in production with a given segment size. Segment size can be dynamically changed, but only by rewriting the data—which consumes bandwidth and impacts performance. Plan this carefully to avoid the redo.

4.2.3 Hot spare drive

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. 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.

Note: There is no definitive recommendation as to how many hot spares you should install, but it is common practice to use a ratio of one hot spare for two to three fully populated expansion enclosures (this proves to be sufficient because disk reliability has really improved nowadays).

A hot spare drive must be at least of the capacity of the configured space on the failed drive (if the failed disk drive is larger than the hot spare, reconstruction is obviously not possible). 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.

4.2.4 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 stand-alone or clustered mode.

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. Storage partitioning adapts the SAN idea of globally accessible storage to the local-storage-minded operating systems.

Storage partitioning lets you map and mask LUNs (that is why it is also referred to as LUN masking). That means once you have assigned that LUN to a host it is hidden to all other hosts connected to the same storage server. Therefore, the access to that LUN is exclusively reserved for that host.

It is a good practice to do your storage partitioning prior to connecting to multiple hosts. Operating systems such as AIX or Windows 2000 like to write their signatures to any device they can access.

> **Note:** The are limitations as to how many logical drives you can map per host. FAStT (with SM 8.4) allows up to 256 LUNs per partition (including the access LUN) and a maximum of two partitions per host. Keep these limitations in mind when planning your installation.

> **Restriction:** Most hosts will be able to have 256 LUNs mapped per storage partition. Windows NT, Solaris with RDAC, NetWare 5.1, and HP-UX 11.0 are restricted to 32 LUNs. If you try to map a logical drive to a LUN that is greater than 32 on these operating systems, the host will be unable to access it. Solaris will require use of Veritas DMP for failover for 256 LUNs.

Heterogeneous host support means that the host systems can run different operating systems. But be aware that all the host systems within a particular storage partition must run the same operating system, as all host systems within a particular storage partition have unlimited access to all logical drives in this partition. Thus, file systems on these logical drives must be compatible with host systems. To ensure this, the best way is running the same operating system on all hosts within the same partition. Some operating systems may be able to mount foreign file systems.

Storage partition topology is a collection of topological elements (default group, host groups, hosts, and host ports) shown as nodes in the topology view of the mappings view. You must define the various topological elements if you want to define specific logical drive-to-LUN mappings for host groups and/or hosts.

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 precise) can access these logical drives until they are mapped.

FAStT Storage Manager 8.4 can support up to 64 storage partitions. For the number of maximum storage partitions for a specific FAStT model see Table 4-3. Note that on some FAStT models the number of partitions also depends on the licences that have been purchased.

Table 4-3 Currently supported maximum storage partitions

FAStT Model	Number of supported partitions
FAStT 200	Up to 16
FAStT 500	Up to 64
FAStT600	1 can upgrade to 4, 8 or 16
FAStT 600 Turbo	8, can upgrade to 16, 64

FAStT Model	Number of supported partitions
FAStT 700	Up to 64
FAStT 900	16, can upgrade to 64

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 that 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 “Configuring storage partitioning” on page 150.

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 does *not* put new logical drives in the default host group by default.

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 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.

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, it is better to 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

When using Storage Manager 8.x 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.

Important: The heterogeneous hosts feature is available only with storage partitioning enabled.

Before connecting the systems always check for the latest interoperability issues on the following Internet site:

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

4.2.5 Performance monitor data

The FAStT Storage Manager has a built-in tool for monitoring and analyzing performance on each logical volume. The performance monitor data can be used to help in tuning the storage subsystem. The data collected includes (see Figure 4-8 on page 86):

- 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

Note: Cache hit percentage is only for reads when read caching is enabled (writes are always cache hits if caching is enabled, thus they are not reported).

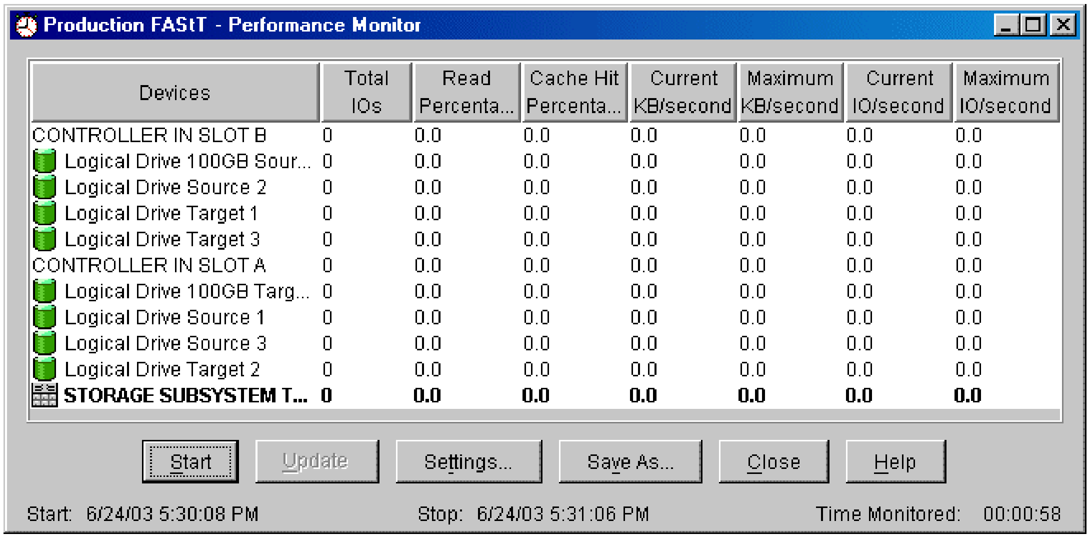

Figure 4-8 Performance Monitor

You launch the Performance Monitor from the SMclient Subsystem Management window by either:

- Selecting the **Performance Monitor** icon
- Selecting the **Storage Subsystem -> Monitor Performance** pull-down menu option
- Selecting the storage subsystem node in the Logical View or Mappings View, then **Monitor Performance** from the right-mouse pop-up menu

The Performance Monitor will not dynamically update its display if any configuration changes occur while the monitor window is open (for example, creation of new logical drives, change in logical drive ownership, and so on). The Performance Monitor window must be closed and then reopened for the changes to appear.

Note: Using the Performance Monitor to retrieve performance data can affect the normal storage subsystem performance, depending on how many items you want to monitor and what the refresh interval is.

4.2.6 Diagnostics

The FAStT Storage Server provides a high level of system availability thanks to its built-in features and reliable components. Nevertheless, a failure may still happen and it is very important that it can be identified and fixed promptly.

The FAStT Storage Server automatically logs all error conditions in its own event log. The event log is stored on reserved areas of the storage subsystem disks. As shown in Figure 4-9 on page 87, the event log is a recording of configuration events and storage subsystem component failures. Use the Event Log Viewer, accessible from the Subsystem Management window, to display the log. The event log can be copied to a local system.

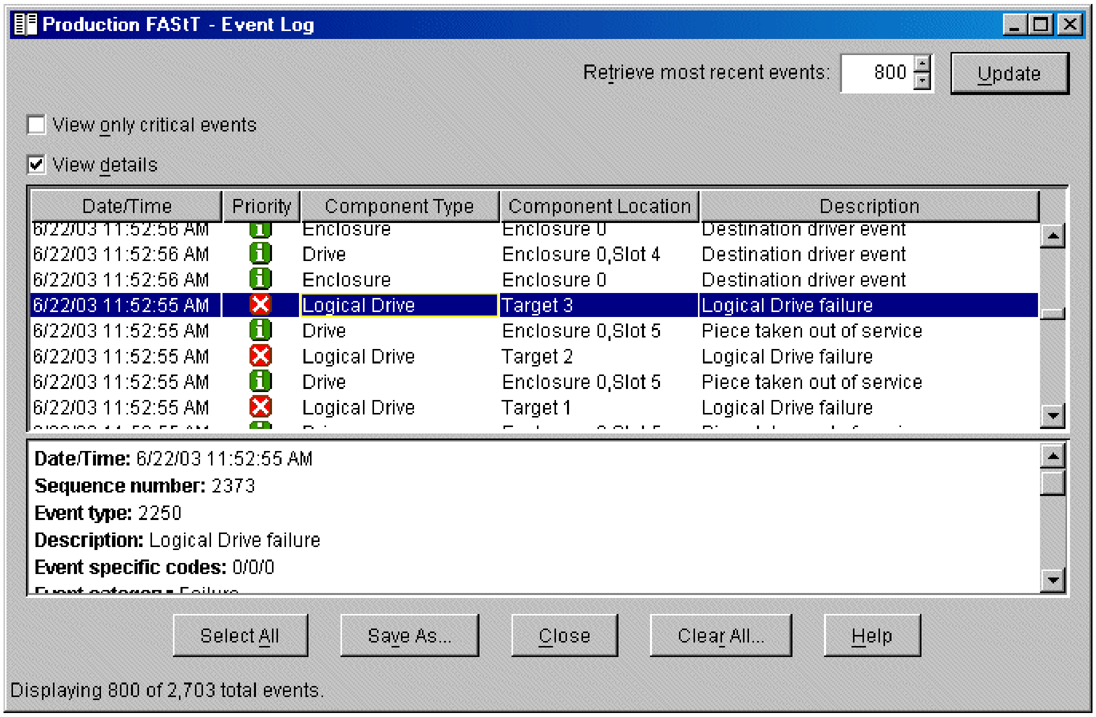

Figure 4-9 Event Log

You can find a complete list of all critical events reported in Appendix B, “Critical event descriptions” on page 307.

The *Recovery Guru* (a feature included with the FAStT Storage Manager) can interpret the event log and present a detailed step-by-step procedure on 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 “FAStT and Storage Manager 8.4 error reporting and diagnostics” on page 175 for more information.

To start the Recovery Guru use the **Storage Subsystem -> Recovery Guru** pull-down menu option in the Subsystem Management window.

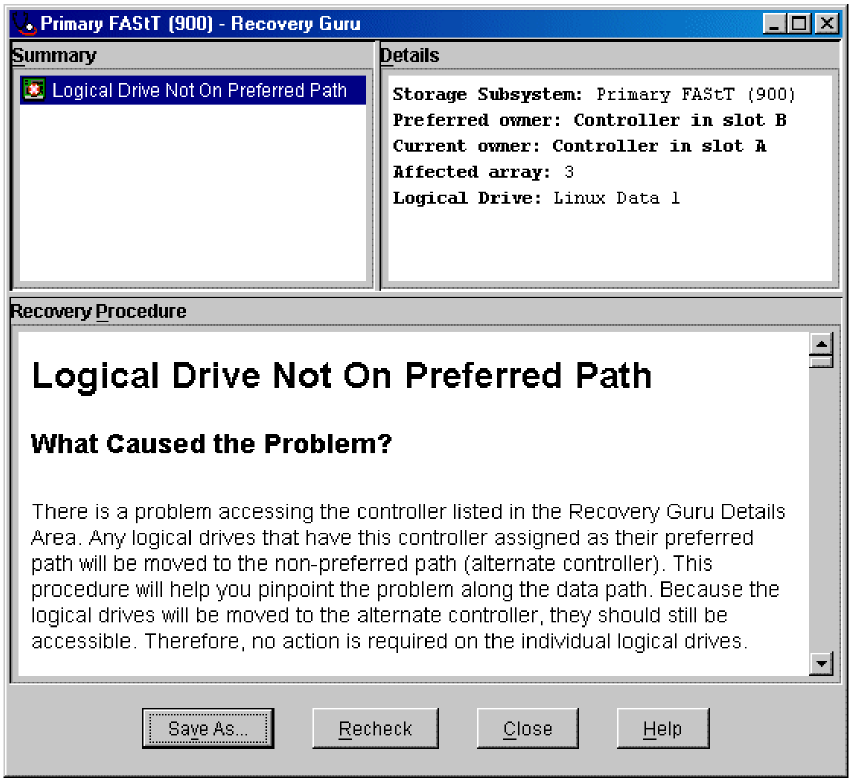

Figure 4-10 Recovery Guru

4.2.7 Password

Executing destructive commands on a storage subsystem can cause serious damage, including data loss. By specifying a password, you can protect against the execution of actions that the controller firmware deems destructive (including actions that change the state of the storage subsystem, such as the creation of logical drives and modification of cache settings, for example).

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.

For additional information about how to set up a password refer to “Initial configuration” on page 141.

4.2.8 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 4-11, or right-click and select **Execute Script**.

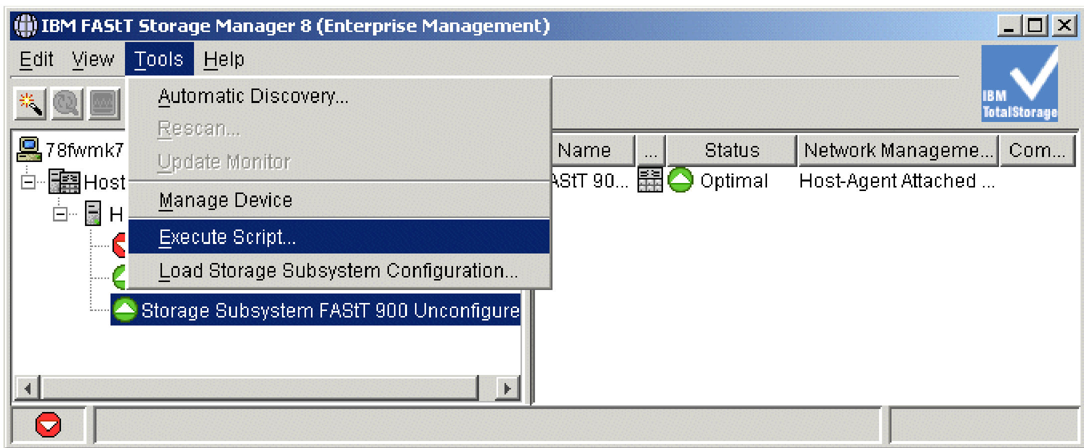

Figure 4-11 Starting the Script Editor

The Script Editor opens as shown in Figure 4-12. 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.

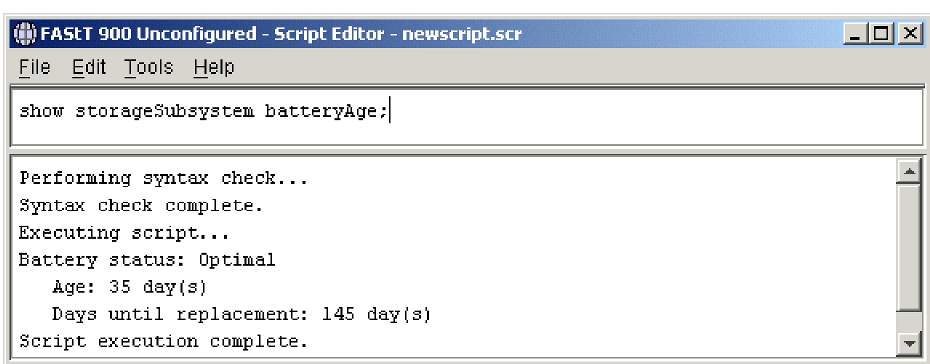

Figure 4-12 The Script Editor

For more information regarding Script Editor refer to "Script editor" on page 322.

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 depending on the storage subsystem management method you are using. For more information regarding the command line interface refer to Chapter 10, “Command Line Interface and Script Editor” on page 311.

4.3 Advanced functions

This section of the book introduces some of the advanced features of the FAStT Storage Manager.

4.3.1 Expanding arrays and logical drives

The ability to increase the available free capacity on an array *(Dynamic Capacity Expansion - DCE)* 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 be non-disruptive and 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 possible to increase the size of the logical drive. This is called *Dynamic Volume Expansion (DVE)*. 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 not supported, the expanded capacity will be unusable and you cannot restore the original logical drive capacity.

For a list of operating systems that support a dynamic increase in logical drive capacity (DVE)—and when logical drive-to-LUN mappings have already been defined—refer to:

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

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 DCE/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 “Maintaining arrays” on page 162 and in “Maintaining logical drives” on page 164.

4.3.2 Changing 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.

Note: Even though RAID migration is a non-disruptive process, we recommend carrying out this migration when I/O activity is at a minimum.

You can find more information in “Maintaining arrays” on page 162.

4.3.3 Defragmenting an array

Defragmenting an array results in the consolidation of all free capacity on the selected array. Doing this 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 defragmentation 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. You can find more information in “Maintaining arrays” on page 162.

4.3.4 Controller ownership

Each logical drive has a preferred controller of ownership. This controller normally handles all I/O requests for this particular logical drive. In other words, each logical drive is owned by one and only one controller. 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. When defining logical drives, the system normally alternates ownership between the two controllers.

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.

Balancing traffic is unfortunately not always a trivial task. For example, if an application requires large disk space to be located and accessed in one chunk, it becomes harder to balance traffic by spreading the smaller volumes among controllers.

Also, typically the load across controllers and logical drives is constantly changing. The logical drives and data accessed at any given time depend on which applications and users are active during that time period. Hence the importance of monitoring the system.

The Performance Monitor provides data 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.

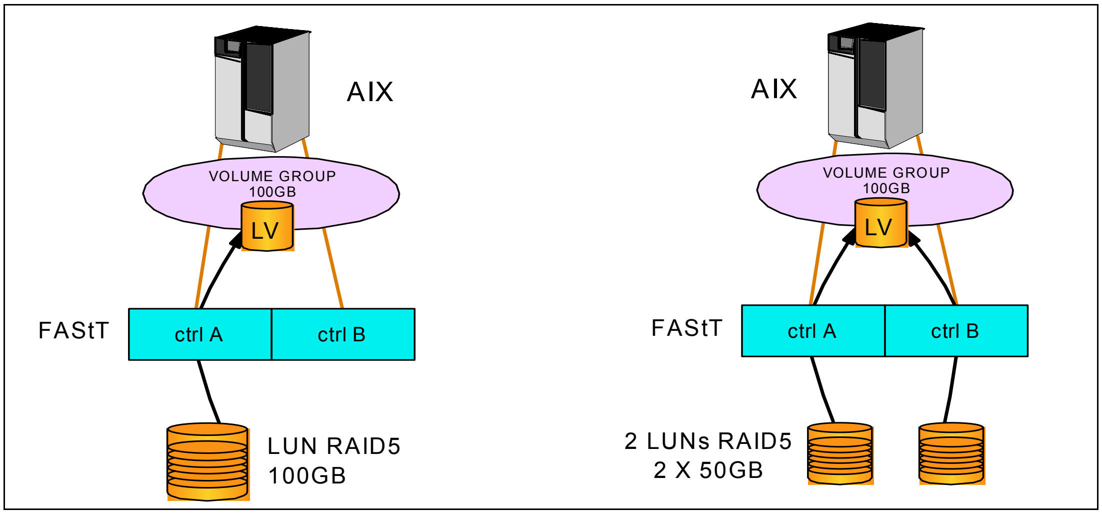

Figure 4-13 Balancing LUNs

Tip: Rules of thumb for LUN assignment and storage partitioning:

- Assign LUNs across all controllers.
- Unless you have special requirements, use the automatic feature (wizard) of Storage Manager to create your LUNs.
- If you have highly utilized LUNs, move them away from other LUNs.

The preferred controller ownership of a logical drive or array is the controller of an active-active pair that is designated to own these logical drives. The preferred controller owner is the controller that currently owns the logical drive or array.

If the preferred controller is being replaced or undergoing a firmware download, ownership of the logical drives is automatically shifted to the other controller, and that controller becomes the current owner of the logical drives. This is considered a routine ownership change and is reported with an informational entry in the event log.

There can also be a forced failover from the preferred controller to the other controller because of I/O path errors. This is reported with a critical entry in the event log, and will be reported by the Enterprise Management software to e-mail and SNMP alert destinations.

Restriction: A secondary logical drive in a remote mirror does not have a preferred owner. Instead, the ownership of the secondary logical drive is determined by the controller owner of the associated primary logical drive. For example, if Controller A owns the primary logical drive in the primary storage subsystem, then Controller A owns the associated secondary logical drive in the secondary storage subsystem. Controller ownership changes of the primary logical drive cause a corresponding controller ownership change of the secondary logical drive.

4.3.5 Cache parameters

Cache memory is an area of temporary volatile storage (RAM) on the controller that has a faster access time than the drive media. This cache memory is shared for read and write operations.

Efficient use of the RAID controller cache is essential for good performance of the FAStT storage server.

The diagram shown in Figure 4-14 on is a schematic model of the major elements of a disk storage system—elements through which data moves (as opposed to other elements such as power supplies). In the model, these elements are organized into eight vertical layers: Four layers of electronic components shown inside the dotted ovals and four layers of paths (that is, wires) connecting adjacent layers of components to each other. Starting at the top in this model, there are some number of host computers (not shown) that connect (over some number of paths) to host adapters. The host adapters connect to cache components. The cache components in turn connect to disk adapters that in turn connect to disk drives.

Here is how a read I/O request is handled in this model. A host issues a read I/O request that is sent over a path (such as a Fibre Channel) to the disk system. The request is received by a disk system host adapter. The host adapter checks whether the requested data is already in cache, in which case it is immediately sent back to the host. If the data is not in cache, the request is forwarded to a disk adapter that reads the data from the appropriate disk and copies the data into cache. The host adapter sends the data from cache to the requesting host.

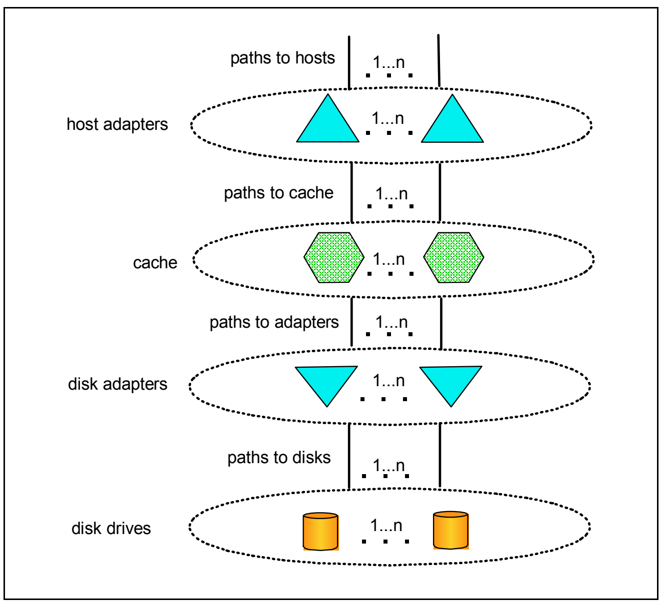

Figure 4-14 Conceptual model of disk caching

Most (hardware) RAID controllers have some form of read and/or write caching. You should obviously plan to take advantage of these caching capabilities, as they enhance the effective I/O capacity of the disk subsystem. The principle of these controller-based caching mechanisms is to gather smaller and potentially nonsequential I/O requests coming in from the host server (for example, SQL Server) and try to batch them with other I/O requests; consequently, the I/O requests are sent as larger (32 KB to 128 KB), and possibly sequential, requests to the hard disk drives. The RAID controller cache arranges incoming I/O requests by making the best use of the hard disks underlying I/O processing ability. This increases the disk I/O throughput.

There are many different settings (related to caching) that come into play. When implementing a FAStT Storage Sever as part of a whole solution you should plan at least one week of performance testing and monitoring to adjust the settings.

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

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

Figure 4-15 shows the default values when using the Create Logical Drive Wizard. Storage Manager allows you to specify cache settings for each logical drive independently, for more flexibility.

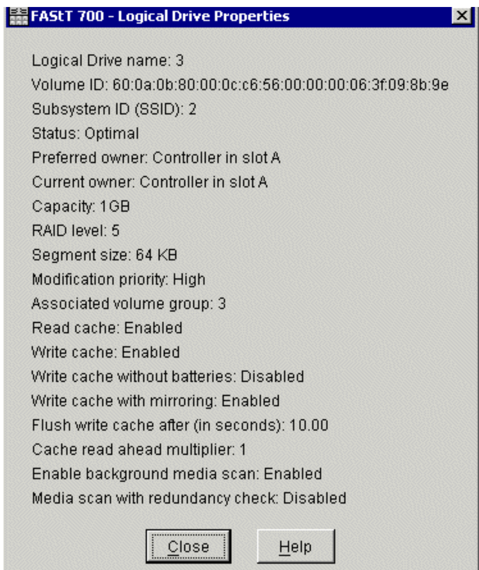

Figure 4-15 Default values use by the Create Logical Drive Wizard

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.

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.

Read caching

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.

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 zero. Otherwise each read request will unnecessarily pre-fetch additional data blocks. Since these data blocks will rarely be needed, the performance is going to be negatively impacted.

For sequential workloads, a good value would be between 1 and 4, depending on the particular environment. When using such a setting, a read request causes pre-fetching of several sequential data blocks into the cache; this speeds up subsequent disk access. This leads to a fewer number of I/O transfers (between disk and cache) required to handle the same amount of data, which is good for performance in a sequential environment. 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 a logical drive to find a proper value.

Write caching

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.

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. We recommend keeping controller write cache mirroring enabled for data integrity reasons in case of a controller failure.

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.

Write-back caching and write-through

If you configure write-through, it means that writing operations do not use cache at all. The data is always going to be written directly to the disk drives. Setting this parameter frees up cache for reading (since the cache is shared for read and write operations).

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 this 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.

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 percent. Very often the stop flushing level is set to 80 percent, too. This means the cache controller will not allow more than 80 percent 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. 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 percent 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 a 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.

4.3.6 Logical drive modification priority

The modification priority defines how much processing time is allocated 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.

The following modification priority rates are available: Lowest, low, medium, high and 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.

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 and 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. At *low* it takes approximately six times as long. At *medium* it takes approximately three and a half times as long. At *high* it takes approximately twice as long.

The synchronization progress bar at the bottom of the Mirroring tab of the Logical Drive Properties dialog displays the progress of a full synchronization.

4.3.7 Media scan

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. This is sometimes known as *data scrubbing*.

4.4 Premium features

Premium features are storage subsystem features that may not be available in the standard configuration of the storage management software. The set of available premium features depends on the model of FAStT Storage Server. Possible premium features are FlashCopy, Remote Volume Mirroring, VolumeCopy, and different sets of Storage Partitioning.

Refer to “Handling premium features” on page 347 for details.

Note: For detailed information on how to enable specific premium features refer to the following sections:

- FlashCopy - “FlashCopy: Step-by-step” on page 202
- VolumeCopy - “VolumeCopy: Step-by-step” on page 287
- Remote Volume Mirroring - “Remote Volume Mirroring: Step-by-step” on page 250

4.4.1 FlashCopy overview

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, the logical drive from which you are basing the FlashCopy, called the *base logical drive*, must be a standard logical drive in the storage subsystem. Typically, you create a FlashCopy so that an application (for example, an application to take backups) 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.

There is a limit as to how many FlashCopy relationships you can establish on a single storage subsystem. For current limitations refer to Table 4-4.

Table 4-4 Maximum number of FlashCopy volumes without remote mirror option enabled

Model	Maximum number of FlashCopy volumes	Max running FlashCopies per volume
FAStT 200	256	4
FAStT 500	256	4
FAStT 600	512	4
FAStT 700	1024	4
FAStT 900	1024	4

Note: When the Remote Mirror option is enabled the maximum number of FlashCopy volumes is further restricted. For specific limitations see Table 4-6 on page 108.

You can also create multiple FlashCopies of a base logical drive and use the copies in write mode 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 only a few seconds. Meanwhile, it creates a physical logical drive called the *FlashCopy repository logical drive* where it stores FlashCopy meta data and copy-on-write data (Figure 4-16). When the controller finishes creating the FlashCopy repository logical drive, I/O write requests to the base logical drive can resume. 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.

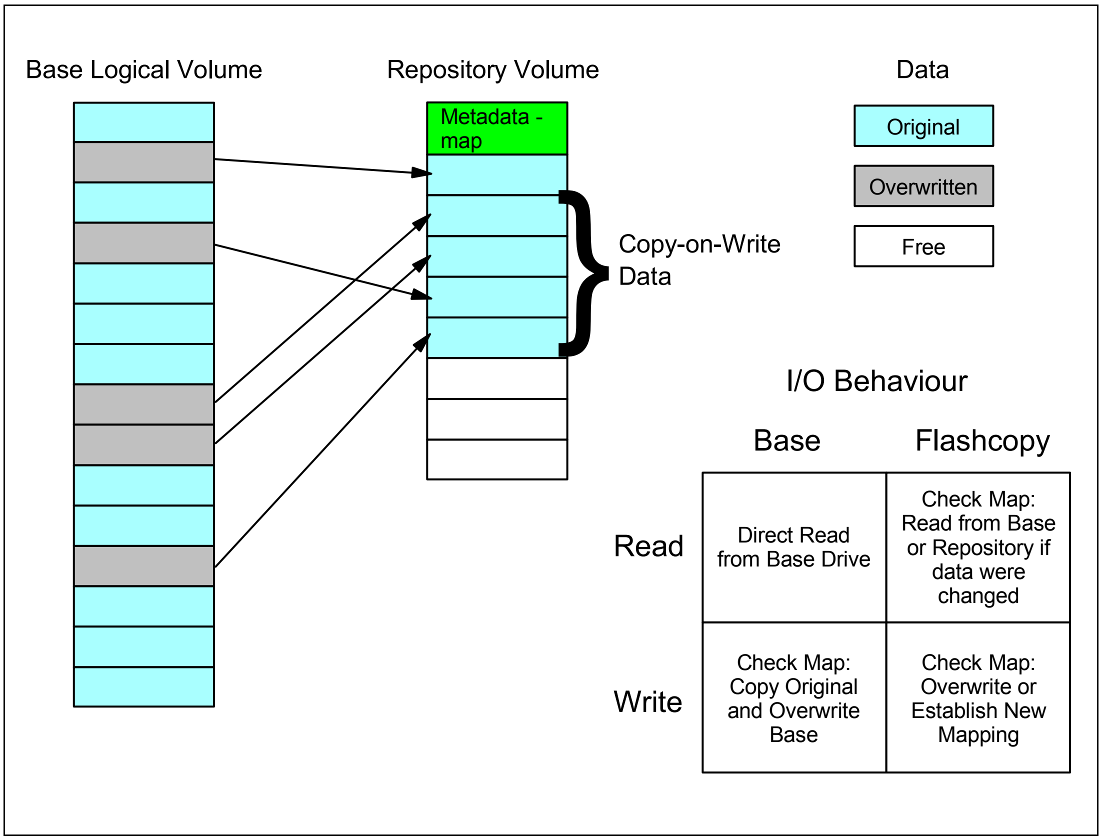

Figure 4-16 FlashCopy read and write schema

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 20 percent). 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, Linux standard disks and AIX disks. Due to operating system restrictions, FlashCopy logical drives of Windows 2000 dynamic disks *must* be mapped to a different host.

See "FlashCopy: How it works" on page 194 for detailed information.

4.4.2 VolumeCopy overview

The VolumeCopy feature (premium feature introduced with SM 8.4) is a firmware-based mechanism for replicating volume data within a storage subsystem. This feature is designed as a system management tool for tasks such as relocating data to other drives for hardware upgrades or performance management, data backup, and restoring snapshot volume data.

A VolumeCopy creates a complete physical replication of one logical drive (source) to another (target) within the same storage subsystem. The target logical drive is also referred to as *clone*.

Because VolumeCopy is a full point-in-time replication, it allows for analysis, mining, and testing without any degradation of the production logical drive performance. It also brings improvements to backup and restore operations, making them faster and eliminating I/O contention on the primary (source) logical drive.

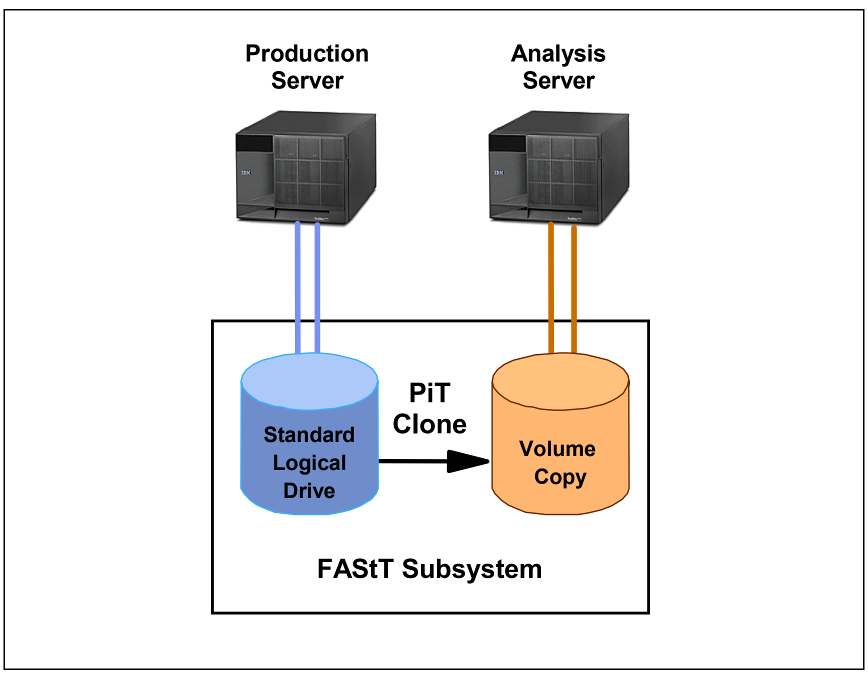

Figure 4-17 VolumeCopy

The VolumeCopy premium feature must be enabled by purchasing a feature key file. As we will see, for efficient use of VolumeCopy, FlashCopy must be installed as well.

There is a limit as to how many VolumeCopy relationships you can establish on a single storage subsystem. For current limitations refer to Table 4-5 on page 103.

Table 4-5 VolumeCopy limitations

Model	Max number of copy requests	Max running copies per array
FAStT 200	N/A	N/A
FAStT 500	N/A	N/A
FAStT 600 (Turbo)	1024	8
FAStT 700	2048	8
FAStT 900	2048	8

Users submit a VolumeCopy request by specifying two compatible volumes. One volume is designated as the source, and the other is the target. The VolumeCopy request is persistent so that any relevant result of the copy process can be communicated to the user. The feature includes a Create Copy wizard, to assist in creating a VolumeCopy, and a Copy Manager, to monitor logical drive copies after they have been created.

While a copy is actively copying data, a background operation managed by the controller firmware reads the source volume and writes the data to the target volume. If the storage array controller experiences a reset, the copy request is restored during the start-of-day processing and the copy process resumes from the last known progress boundary. During the copy process, the copy boundary of the background operation is available as a percentage of copied data with respect to the source volume capacity.

After submitting a copy request, host-initiated I/O operations to the associated volumes can be affected, depending on the state of the copy request. Read I/O requests to the source volume are unaffected by the copy request regardless of the state of the copy request.

When a background process is reading the source volume and writing the data to the target volume, host I/O access to the source and target volumes are restricted:

- All write requests to a source volume of a copy request at the time of copying are rejected.
- All write and read requests to the target volume while the copying is in progress are rejected.

These restrictions are necessary to ensure the integrity of the point-in-time copy. If the volume being copied is large, this can result in an extended period of time without the ability for a production application to make updates or changes to the data.

Important: To limit restrictions on the source volume, in practice VolumeCopy must only be used in conjunction with FlashCopy.

As illustrated in Figure 4-18 on page 104, FlashCopy which allows a point-in-time copy to be made while maintaining read/write access, enables a complete copy to be created without interrupting the I/O activity of the production volume.

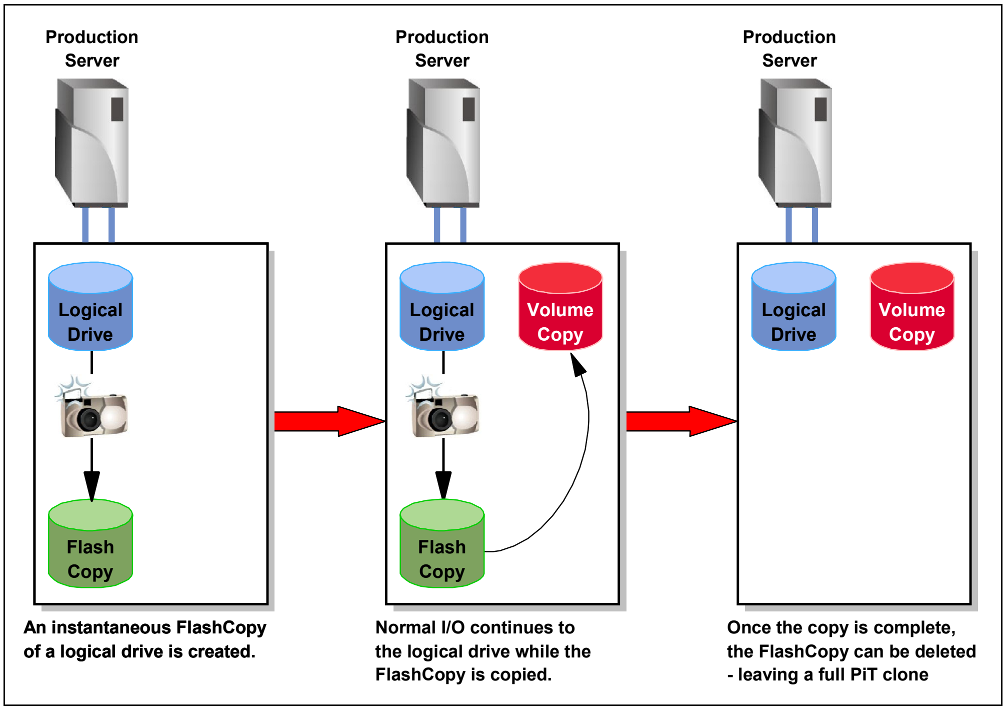

Figure 4-18 VolumeCopy integration with FlashCopy

When a copy request is established, controller volume ownership may be changed by the controller firmware because the source and target volumes must be owned by the same controller when the copying is in progress. When the background copy process completes or is stopped the target volume ownership is restored to the original owning controller. If the target volume ownership is moved from the defined preferred owner for the copy process, the normal Needs-Attention for a volume not on the preferred path will be blocked. Auto Volume Transfer (AVT) driven ownership change requests to the target volume are ignored. If the ownership of the source volume is changed during the copy process, the ownership of the target volume is also changed.

Volume and array reconfiguration includes operations such as changing a volume segment size (DSS), changing an array RAID level (DRM), expanding an array (DCE), expanding a volume (DVE), and defragmenting array free space. While these operations do not affect data accessibility, there is no compelling reason to allow reconfiguration operations to operate concurrently with background copy requests. Therefore, if a reconfiguration operation is running on a source or target candidate, submitting a VolumeCopy request is disallowed.

A copy request is submitted by highlighting a copy source volume and invoking a menu option to create a copy request. The graphical user interface guides the user through the steps of choosing an existing volume for the copy target and specifying all required parameters. The copy request is immediately started after submitting the request.

The Re-Copy menu option is used to begin the data transfer process between a source volume and a target volume that already have an entry in the Copy Manager table. The copy process is always a full copy, even if the process was stopped and then re-started.

The Stop Copy option allows the user to stop a copy that is in progress. It can also be used to remove the Needs Attention condition if a copy process has failed and the user wishes to maintain the copy relationship between the affected source and target.

Copy requests are managed through a separate copy management window where all copy requests are shown as source and target pairs with associated state information. The copy management window will automatically refresh when the information for any copy pair in the table changes, and will poll the storage array at regular intervals and refresh its display accordingly to ensure that the data it displays is as accurate as possible.

The background copy operation priority is user defined to one of five priorities. The lowest priority has the lowest impact on overall I/O performance; the highest has a higher impact on I/O performance. The priority setting can be changed by highlighting a copy request in the copy management display and invoking a menu option to change the copy priority. It can also be set on the original copy request and re-copy operations.

A copy request is removed by highlighting the source-target pair in the Copy Management window and invoking a menu option to remove the copy request.

Refer to the Chapter 9, "VolumeCopy" on page 281, for procedures and more information about the VolumeCopy feature.

4.4.3 Remote Volume Mirroring (RVM) overview

In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Mirror option enables you to promote a second storage subsystem to take over responsibility for normal input/output (I/O) operations. The storage subsystems participating in a Remote Mirroring are called primary and secondary storage subsystems, or local and remote storage subsystems. These names are used interchangeably to describe remote mirror setups or concepts. The names do not refer to the location of storage subsystems or the role storage subsystems have in a remote mirror relationship. The Remote Volume Mirroring premium feature is used for online, real-time data replication between storage subsystems over a remote distance.

You must purchase keys and enable the feature on both primary and secondary storage subsystems.

The primary and secondary role in a remote mirror setup is implemented at the logical drive level instead of at the storage subsystem level. A storage subsystem can have all logical drives participating in a remote mirror relationship be in either a primary or secondary role only. 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 role, it counts towards the maximum number of mirror logical drive pairs that can be defined in a storage subsystem.

Note: The IBM FAStT200 (machine type 3542) and IBM FAStT600 base or Turbo (machine type 1722) storage subsystems do not support the Remote Mirror option.

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*.

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 fail over 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* (see Figure 4-19).

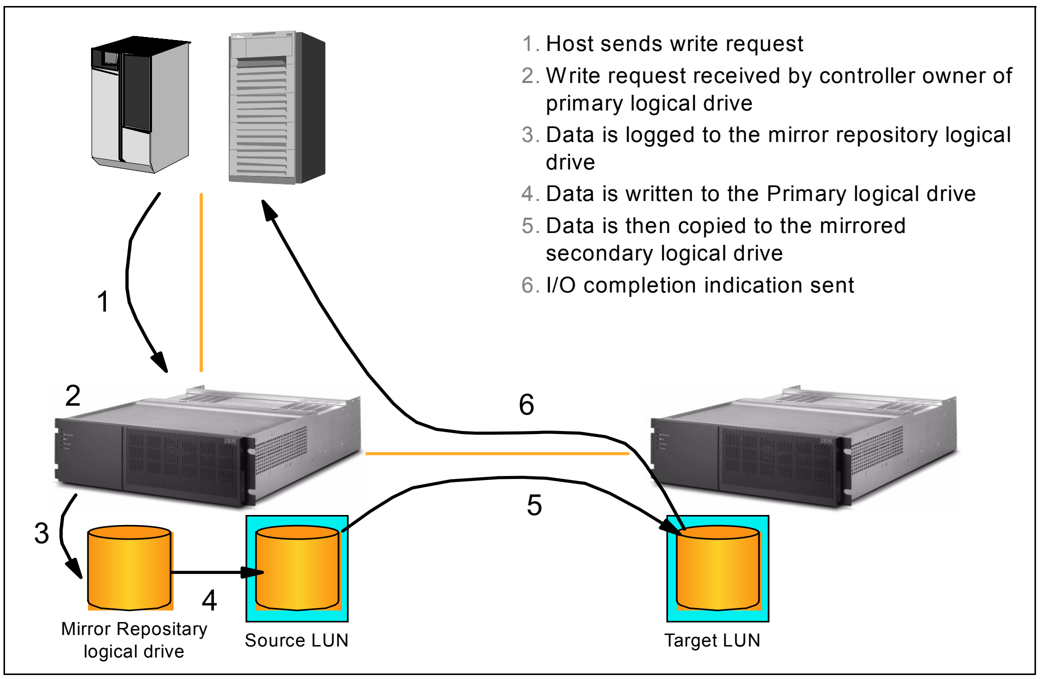

Figure 4-19 Remote volume mirroring

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 site (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 site.

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 site. 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 site, 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 site. 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 site are controlled by the primary site, 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 site 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.

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:

- 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.
- RAID level, caching parameters, and segment size can be different on the two mirrored logical drives.

There is a limit as to how many logical drives can be created in a single storage subsystem and how many FlashCopy relationships can be defined when the Remote Mirror option is enabled. Table 4-6 shows the limitations for each FAStT model.

Table 4-6 Logical drives limitations using Remote Mirror option

Model	Maximum logical drives with RM enabled	Maximum of remote mirror pairs	Maximum FlashCopy relations with RM enabled
FAStT200	N/A	N/A	N/A
FAStT500	128	8	64
FAStT 600	N/A	N/A	N/A
FAStT 700	256	32	128
FAStT 900	256	32	128

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

5

Step-by-step procedures using Storage Manager 8.4

This chapter explains the installation of the device drivers and other required software for several host server operating system environments. This is followed by a step-by-step procedure on how to install the FAStT Storage Manager client and how to use it to configure the Storage Server, define arrays and volumes, and enable storage partitioning. We also discuss the different ways to monitor the FAStT Storage Server.

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.

5.1 Installing the HBA device drivers and host software

This section will describe the installation of the device drivers and host software required on the host server(s) to which the FAStT Storage Server is attached.

We have already explained that the host software package consists of various elements; some elements might not be required (or available) depending on the host operating system:

- FAStT Storage Manager client (SMclient) and Runtime environment (SMruntime)
- FAStT RDAC
- FAStT Management Suite Java (MSJ)
- FAStT Storage Manager agent (SMagent)
- FAStT Utilities

The FAStT Storage Manager client and runtime elements are available for all operating systems (the runtime element may be incorporated into the client package).

The RDAC package provides redundancy within Fibre Channel paths and is available for Windows, AIX and Solaris.

Linux uses the FAStT MSJ, which allows the definition of a preferred path and an alternate path for each logical drive (achieving load balancing and path protection).

The FAStT agent is only available for systems that allow in-band management.

5.1.1 Installing the HBAs and updating their BIOS

If the HBAs are not already installed, shut down your system and install the adapters.

At this stage it is best to either leave the cards disconnected from the SAN or have the SAN switched off.

The very first step is to make sure that you have the appropriate BIOS level for the particular HBAs you are using. You can obtain the latest BIOS update files at:

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

Create the bootable diskette and use it to boot the server. You should end up with a DOS `C:\` prompt. The update will do all like HBAs in the server at the same time. (Dissimilar HBAs need to be updated separately.)

- Type the following to update the HBA BIOS:

  ```
  flasutil /f /l
  ```

- Type the following to set the defaults:

  ```
  flasutil /u
  ```

Remove the boot diskette and restart the server.

Note: The FAStT adapter BIOS update program will only update like adapters. If you have a server that contains dissimilar adapters, you must manually update each adapter at a time. Type 'flasutil" without any options to get to the list of adapters.

Important: For Linux servers with Qlogic adapters you need to update the following settings in the QLogic BIOS:

- Host adapter settings

  Loop reset delay - 8

- Advanced adapter settings
  - LUNs per target - 0
  - Enable Target Reset - Yes
  - Port down retry count - 12

For Netware you will need to update the following setting in the QLogic BIOS: Advanced adapter settings: LUNs per target - 32.

5.1.2 Install the device driver for the HBAs

To update the device driver, follow the instructions supplied with the device driver corresponding to your HBA and choice of operating system; these are maintained at and can be downloaded from:

`http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-50176`

5.1.3 Install FAStT software on a Windows host server

This section covers the installation of the FAStT Management software components for a Windows 2000 and Windows Server 2003 host environment.

Important: In preparation of this publication we used pre-release code of Storage Manager 8.4 to experiment and document the installation under different software platforms. Beware that some instructions might be different when SM 8.4 is made generally available.

Note: Storage Manager 8.4 for Windows 2000 will not install on Windows Server 2003; you must use the appropriate package.

The host software for Windows includes the following packages:

- FAStT SMclient
- FAStT RDAC driver
- FAStT SMagent
- FAStT Utilities

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 Microsoft Virtual Machine

For systems running Windows Server 2003 32-bit editions, the Microsoft Virtual Machine is required to support the following Storage Manager 8.4 components or options:

- SMagent

- Event Monitor option of the SMclient

To install Microsoft Virtual Machine on management stations and host computers perform the following steps:

1. Close all programs.
2. Insert the IBM FAStT Storage Manager Version 8.4 CD into the CD-ROM drive.
3. Click **Start -> Run**. The Run window opens.
4. Click **Browse**. The Browse window opens.
5. Select the CD-ROM drive.
6. Select the folder that pertains to your operating system then click **Open**.
7. Select the **MSJavx86-3805.exe** file. Click **Open** and follow the instructions on the screen. The Setup Complete window opens.
8. Click **No, do not restart the system at this time**. Click **Finish**.
9. Continue with "Installing the SMclient" on page 48.

Installing the SMclient

Note that the SMclient is only required on the host server attaching to the FAStT Storage Server when you use in-band management.

To install SMclient on Windows operating systems, perform the following steps:

1. Insert the IBM FAStT Storage Manager Version 8.4 CD into the CD-ROM drive.
2. Click **Start-> Run**. The Run window opens.
3. Click **Browse**. The Browse window opens.
4. Select the CD-ROM drive.
5. Select the folder that pertains to your operating system and click **Open**.
6. Select the displayed .exe file. Click **Open**. The InstallShield Wizard window opens, then the Welcome window opens.
7. Click **Next**. The End User License Agreement window opens.
8. Read the license agreement and click **Yes**. The Event monitor window opens.
9. At the prompt, click **Yes** if you want to install the Event Monitor. The Choose Destination Location window opens.
10. Click **Browse** to change the destination location.
11. Click **Next** to start the installation. When the installation is complete, the Operation Complete window opens.
12. Click **Finish**

At the end of step 8, the window shown in Figure 5-1 on page 113 is displayed, asking if you want to install the Event Monitor service.

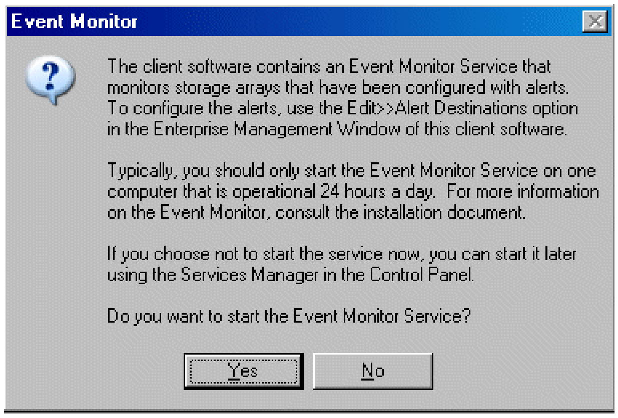

Figure 5-1 Installing the FAStT Storage Manager Event Monitor

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. (When installing the client software on a management workstation, which is not directly attached to the FAStT Storage Server or is not operated 24 hours, you do not need to install this service.)

Perform the following procedures depending on your operating system:

- For Windows 2000:
  a. Read the license agreement and click **Yes**. The Event monitor window opens.
  b. At the prompt, click **Yes** if you want to install the Event Monitor. The Choose Destination Location window opens.
  c. Click **Browse** to change the destination location.
  d. Click **Next** to start the installation. When the installation is complete, the Operation Complete window opens.
  e. Click **Finish**.
- For Windows Server 2003:
  a. Read the license agreement and click **Yes**. The Choose Destination Location window opens.
  b. Click **Browse** to change the destination location. Click **Next** to continue with the installation.
  c. The Event monitor window opens. Click **Yes** if you want to install the Event Monitor. Otherwise, click **No**.
  d. The Setup Status window opens showing the installation progress. When the installation is complete, the Operation Complete window opens.
  e. Click **Finish**.

Installing the RDAC driver

To install RDAC on Windows operating systems, perform the following steps:

1. Insert the IBM FAStT Storage Manager Version 8.4 CD into the CD-ROM drive.
2. Click **Start -> Run**. The Run window opens.
3. Click **Browse**. The Browse window opens.

4. Select the CD-ROM drive.
5. Select the directory that pertains to your operating system.
6. Select the displayed .exe file. Click **Open**. The InstallShield Wizard window opens then the Welcome window opens.
7. Click **Next**. The End User License Agreement window opens.
8. Click **Yes** to start the installation. Perform the following steps depending on your operating system:
   a. For Windows 2000, after you install RDAC, the InstallShield Wizard Complete window opens. Continue with step 10.
   b. For Windows Server 2003, the Installation/Removal status window opens displaying the RDAC installation progress. Continue with step 9.
9. After the installation is complete, the Operation Complete window opens.
10. Select **Yes I want to restart my computer now**. Click **Finish**.
11. Continue with “Verifying the RDAC installation” on page 114.

Verifying the RDAC installation

The verification procedure is given for Windows 2003 and Windows Server 2003 next.

Windows 2000

To verify that you have installed RDAC correctly, perform the following steps:

1. Click **Start -> Programs -> Administrative Tools -> Computer Management**.The Computer Management window opens. Go to the \System Tools\System Information\Software environment\drivers directory.
2. Scroll through the list of device drivers until you find rdacfltr.
3. Verify that rdacfltr is displayed with the state type Running and status OK. If it is not, repeat the procedure outlined in “Installing the RDAC driver” on page 113; otherwise, continue as follows.
   a. If you are managing the storage subsystems using the host-agent method, continue with “Installing SMagent” on page 115.
   b. If you are managing your storage subsystems using the direct-management method, go to “Installing SMutil” on page 116.

Windows Server 2003

To verify that you have installed RDAC correctly, perform the following steps:

1. Click **Start -> All Programs -> Accessories -> System Tools -> System information**. The System Information window opens. Go to the \System Summary\Software Environment\System Drivers list.
2. Scroll through the list and select one of the following device drivers based on the host computer processor type.
3. Verify that either rdacbus32 or rdacbus64 displays with the state type Running and status OK. If it does not display, repeat the procedure in “Installing the RDAC driver” on page 113; otherwise, continue with installing “Storage Manager 8.3 Utility” on page 56.
   – If you are managing the storage subsystems using the host-agent method, continue with “Installing SMagent” on page 115.
   – If you are managing your storage subsystems using the direct-management method, go to “Installing SMutil” on page 116.

Installing SMagent

This section describes how to install SMagent software on Windows operating systems. You must install the SMagent software if you want to manage the storage subsystem using the host-agent management method.

> **Note:** The SMagent software package is not available for the Windows Server 2003 64-bit editions.

Before you install the software, ensure that the following conditions are met:

- You have administrator privileges on the host computer.
- The host computer has at least 1 MB of available disk space.
- RDAC is installed on each host computer. If not, refer to "RDAC" on page 51.
- For Windows Server 2003 make sure Microsoft Virtual Machine is installed; refer to "Installing the Microsoft Virtual Machine" on page 111.
- All storage subsystems are connected to the host computer and are powered on.

To install SMagent on Windows operating systems, perform the following steps:

1. Close all programs.
2. Insert the IBM FAStT Storage Manager Version 8.4 CD in the CD-ROM drive.
3. Click **Start -> Run**. The Run window opens.
4. Click **Browse**. The Browse window opens.
5. Select the CD-ROM drive.
6. Select the folder that pertains to your operating system, then click **Open**.
7. Select the displayed .exe file. Click **Open**. The InstallShield window opens, showing the status of the file unpacking. The Welcome window opens.
8. Click **Next**. The End User License Agreement window opens.
9. Read the license and click **Yes** to begin the installation. The software installation ends, and the Operation Complete window opens.
10. Click **Finish**.

Verifying the SMagent installation

This section provides instructions on how to verify that you have installed SMagent correctly on Windows operating systems.

Windows 2000

To verify that you have installed the SMagent correctly, perform the following steps:

> **Note:** Make sure that all storage subsystems are connected to the host computer and are powered on.

1. Click **Start -> Programs -> Administrative Tools -> Services**. The Services window opens.
2. Scroll through the list of services until you find IBM FAStT Storage Manager 8 Agent.
3. If IBM FAStT Storage Manager 8 Agent does not display with a startup type of Started, click **Start**.

4. Perform one of the following steps:
   – If you plan to use the host-agent software to manage one or more storage subsystems, go to “Installing SMutil” on page 116.
   – If you do not plan to use the host-agent software to manage one or more storage subsystems, perform the following procedure to disable the host-agent service.
5. Click **Start -> Programs -> Administrative Tools -> Computer Management**. The Computer Management window opens.
6. Click **Services**.
7. From the list of displayed services, right-click **IBM FAStT Storage Manager 8 Agent**.
8. Click **Properties -> Startup Type -> Manual**.
9. Click **OK**.
10. Continue with “Installing SMutil” on page 116.

Windows Server 2003

To verify that you have installed the SMagent correctly, perform the following steps.

Note: Make sure that all storage subsystems are connected to the host computer and are powered on.

1. Click **Start -> Administrative Tools -> Services**. The Services window opens.
2. Scroll through the list of services until you find **IBM FAStT Storage Manager 8 Agent**.
3. If IBM FAStT Storage Manager 8 Agent does not display with a startup type of Started, click **Start**.
4. Perform one of the following steps: If you plan to use the host-agent software to manage one or more storage subsystems, see “Installing SMutil” on page 116. If you do not plan to use the host-agent software to manage one or more storage subsystems, perform the following steps to disable the host-agent service.
5. **Click Start -> Administrative Tools -> Computer Management**. The Computer Management window opens.
6. Click **Services**.
7. From the list of displayed services, right-click IBM FAStT Storage Manager 8 Agent.
8. Click **Properties** from the pull-down menu. The IBM FAStT Storage Manager 8 Agent Properties window opens.
9. In the Startup Type list, select **Manual**. Click **OK**. The IBM FAStT Storage Manager 8 Agent Properties window closes.

Installing SMutil

This section describes how to install SMutil. This component contains utilities that will register and map new logical drives to the operating systems.

To install SMutil on Windows operating systems, perform the following steps:

1. Close all programs and insert the IBM FAStT Storage Manager Version 8.4 installation CD in the CD-ROM drive.
2. Click **Start -> Run**. The Run window opens.
3. Click **Browse**. The Browse window opens.
4. Select the CD-ROM drive.

5. Select the folder that pertains to your operating system then click **Open**.
6. Select the displayed .exe file. Click **Open**. The InstallShield Wizard window opens, displaying the status of the files unpacking. The Welcome window opens.
7. Click **Next**. The End User License Agreement window opens.
8. Click **Yes** to begin the installation. After the installation is complete, the Operation Complete window opens.
9. Click **Finish**.

Verifying the SMutil installation

To verify that you have installed the SMutil correctly on Windows operating systems perform the following steps:

1. Go to the following directory: *installation_directory*\IBM_FAStT\Util, where *installation_directory* is the directory where you installed the SMutil.
2. Verify that the directory contains the following files:
   - hot_add.exe
   - SMdevices.bat
   - SMflashcopyassist.exe

5.1.4 Install FAStT software on an AIX host server

This section covers the installation of the FAStT Management software components attaching to a host server running the AIX operating system.

Because AIX is only supported with out-of-band management, the host software consists of only one package—the FAStT Storage Manager 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.

Since 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 “Configuring storage partitioning” on page 150 to learn how to define mappings to a logical drive.

Before installing SMclient, make sure that the following conditions are met:

- The AIX host on which you are installing SMruntime meets the minimum hardware and software requirements described in “*Installation and Support Guide for AIX, HP-UX, and Solaris*” for the release of Storage Manager you are installing.
- You have prepared the correct filesets for an AIX system. You can see the list of required filesets in “*Installation and Support Guide for AIX, HP-UX, and Solaris*”, which is provided with the machine or can be downloaded from the Internet site along with FAStT Storage Manager code. Go to:

  `http://www.ibm.com/storage/techsup.htm`

  Look for the FAStT Storage Server link under Technical Support.

To properly install the client software under AIX, you must first install SMruntime, followed by SMclient. (SMruntime provides the Java runtime environment required to run the SMclient.)

Installing SMruntime

You may need to adjust the following instructions for the specifics of your installation. No restart is required during the installation process.

1. Install SMruntime by typing the following command:

```
# installp -a -d /complete path name/SMruntime.aix-08.30.65.00.bff SMruntime.aix.rte
```

2. Verify that the installation was successful by typing the following command:

```
# lslpp -ah SMruntime.aix.rte
```

The verification process should return a table that describes the software installation, including the install package file name, version number, action, and action status:

```
# lslpp -ah SMruntime.aix.rte

Fileset         Level     Action       Status       Date         Time
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  SMruntime.aix.rte
              8.40.6500.0   COMMIT       COMPLETE     06/20/03     14:10:33
              8.40.6500.0   APPLY        COMPLETE     06/20/03     14:10:33
```

Installing SMclient

You may need to adjust the following instructions for the specifics of your installation. No restart is required during the installation process.

1. Install SMclient by typing the following command:

```
# installp -a -d /complete path name/SMclient.aix-08..33.G5.03.bff SMclient.aix.rte
```

2. Verify that the installation was successful by typing the following command:

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

The verification process should return a table that describes the software installation, including the install package file name, version number, action, and action status:

```
[# lslpp -ah SMclient.aix.rte
  Fileset         Level     Action       Status       Date         Time
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  SMclient.aix.rte
             98.40.6500.0   COMMIT       COMPLETE     06/20/03     14:14:28
             98.40.6500.0   APPLY        COMPLETE     06/20/03     14:14:28
```

Performing the initial configuration on AIX hosts

Complete the installation by defining logical drives. For instructions on how to do this please refer to “Creating arrays and logical drives” on page 144. Logical drives, partitioning, and all other related tasks can be done also from the AIX Storage Manager client. The look and feel of the interface are similar on all operating systems currently supported (see Figure 5-2 on page 119).

To start the client, issue the following command:

```
# /usr/SMclient/SMclient
```

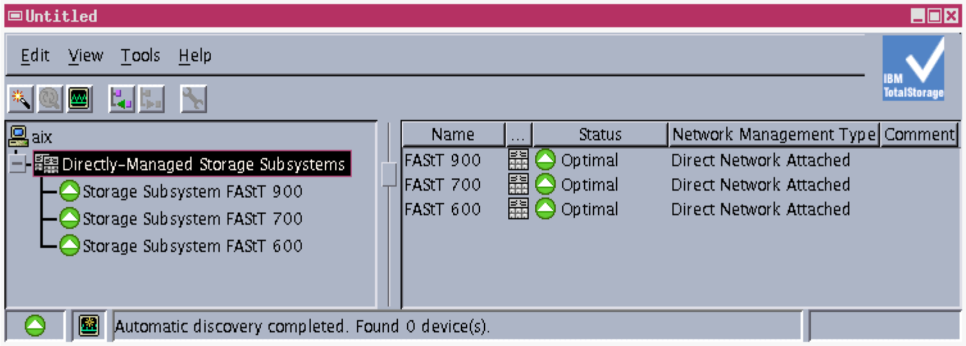

Figure 5-2 Storage Manager client 8.4 for AIX

After you set up an AIX host group, perform the following steps to verify that the host ports match the AIX host:

1. Type the following command:

```
# lsdev -Cc adapter | grep fcs
```

   A list that contains all the HBAs that are in the system is displayed, as shown in the following example:

```
# lsdev -Cc adapter | grep fcs
fcs0    Available 20-58    FC Adapter
fcs1    Available 20-60    FC Adapter
```

2. Identify the fcs number of the HBA that is connected to the FAStT.
3. Type the following command:

```
# lscfg -vl fcs? |grep Network
```

   Where *fcs?* is the fcs number of the HBA that is connected to the FAStT.

   The network address number of the HBA is displayed, as in the following example:

```
# lscfg -vl fcs0 | grep Network
     Network Address.............10000000C926B08F
```

4. Verify that the network address number matches the host port number that displays in the host partition table of the FAStT SMclient.
5. Repeat this procedure to verify the second or more host ports.

Installing the RDAC driver

You need the following filesets:

- devices.fcp.disk.array.rte - RDAC software
- devicess.fcp.disk.array.diag - RDAC software
- devices.fcp.disk.rte - FC Disk Software
- devices.common.IBM.fc.rte - Common FC Software

Depending on the HBA, you need:

- devices.pci.df1000f7.com - Required for 6227 and 6228 adapter
- devices.pci.df1000f7.rte - Required for 6227 adapter
- devices.pci.df1000f7.diag - Required for 6227 adapter
- devices.pci.df1000f9.rte - Required for 6228 adapter
- devices.pci.df1000f9.diag - Required for 6228 adapter

Before installing the RDAC driver always check prerequisites for a list of required driver version level filesets on the AIX system. Prerequisites can be found in “*Installation and*

Support Guide for AIX, HP-UX, and Solaris" provided with the machine or with the latest Storage Manager software.

Use the **lslpp** command to verify that the correct driver versions are installed:

```
# lslpp -ah devices.fcp.disk.array.rte
```

The RDAC driver must be installed on *all* AIX hosts that will be attached to a FAStT storage subsystem. The RDAC driver creates the following devices that represent the FAStT storage subsystem configuration:

- dar (disk array router): Represents the entire array, including the current and the deferred paths to all LUNs (hdisks on AIX).
- dac (disk array controller devices): Represent a controller within the storage subsystem. There are two dacs in the storage subsystem.
- hdisk: These devices represent individual LUNs on the array.

When these devices are configured, the Object Data Manager (ODM) is updated with default parameters. In most cases and for most configurations, the default parameters are satisfactory. However, there are some parameters that can be modified for maximum performance and availability. See *"Installation and Support Guide for AIX, HP-UX, and Solaris"* provided with the Storage Manager software.

After the FAStT storage subsystem has been set up, volumes have been assigned to the host, and the RDAC driver has been installed, you must verify that all of your FAStT device names and paths are correct and that AIX recognizes your dars, dacs, and hdisks.

You must do his before you mount file systems and install applications. Type the following command to probe for the new devices:

```
# cfgmgr -v
```

Next, use the **lsdev -Cc disk** command to see if the RDAC software recognizes each FAStT volume correctly:

- FAStT200 volume as a "3542 (200) Disk Array Device"
- FAStT500 volume as a "3552 (500) Disk Array Device"
- FAStT600 volume as a "1722 (600) Disk Array Device"
- FAStT700 volume as a "1742 (700) Disk Array Device"
- FAStT900 volume as a "1742-900 Disk Array Device"

The following example illustrates the results of the command for a set of FAStT700 LUNs.

```
#lsdev -Cc disk
hdisk0 Available 10-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 20-58-01     1742     (700) Disk Array Device
hdisk2 Available 20-60-01     1742     (700) Disk Array Device
```

After the operating system device names are found, those names must be correlated to the preferred and failover paths of the FAStT device, and then from each path to its associated logical drive.

AIX provides the following commands to help you determine the FAStT configuration, and to get information about device names and bus numbers:

- **lsdev**

  Displays devices and their characteristics. The **lsdev** command shows the state of the devices at startup time, or the last time that the **cfgmgr -v** command was run.

- **lsattr**

  Displays device attributes and possible values. Attributes are only updated at startup time, or the last time that the **cfgmgr -v** command was run.

- **fget_config**

  Displays controllers and hdisks that are associated with a specified FAStT (dar). The **fget_confg** command shows the current state and volume (hdisk) ownership.

There are several ways to correlate a system's configuration and monitor the state of FAStT storage subsystems.

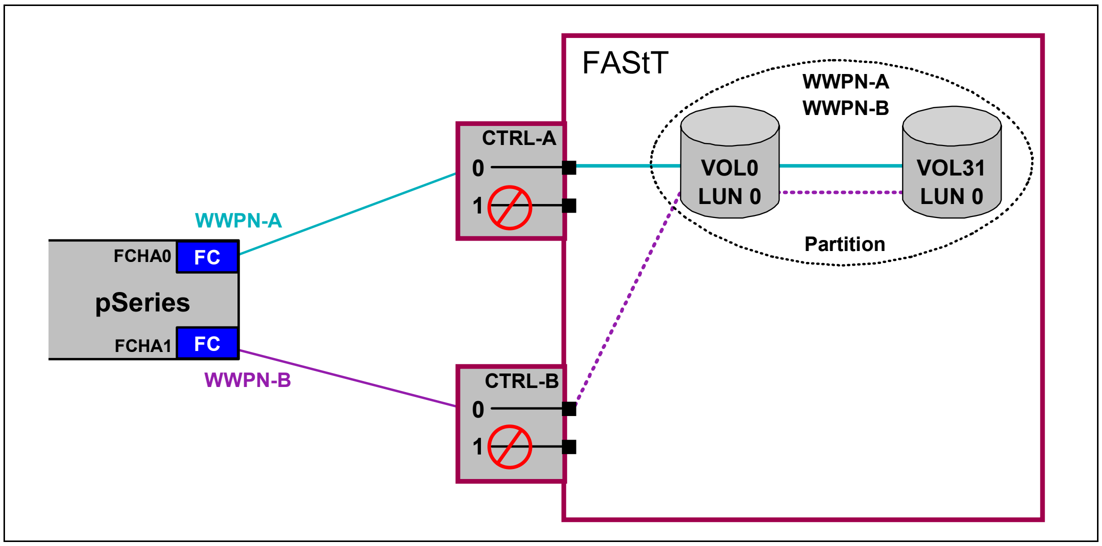

Figure 5-3 AIX host on FAStT

Example 5-1 uses the **lsdev** command to show the status of the dar, which represents the FAStT storage subsystem illustrated by Figure 5-2 on page 119. This example shows dar as a machine type 1742, which is a FAStT700. It is in the Available state, which is the state at the time the device was last configured by AIX.

The example also shows the status of two dacs, which represent the FAStT storage subsystem controllers. The third column shows the location code. In this example, each dac has its own location or path, which are represented by the values 20-58-01 and 20-60-01. Each AIX system has its own set of location codes that describe the internal path of that device, including bus and host adapter locations. See the service manual for your system type to identify device locations.

Example 5-1 Status of dar

```
# lsdev -C | grep dar
dar0       Available             1742     (700) Disk Array Router

# lsdev -C | grep dac
dac0       Available 20-58-01    1742     (700) Disk Array Controller
dac1       Available 20-60-01    1742     (700) Disk Array Controller
```

The example below uses the **lsdev** command to show the status and location codes of two FAStT700 hdisks. Notice that the location codes of the hdisk1 correspond to the same location code of dac0 in the previous example, and that the even hdisk2 corresponds to the same location code of dac1. This means that the preferred paths for I/O for hdisk1 is through dac0. The failover path would be through dac1. Conversely, the preferred path for the hdisk2 would be through dac1, and failover path through dac0.

Example 5-2 Status and location codes

```
#lsdev -Cc disk
hdisk0 Available 10-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 20-58-01     1742     (700) Disk Array Device
hdisk2 Available 20-60-01     1742     (700) Disk Array Device
```

The **fget_config** command displays the state of each controller in a FAStT array, and the current path that is being used for I/O for each hdisk.

Example 5-3 shows that both controllers (dac0 and dac1) are in the Active state. This is normal when the FAStT storage subsystem is configured correctly. Other possible states could be:

- NONE: The controller is not defined or is offline.
- RESET: The controller is in the reset state.

Example 5-3 Controller status

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac1-hdisk2
```

The **lsattr** command provides detailed information about a volume, including information that allows you to map the system device name to the logical volume on the FAStT storage subsystem. In Example 5-4 we run the **lsattr** command on the LUN named hdisk1. It provides the following information: It is a 36 GB LUN of type RAID 5, with a LUN ID of 0, and an IEEE volume name of 600A0B80000CD96D000000063EF70D0C. You can make a quick identification by locating the LUN ID on the far right side of the Mappings View tab.

Example 5-4 Volume information

```
# lsattr -El hdisk1
cache_method  fast_write                       Write Caching method                 True
ieee_volname  600A0B80000CD96D000000063EF70D0C IEEE Unique volume name              False
lun_id        0x0000000000000000               Logical Unit Number                  False
prefetch_mult 1                                Multiple of blocks to prefetch on read True
pvid          none                             Physical volume identifier           False
q_type        simple                           Queuing Type                         False
queue_depth   10                               Queue Depth                          True
raid_level    5                                RAID Level                           False
reassign_to   120                              Reassign Timeout value               True
reserve_lock  yes                              RESERVE device on open               True
rw_timeout    30                               Read/Write Timeout value             True
scsi_id       0x11100                          SCSI ID                              False
size          36864                            Size in Mbytes                       False
write_cache   yes                              Write Caching enabled                True
```

You can make a more precise correlation which LUN this is using the distinctive ieee_volname attribute. The value of this attribute on the AIX host is the same as the Unique Logical Drive Identifier on the FAStT storage subsystem. The Unique Logical Drive Identifier can be found in a Storage Manager window by right-clicking **Logical Drive Name -> Properties**. Look for the Volume ID, Capacity, and RAID level properties.

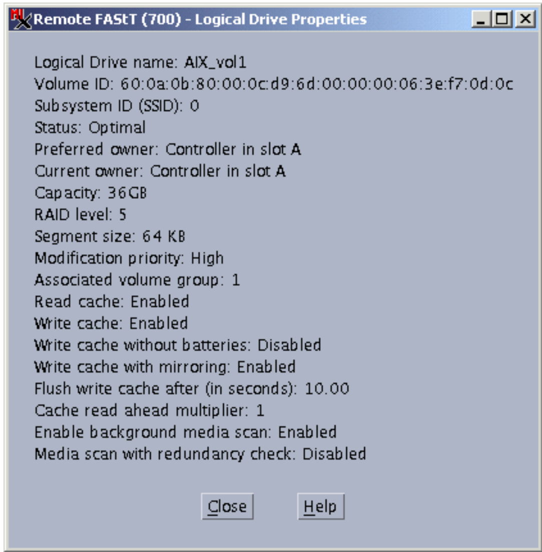

Figure 5-4 Logical Drive Properties

Changing ODM attribute settings in AIX

You can actually change the ODM attributes for the RDAC driver and FAStT, and here we describe the settings that can be used for best performance and availability on hdisk devices.

This section lists the attribute settings that you should use for hdisk devices and shows how to set them using the **`chdev -l`** command. To make the attribute changes permanent in the Customized Devices object class, use the -P option. Some attributes can be changed from both the SMclient and through AIX. To make these changes permanent, the modifications must be made by changing the AIX ODM attribute.

The specific hdisk attributes affected are (see also Example 5-4 on page 122):

- write_cache

  Indicator that shows whether write-caching is enabled on this device (yes) or not (no).

- cache_method

  If write_cache is enabled, the write-caching method of this array is set to one of the following:

  – default - Default mode; the word *default* is not seen if write_cache is set to yes.

  – fast_write - Fast-write (battery-backed, mirrored write-cache) mode.

  – fw_unavail - Fast-write mode was specified but could not be enabled; write-caching is not in use.

  – fast_load - Fast-load (non-battery backed, non-mirrored write-cache) mode.

  – fl_unavail - Fast-load mode was specified but could not be enabled.

- prefetch_mult

  Number of blocks to be prefetched into read cache for each block read (value: 0–100).

- queue_depth

  Number that specifies the depth of the queue based on system configuration. Reduce this number if the array is returning a BUSY status on a consistent basis (value: 1–64)

Attention: Controller cache mirroring should not be disabled while cache_write is enabled. If this condition exists, the RDAC software will automatically re-enable it the next time the system is restarted, or when **`cfgmgr`** is run.

If the changes are made through the SMclient, they will operate properly until you either restart the host or restart **`cfgmgr`**. To make the changes permanent, you must use SMIT or the **`chdev -P`** command.

Setting the queue depth for hdisk devices

Setting the queue_depth attribute to the appropriate value is important for system performance. For large FAStT configurations with many volumes and hosts attached, this is a critical setting for high availability.

Use the following formula to determine the maximum queue depth for your system:

*512 / (number-of-hosts * LUNs-per-host)*

For example, a system with four hosts, each with 32 LUNs (the maximum number of LUNs per AIX host), would have a maximum queue depth of 4:

*512 / (4 * 32) = 4*

In this case, you would set the queue_depth attribute for hdiskX as follows:

```
#chdev -l hdiskX -a queue_depth=4 -P
```

Attention: If you do not set the queue depth to the proper level, you might experience a loss of your file system.

AIX restrictions

When connecting the FAStT to AIX system there are some restrictions/guidelines that you must take into account. This does not mean that other configurations will not work, but you might end up with unstable or unpredictable configurations that are hard to manage and troubleshoot.

The following restrictions apply to FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 storage servers:

- F-RAID Manager is not supported.
- The maximum number of partitions per AIX host, per FAStT storage server, is two.
- Each AIX host can support two or four host bus adapters and up to two FAStT storage partitions, each requiring two HBA connections per FAStT storage server. Additional HBA pairs can be added to support additional FAStT storage servers.
- Zoning must be implemented. If zoning is not implemented in a proper way devices might appear on the hosts incorrectly (too many, not enough, odd errors, etc.). Follow this rule when implementing the zoning: Multiple FC adapters in the same host must be in separate zones. Each Fibre Channel adapter must be zoned with one particular controller port on the FAStT.
- Single HBA configurations are allowed only in configuration with a SAN switch, but each single HBA configuration requires that both controllers in the FAStT be connected to a switch within the same SAN zone as the HBA. See Figure 5-5 on page 125.
- Direct-attach configurations are restricted to single-initiator configurations only. Only one connection to each FAStT mini-hub is allowed. That means we can connect directly to the

FAStT 500/700/900 two AIX servers only. More servers can be connected through the switch. See Figure 5-3 on page 121.

FAStT200 and FAStT600 storage servers do not have mini-hubs so we can connect one AIX system directly to the FAStT200 and two AIX systems directly to the FAStT600.

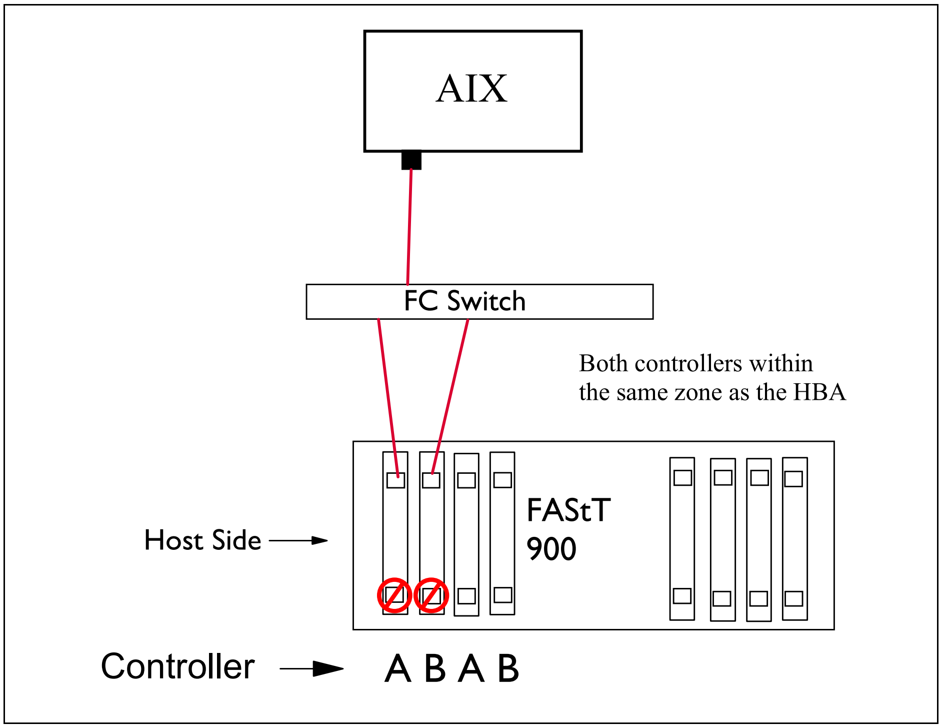

Figure 5-5 Single HBA configuration in AIX environment

- Single-switch configurations are allowed, but each HBA and FAStT controller combination must be in a separate SAN zone.
- All volumes that are configured for AIX must be mapped to an AIX host group. Connecting and configuring to volumes in the default host group is not allowed.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.
- Booting up from a FAStT device is supported only on AIX 5.1 and 5.2. When you start from a FAStT device, both paths to the boot device must be up and operational. Path failover is not supported during the AIX boot process. Once the AIX host has started, the failover operates normally.
- Dynamic Volume Expansion (DVE) is only supported in AIX 5.2. At AIX 5.2, a new option was added to the **chvg** command. After a successful DVE of a volume on the FAStT Strorage Server, the user has to run **chvg -g <vgname>** in order for the volume group to recognize the additional physical partitions that were added by increasing the size of the volume. The -g option is not available on AIX Version 4.3.3 and 5.1.

  For more info and the exact procedure on how to expand a volume in AIX refer to *"Installation and Support Guide for AIX, HP-UX, and Solaris"*.

  In any case, before using DVE for operating systems that have an LVM installed or integrated (like AIX) you should first evaluate if it is worth using. Using DVE facilitates the administration of your FAStT (because you end up with less volumes), but bigger volumes will not give you a gain in performance. Do not forget that when using LVM you can expand your file systems by adding new volumes to the volume groups or disk groups, and this gives you more communication paths to your LUNs.

5.1.5 Install FAStT software on a Linux xSeries host server

This section describes the installation of device drivers and host software (based on Storage Manager 8.4) on xSeries server running Linux Red Hat Enterprise Server AS 2.1.

We successively describe:

- Installing FAStT Management Suite Java (MSJ)
- Installing the HBA device driver
- Installing the FAStT host software

Note that there is no Storage Manager agent for Linux, hence no in-band management available for Linux hosts.

Linux (Red Hat Enterprise Server AS 2.1)

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

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 drivers 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.

You can find the Linux Red Hat updates from the following URLs:

http://www.redhat.com

http://updates.redhat.com

Be sure to log onto the server 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

Note: you can find more detailed information on FAStT MSJ, including installation, in "FAStT MSJ" on page 128.

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. 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 Bus Adapter driver

The device driver enables the operating system to communicate with the host bus adapter.

Before installing drivers, make sure that your system meets the requirements stated in the readme file of the driver package regarding kernel versions, host adapter settings, and so on. You may have to install the kernel headers and kernel source for the supported version of kernel before you install the drivers. These can also be downloaded from:

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

1. Download and install the drivers as described in the readme that is contained with the driver package.
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=128
   ```

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 ramdisk to include the module at boot time. Because the exact names vary from distribution to distribution, we use generic names here:

   ```
   /sbin/mkinitrd /boot/initrd-image 2.4.9-3smp
   ```

   The name 2.4.9-3smp represents the exact kernel version that the system is running. This is the subdirectory name within the /lib/modules directory from which the kernel modules were taken. This may differ on the version of Linux that you are running.
5. Reboot the system and the driver loads automatically with the new settings.

Installing the FAStT host software

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

1. FAStT Runtime Environment (SMruntime)
2. FAStT Client (SMclient)
3. FAStT Utilities (Smutil)

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

Installing SMruntime

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 and enter the following command:

```
rpm -ivh SMruntime<version>.rpm
```

Installing SMclient

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

```
rpm -ivh SMclient<version>.rpm
```

The client requires an X-Windows environment. To start the Storage Manager client, type the following command:

```
SMclient
```

Under Linux, the Event Monitor is always installed. The Event Monitor daemon starts automatically after the next reboot. If you do not 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.

Installing SMutil

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

```
rpm -ivh SMutils<version>.rpm
```

5.2 FAStT MSJ

FAStT MSJ is a Java-based utility that can be loaded to preform some management tasks and diagnostics on Linux, Microsoft Windows, and Novell Netware host systems. In Linux, it is also used with the QLRemote agent to define preferred and alternate paths for logical drives.

FAStT MSJ consists of two components:

- An agent that is installed on all servers that have one or more host bus adapters
- A graphical interface 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 3
- LINUX advanced server 2.1
- Novell NetWare 6.0 SP 2

The Graphical interface is supported on the following operating systems:

- Windows NT with service pack 6 a
- Windows 2000 with service pack 2
- LINUX advanced server 2.1

FAStT MSJ supports the following adapters on a Windows or NetWare system:

- IBM FAStT Fibre Channel Host Bus Adapters (2200)
- IBM FAStT FC-2 Host Bus Adapter (2310)

- IBM FAStT 133 FC-2 Host Bus Adapter (2340)
- IBM FAStT 133 FC-2 Dual Port Host Bus Adapter(2342)

FAStT MSJ supports the following adapters on a Linux System:

- IBM FAStT Fibre Channel Host Bus Adapters (2200)
- IBM FAStT FC-2 Host Bus Adapter (2300)
- IBM FAStT FC-2 133 Host Bus Adapter (2340)
- IBM FAStT FC-2 133 Dual Port Host Bus Adapter(2342)

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 publication was written, is Version 2.0, Release 40, for Windows NT4, Windows 2000, NetWare, and Linux. It is available from the following Web site:

http://www-3.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-39194

This section will describe installing FAStT MSJ for both Windows and the Linux environment. In the Linux environment you need to configure FAStT MSJ for multi-path fail-over.

5.2.1 Installing FAStT MSJ for Windows

This section explains how to install FAStT MSJ. FAStT MSJ supports the following two configurations:

- Locally

  If you want to manage the Host Bus Adapters in a stand-alone system, install both the GUI and agent on the local system.

- Remotely

  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 licensing conditions and click **Next** again.
4. The Choose Product Features window opens as shown in Figure 5-6 on page 130. Select the options you want to install on this host and click **Next**.

Figure 5-6 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.

5.2.2 Installing FAStT MSJ for Linux

Under Linux, 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.

Installing and using FAStT MSJ

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 5-7 on page 131.

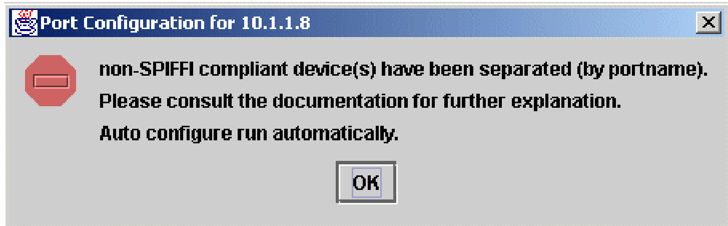

Figure 5-7 FAStT MSJ error message for the access logical drive

If you receive this error, you will need to remove the mapping to LUN31 for this storage partition.

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 its client component by entering the following command in a command window (you must be logged on as 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.
4. Highlight one of the adapters under the HBA tab, and click **Configure**, as shown in Figure 5-8.

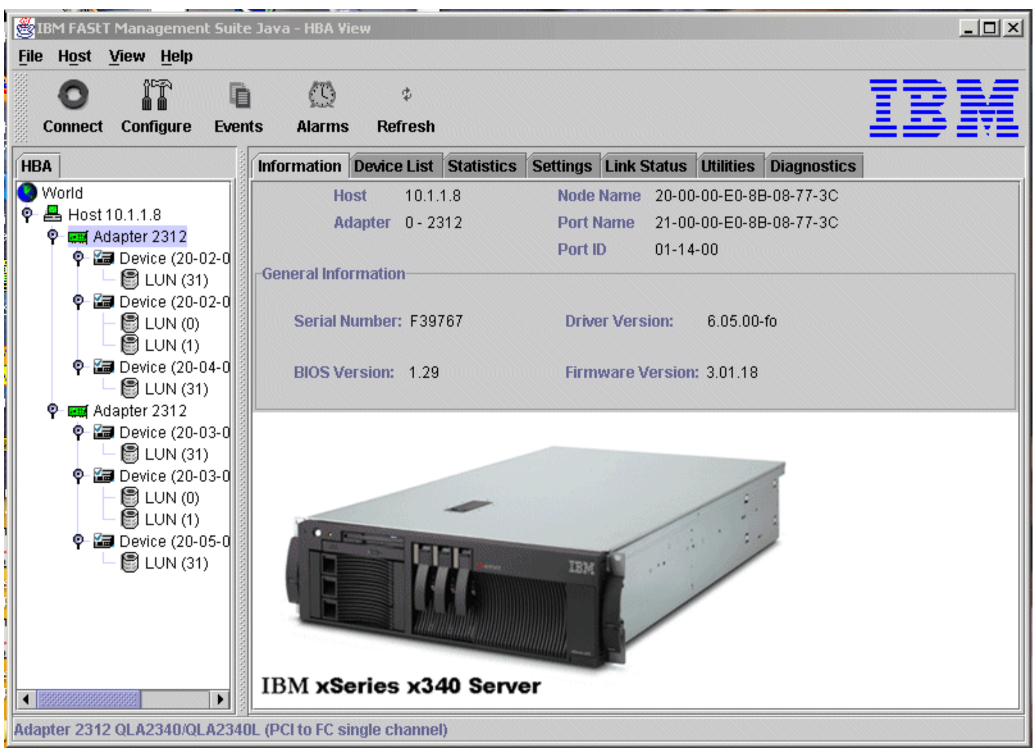

Figure 5-8 FAStT MSJ connected to a Linux host

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 5-9 on page 132) stating an invalid configuration.

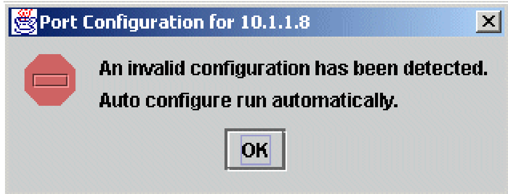

Figure 5-9 FAStT MSJ invalid configuration detected

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 5-10). This leads to a configuration where all LUNs are enabled and the first adapter is used as the preferred path.

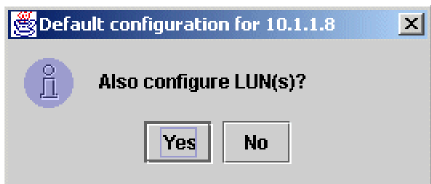

Figure 5-10 FAStT MSJ automatic LUN configuration

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 FAStTs 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 5-11). Configure both adapters: One as visible and the other as hidden.

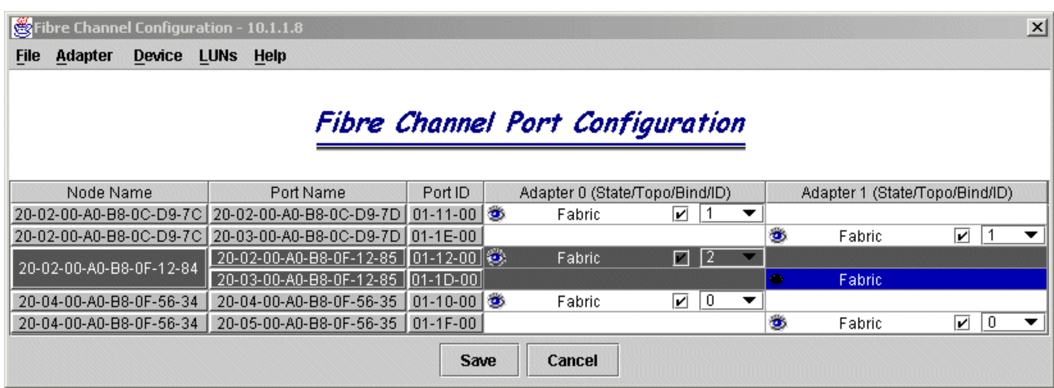

Figure 5-11 FAStT MSJ Fibre Channel port configuration

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.

Even if one adapter is hidden to the operating system, it does not 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 5-10 on page 132), 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 Logical drive not on preferred path, 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 5-12

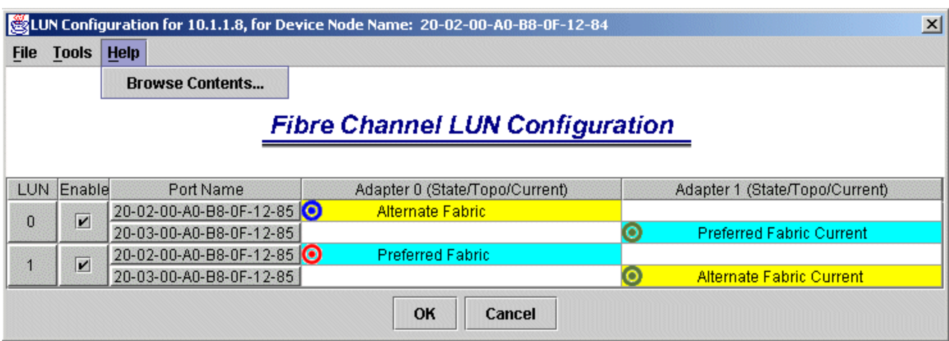

Figure 5-12 FAStT MSJ LUN configuration

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 by 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 5-13 on page 134).

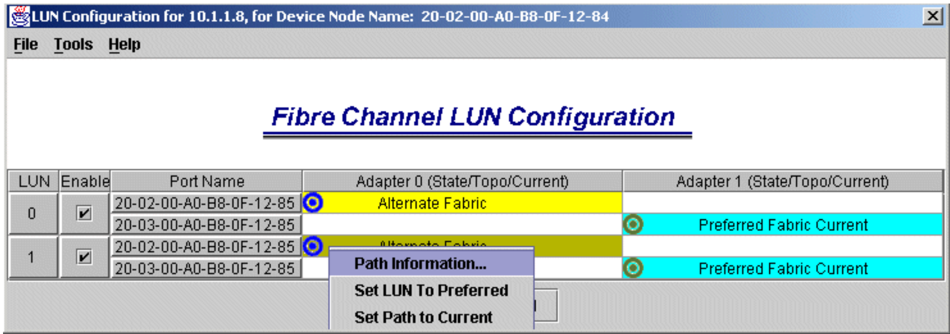

Figure 5-13 FAStT MSJ load balancing

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 5-11 on page 132). A confirmation appears if the configuration saved successfully.

   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 an 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.

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

For additional information on how to run the IBM FAStT MSJ Diagnostics and utilities refer to Chapter 6, “Troubleshooting and diagnostics” on page 169.

5.3 Preparing the FAStT Storage Server

This section explains how to configure the FAStT Storage Server. This includes:

- Configuring TCP/IP addresses on the FAStT Storage Server
- Initializing the Storage Manager client
- Updating the controller microcode
- Updating the drive and ESM microcode

5.3.1 Network setup of the controllers

Tip: With Version 8.4 of the Storage Manager client—and assuming you have the appropriate firmware level for the controllers—it is also possible to set the network settings using the client graphical front-end. Please refer to "User control of network parameters" on page 15.

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

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 have destructive effects (causing loss of data or even affecting the functionality of your FAStT).

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 use 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.
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 onto 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 5-5.

Example 5-5 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 5-6.

Example 5-6 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.

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.

5.3.2 Starting the FAStT Storage Manager client

We know already that the FAStT Storage Manager client can be used for either in-band or out-of-band management of the Storage Server. In-band management uses the Fibre Channel to communicate with the IBM FAStT Storage Server, while out-of-band management uses the TCP/IP network to communicate with the IBM FAStT Storage Server. On host platforms that support both methods, it is possible to use them on the same machine if you have a TCP/IP connection and also a Fibre Channel connection to the FAStT Storage Server.

Note: When you install FAStT SMclient on a stand-alone host and manage the storage server 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.

In our illustration we assume out-of-band management from a machine that only has a TCP/IP connection to the FAStT Storage Server.

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 5-14).

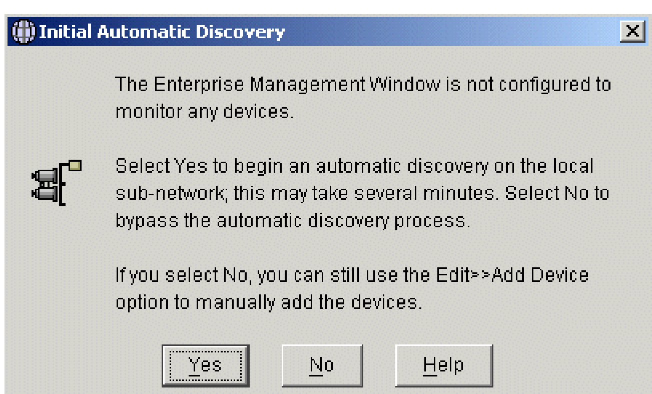

Figure 5-14 Initial Automatic Discovery

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 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 **Tools -> Rescan**. Afterward, all FAStT Storage Servers are listed in the Enterprise Management window, as shown in Figure 5-15.

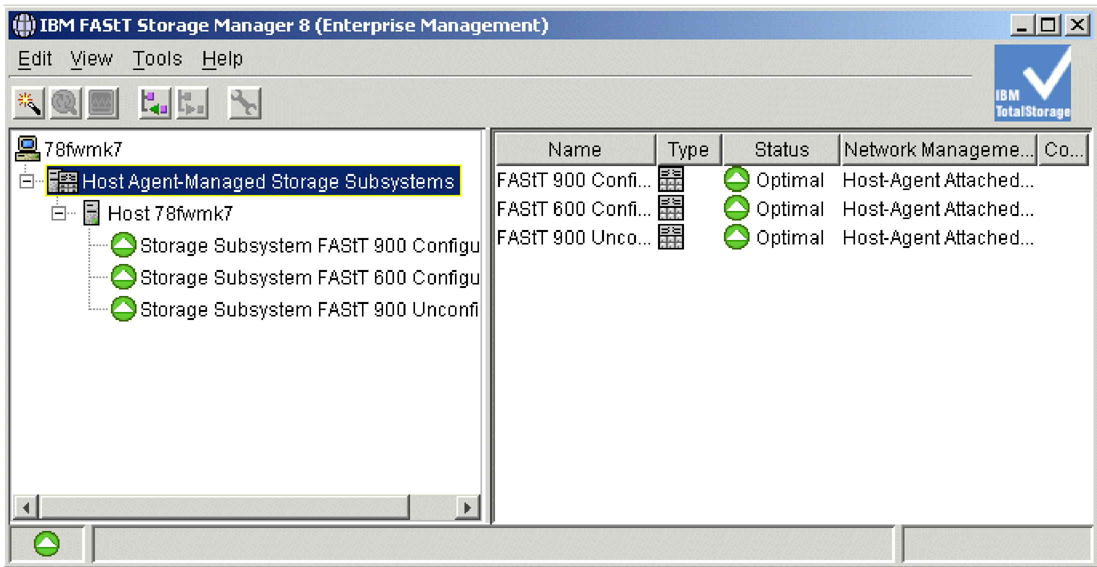

Figure 5-15   The Enterprise Management window (partial view)

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 5-16). If you add a FAStT Storage Server that is directly managed, be sure to enter both IP addresses—one per controller.

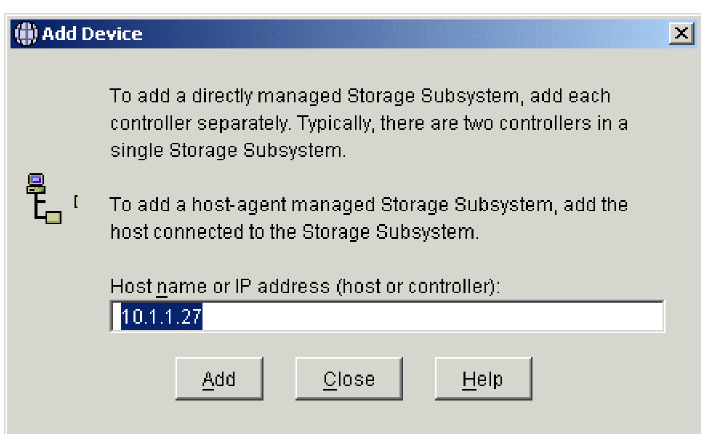

Figure 5-16   Adding a device manually

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.

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 5-17).

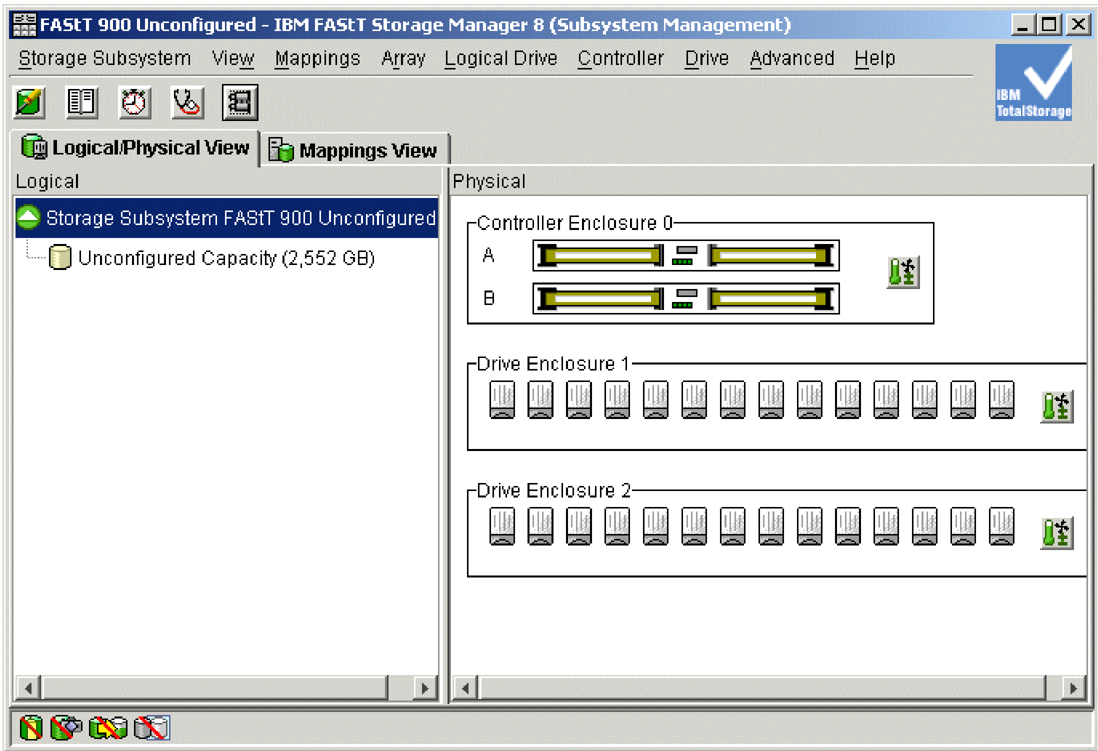

Figure 5-17 First launch of the Subsystem Management window

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 5-18).

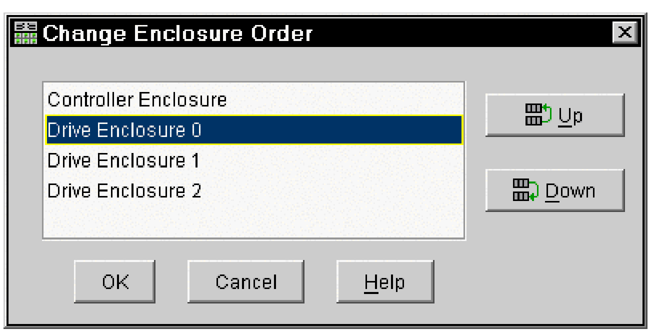

Figure 5-18 Adapting the enclosure order

5.3.3 Updating the controller microcode

Updating the firmware (controller microcode) of your FAStT Storage server may be required when installing a new version of the Storage Manager software; it is also necessary to upgrade from time to time to keep up with the latest fixes and enhancements. For full instructions on how this is done you should refer to the documentation that is supplied with the firmware or refer to:

http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/disk

Then click your FAStT model and click the appropriate version of the device driver matrix for Storage Manager.

Or use the following links for SM8.3 and SM8.4 respectively:

http://www-1.ibm.com/support/docview.wss?rs=506&uid=psg1MIGR-50177
http://www-1.ibm.com/support/docview.wss?rs=506&uid=psg1MIGR-52950

We have also included step-by-step instructions for this process in “Updating controller microcode” on page 335.

5.4 Configuring the IBM FAStT Storage Server

This section describes the sequence of steps to follow for the configuration of a FAStT Storage Server:

1. Configuration planning
2. Naming your Storage Server, synchronizing the clock (Storage Manager 8.4 has a warning if these are not synchronized), and setting the password
3. Defining a hot-spare
4. Creating arrays and logical drives
5. Configuring storage partitioning
6. Monitoring and alert options

5.4.1 Planning your 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 need to apply in advance.

On one hand, 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 applications using the FAStT Storage Server.

On the other hand, you also need to define the layout of the attached hosts with their host bus adapters and the mappings of the logical drives to specific host groups or hosts. You should prepare a mapping table, similar to the one presented in Table 5-1 on page 141, where you keep all necessary information regarding the storage partitioning.

Table 5-1 Sample plan for storage partitioning

Host group	Host name	Port name	WWN	OS type
Windows 2000	Windows_Host	MailAdp_A	200000E08B28773C	Windows 2000 Non-Clustered
		MailAdp_B	200000E08B08773C	
Linux	Linux_Host	LinAdp_A	200100E08B27986D	Linux
		LinAdp_B	200000E08B07986D	
RS6000	AIX_Host	AIXAdp_A	20000000C926B6D2	AIX
		AIXAdp_B	20000000C926B08	

The name of a given 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.

5.4.2 Initial configuration

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.

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 5-19).

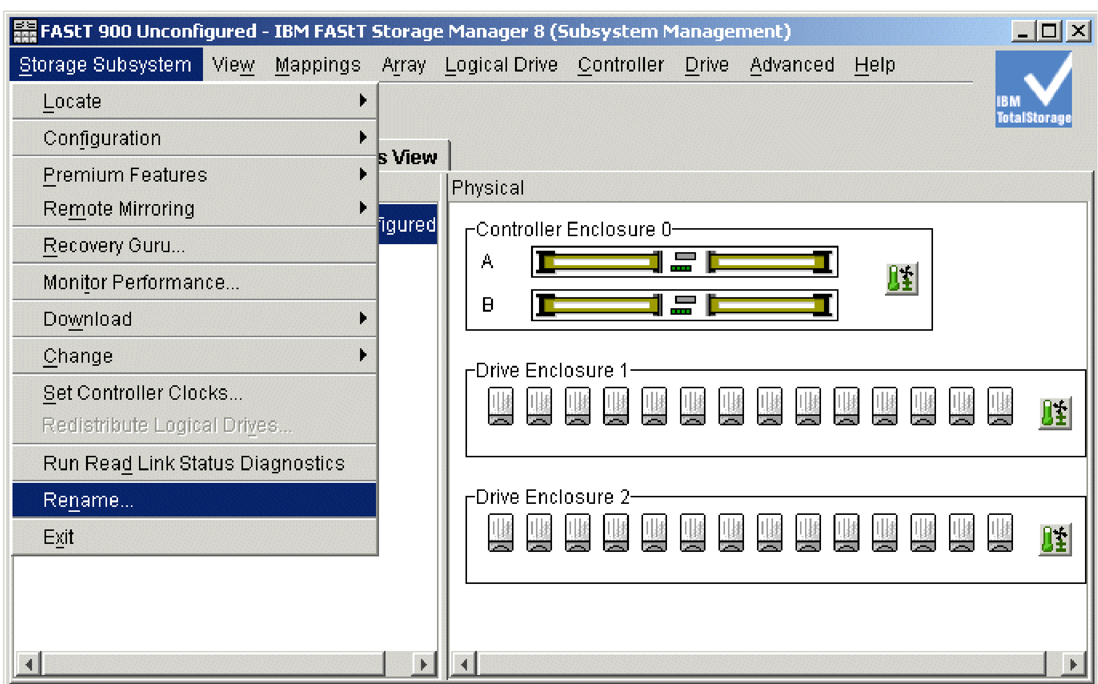

Figure 5-19 Renaming the FAStT Storage Server

2. Enter a new name for the subsystem in the Rename Storage Subsystem dialog (Figure 5-20 on page 142).

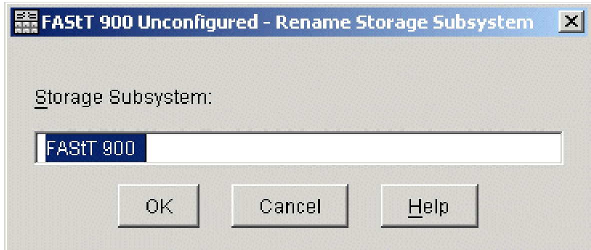

Figure 5-20 Renaming FAStT

3. Since the FAStT Storage Server stores its own event log, synchronize the controller clocks with the time of the host system. If you have not already set the clocks on the Storage Servers, set them now. Be sure that your local system is working using the correct time. Then click **Storage Subsystem -> Set Controller Clock** (Figure 5-21).

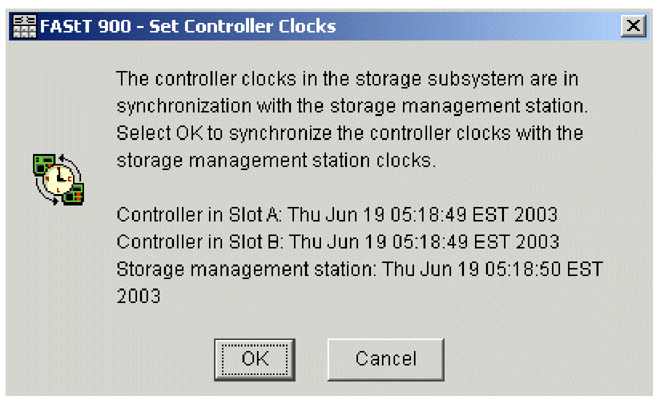

Figure 5-21 Setting the controller clock

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 5-22 on page 143). This password is then stored on the FAStT Storage Server. It is used if you connect through another SMclient or the FAStT Field Tool. It does not matter whether you are using in-band or out-of-band management.

Figure 5-22 Setting a password on the FAStT Storage Server

5.4.3 Defining hot spare drives

The concept of hot spare drive is explained in “Hot spare drive” on page 81.

To define a hot spare drive, highlight the drive you want to use. Click **Drive -> Hot Spare -> Assign** as shown in Figure 5-23.

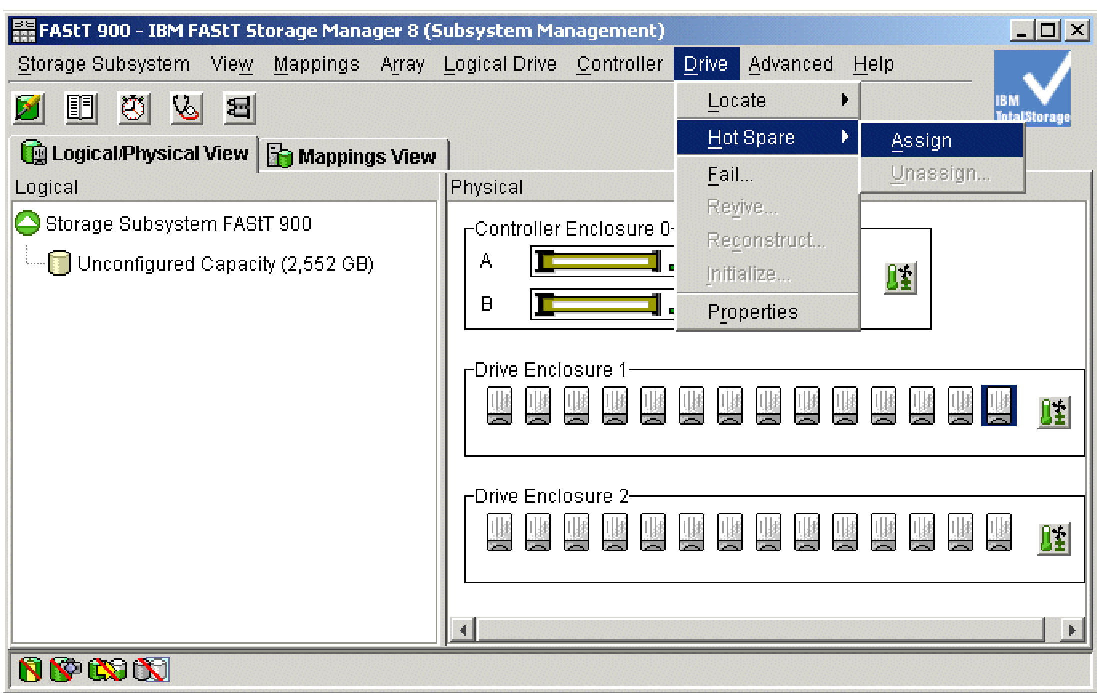

Figure 5-23 Defining a hot spare drive

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**.

5.4.4 Creating arrays and logical drives

At this stage of the process, the storage subsystem is installed, upgraded to the newest microcode level, and at least one hot spare drive is defined. Arrays and logical drives can now be configured. Refer to “FAStT management concepts and basics” on page 75 for guidance on how to divide the available storage capacity into arrays or logical drives, and which restrictions apply, to avoid improper or inefficient configurations.

You can define logical drives from unconfigured-capacity or free-capacity nodes on the storage subsystem.

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:

- 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.

The illustration that follows assume a blank FAStT (unconfigured capacity) with the hot spare drive that we have just defined in the previous section. The same steps are necessary to define additional logical drives on an existing configuration (free capacity).

1. In the Subsystem Management window (Figure 5-24 on page 145), right-click the unconfigured capacity and select **Create Logical Drive**.
   a. If this is the first time you are creating a logical drive from this instance of the Storage Manager client, the window shown in Figure 5-25 on page 145 appears.

      The current default host type might not reflect your desired choice, but once specified it will stick to the last type used from here on. You should only change this type if you attach hosts running the same operating system.

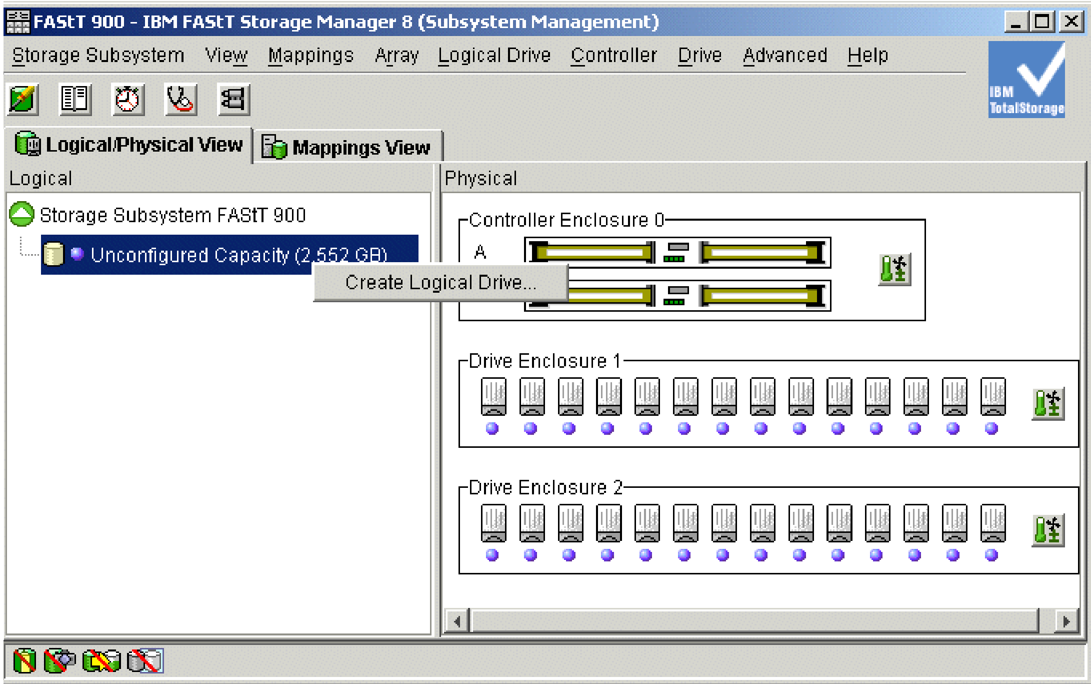

Figure 5-24 Creating a logical drive

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

c. Click **OK**.

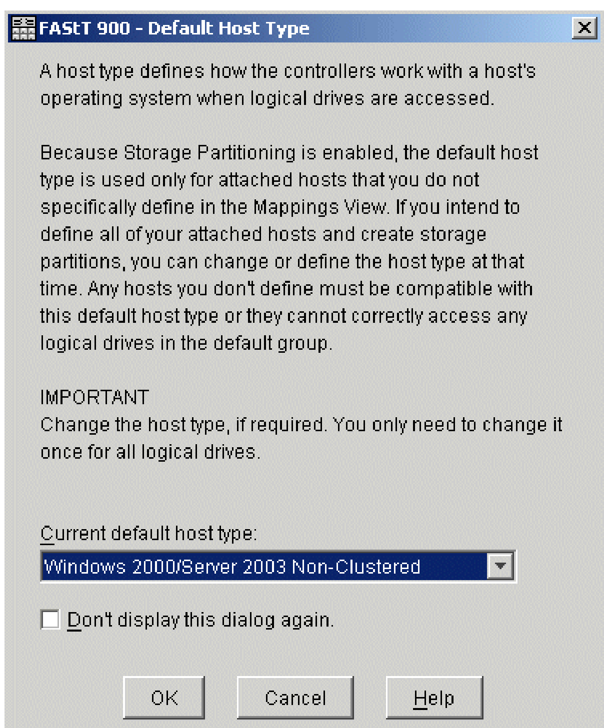

Figure 5-25 Default Host Type window

2. The Create Logical Drive Wizard window displays (see Figure 5-26).

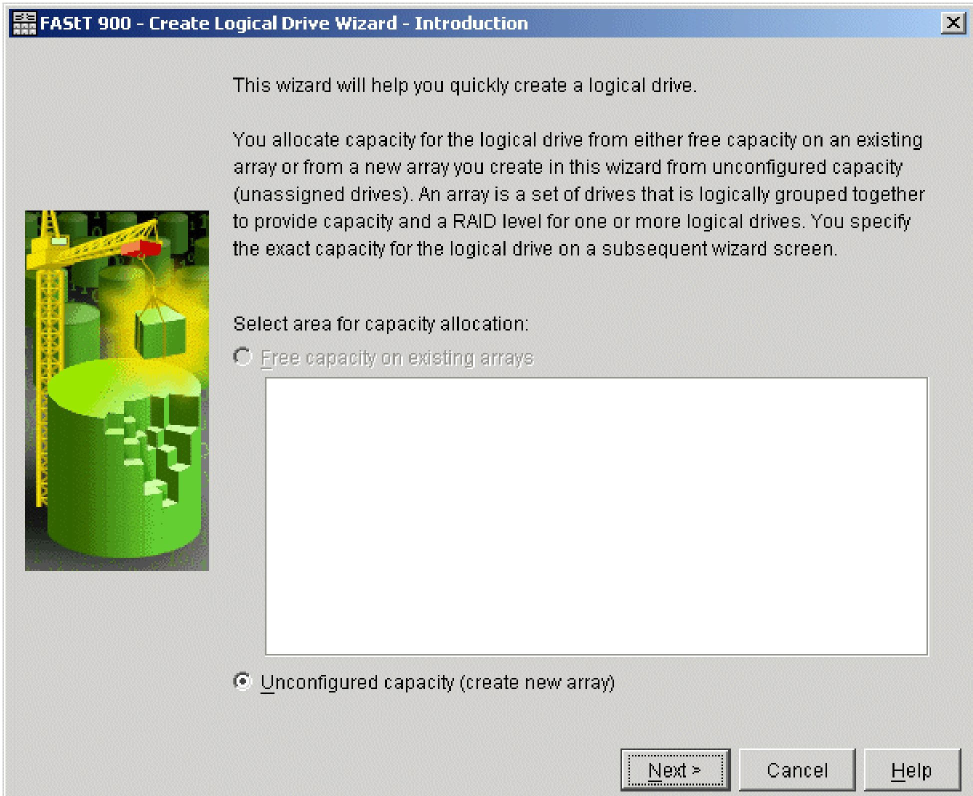

Figure 5-26 Create Logical Drive Wizard

The wizard leads you through the creation of your logical drives. Since you do not have any configured arrays yet, you need to select the **Unconfigured capacity** option to create a new array. Then click **Next**.

3. The dialog shown in Figure 5-27 on page 147 displays; specify the RAID level and number of drives to use for the array. 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 understand performance and availability impacts 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. In this case, you only have to specify 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 optimize performance and availability.

   You need to define the desired RAID level first, as discussed in “RAID levels” on page 78, before you choose the number of drives you want to use. Click **Next**.

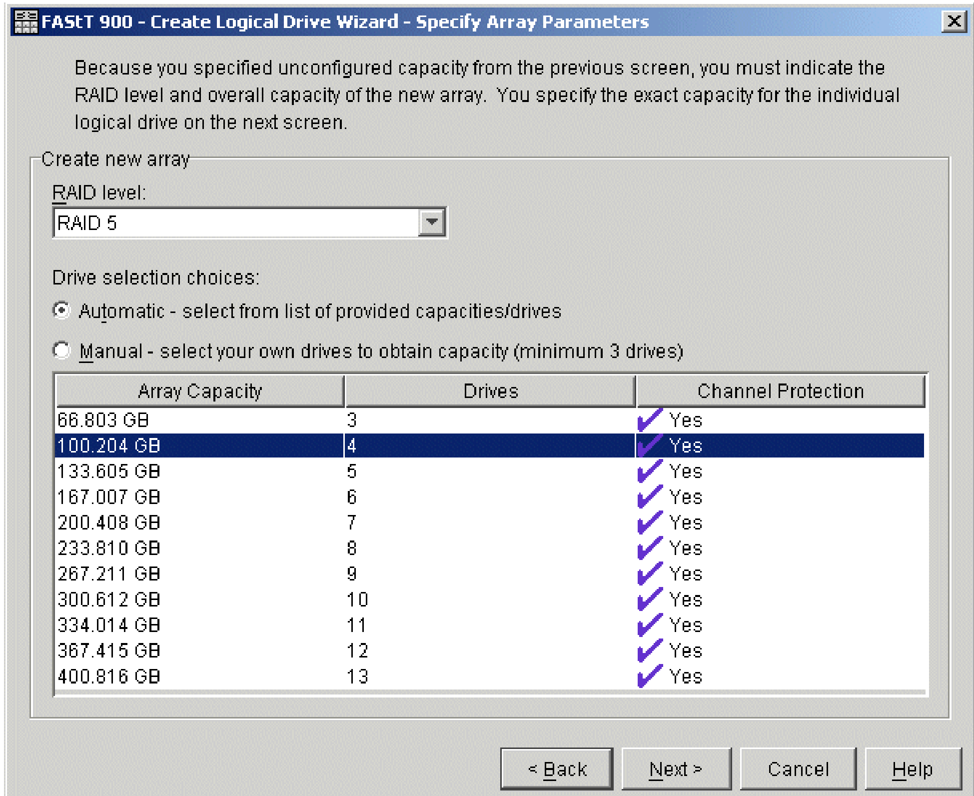

Figure 5-27 Choosing the drives used for the new array

4. Define the logical drive (Figure 5-28 on page 148). 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 array, enter the desired size.

   b. Assign a name to the logical drive, which complies to your mapping table (Table 5-1 on page 141).

   c. If you want to change advanced logical drive settings as segment size or cache settings, select the **Customize settings** option and proceed to step 5.

   Otherwise, click **Finish** to complete the creation of the logical drive and proceed to step 6.

   The newly created logical drive is not mapped automatically and 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.

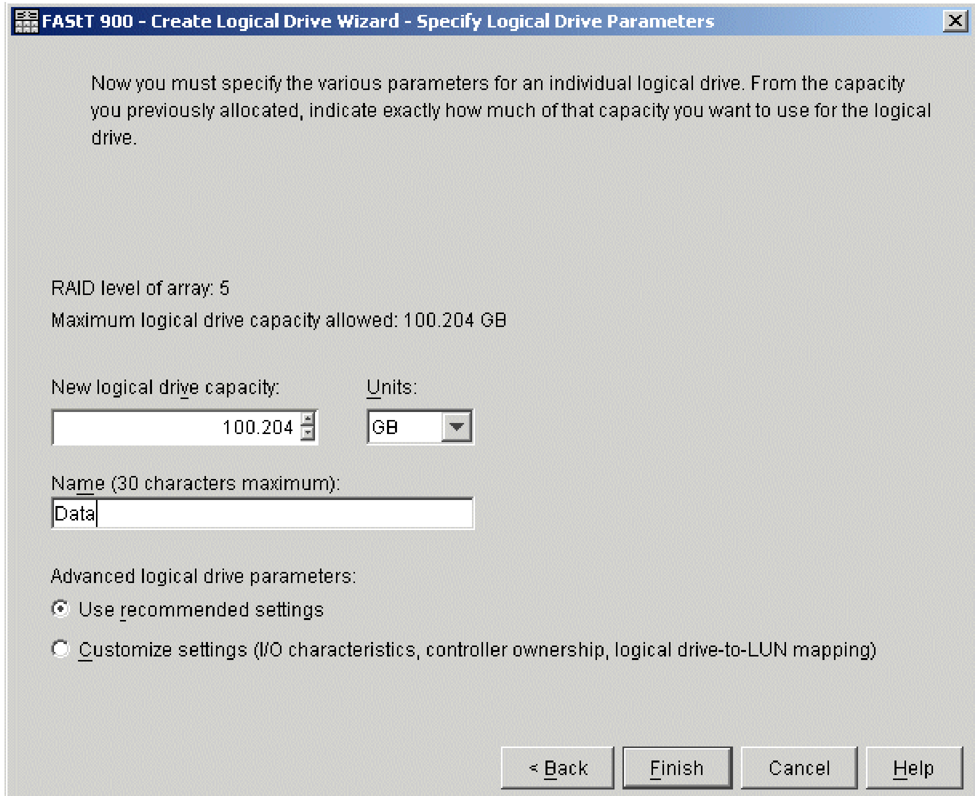

Figure 5-28 Logical drive parameters

5. On the Specify Advanced Logical Drive Parameters panel (see Figure 5-29 on page 149) you can define the logical drive to exactly fit 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**. Table 5-2 shows the defaults of a logical drive when it is created.

Table 5-2 Logical Drive defaults

	File system	**Database**	**Multimedia**
Segment size	64K	64K	128K
Modification priority	High	High	High
Read cache	Enable	Enable	Enable
Write cache	Enable	Enable	Enable
Write cache with out batteries	Disable	Disable	Disable
Write cache with mirroring	Enable	Enable	Enable
Flush write cache after	10 seconds	10 seconds	10 seconds
Cache read ahead multiplier	1	0	8
Enable background media scan	Enable	Enable	Enable

	File system	Database	Multimedia
Media scan with redundancy check	Disable	Disable	Disable

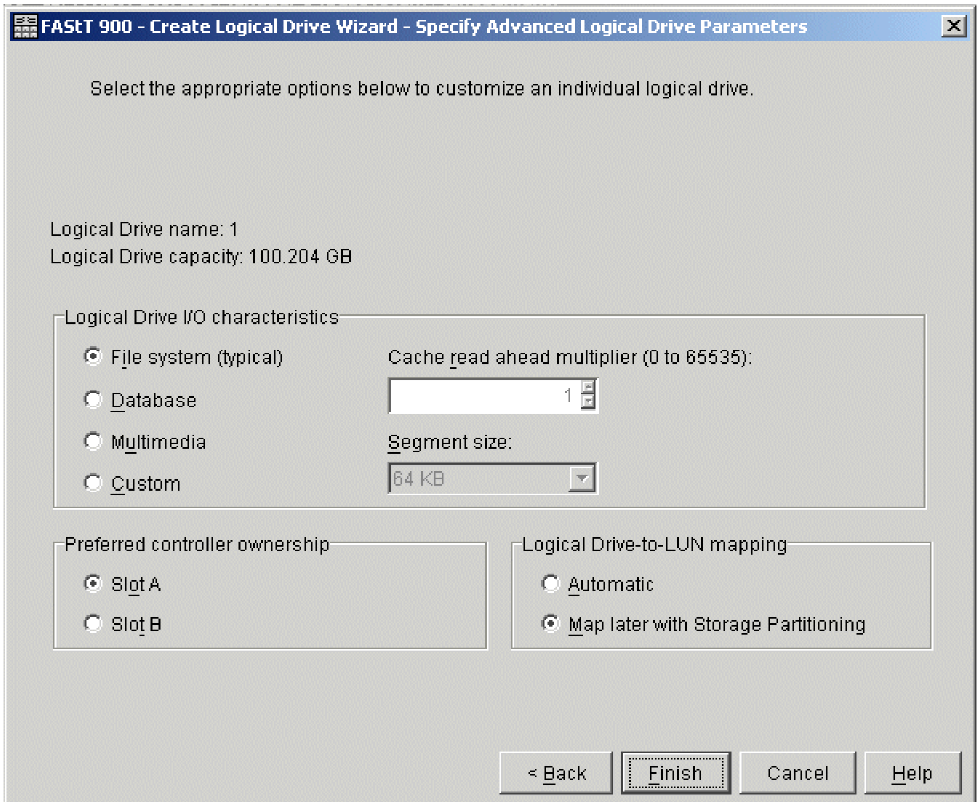

Figure 5-29 Advanced logical drive settings

b. The segment size is chosen according to the usage pattern. For custom settings, you can directly define the segment size. With SM8.4 a segment size of 512KB is possible and must be set through the command line interface.

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 (distribute) your logical drives between 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 “Performance monitor data” on page 85 and “Controller ownership” on page 93).

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 “Configuring storage partitioning” on page 150. 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. You will then be asked if you want to continue creating logical drives or not. Click **No** if you have finished creating logical drives or wish to finish at a later time (see Figure 5-30).

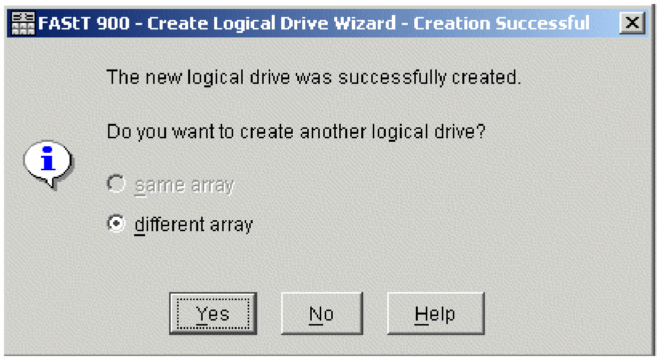

Figure 5-30 Continue creating logical drives

6. Click **Finish** and you are presented with the following screen.

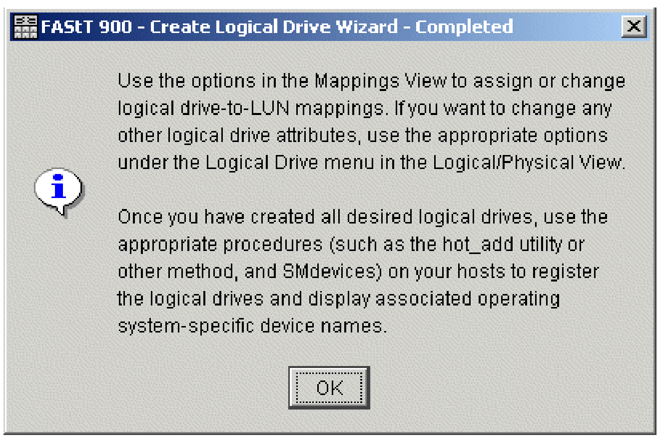

Figure 5-31 Logical Drive creation completed

7. 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 initialized and immediately accessible.

If you left unconfigured capacity inside the array, you can later define another logical drive in this array. Simply highlight this capacity, right-click, and choose **Create Logical Drive**. Follow the steps that we outlined in this section, except for the selection of drives and RAID level (since the array is already defined).

5.4.5 Configuring storage partitioning

We have explained the concept of storage partitioning in “Storage partitioning” on page 82.

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. With Storage Manager 8.4, you can now have up to 256 LUNs assigned to a single storage partition. You may have 2048 LUNs per storage server.

To configure storage partitioning, follow these steps:

1. Choose the **Mappings View** in the Subsystem Management window. You are presented with the following information box when you first enter storage partitioning (Figure 5-32).

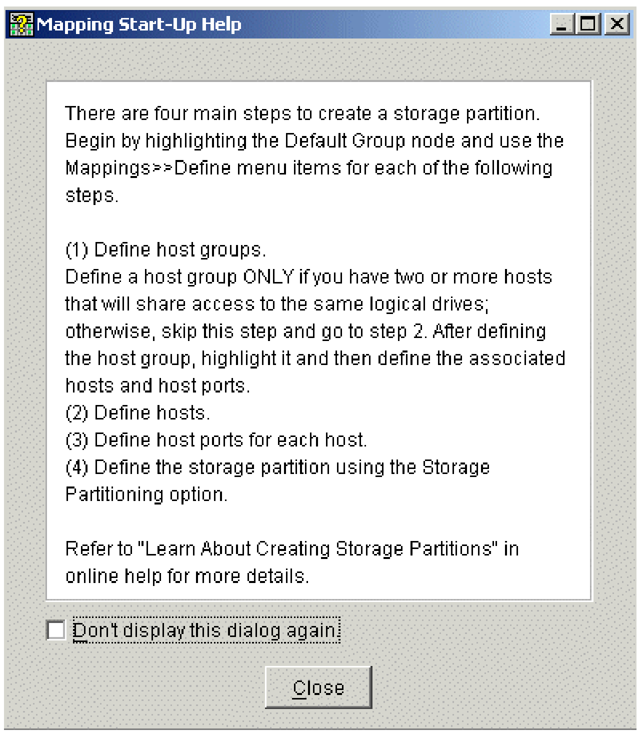

Figure 5-32 Mapping Start Up Help

2. 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.
3. Define the host groups. Highlight the **Default Group**, right-click, and choose **Define Host Group** as shown in Figure 5-33 on page 152.

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

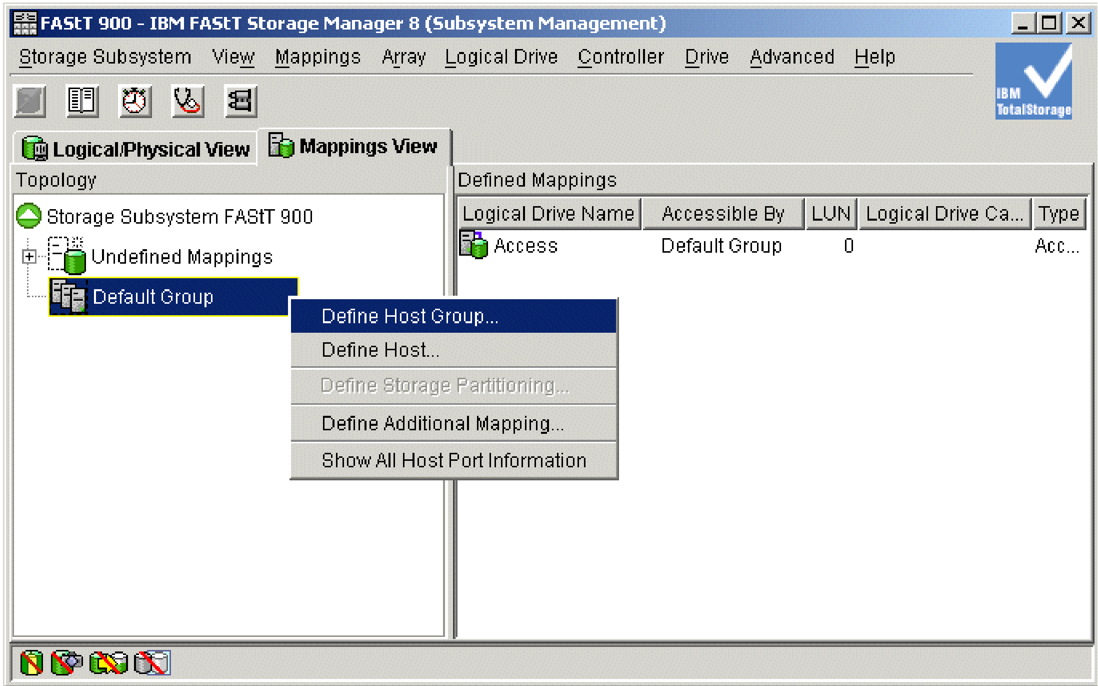

Figure 5-33 Define Host Group

4. The Define Host Group dialog (Figure 5-34) 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 anyway. Otherwise the addition of new systems is not possible without disrupting the already defined mappings in the default host group.

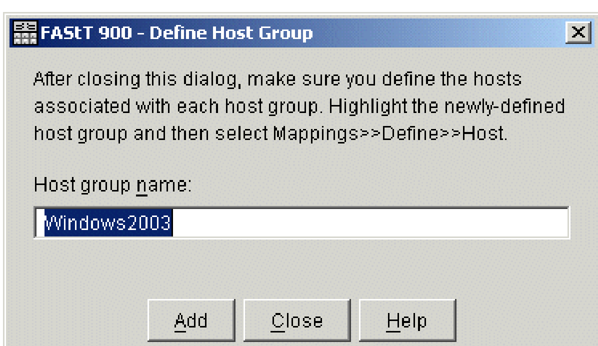

Figure 5-34 Entering the host group name

5. The hosts groups are defined and the hosts in these groups can now be defined according to the example in Table 5-1 on page 141. As shown in Figure 5-35 on page 153, highlight the group for which you want to add a new host. Right-click and choose **Define Host**.

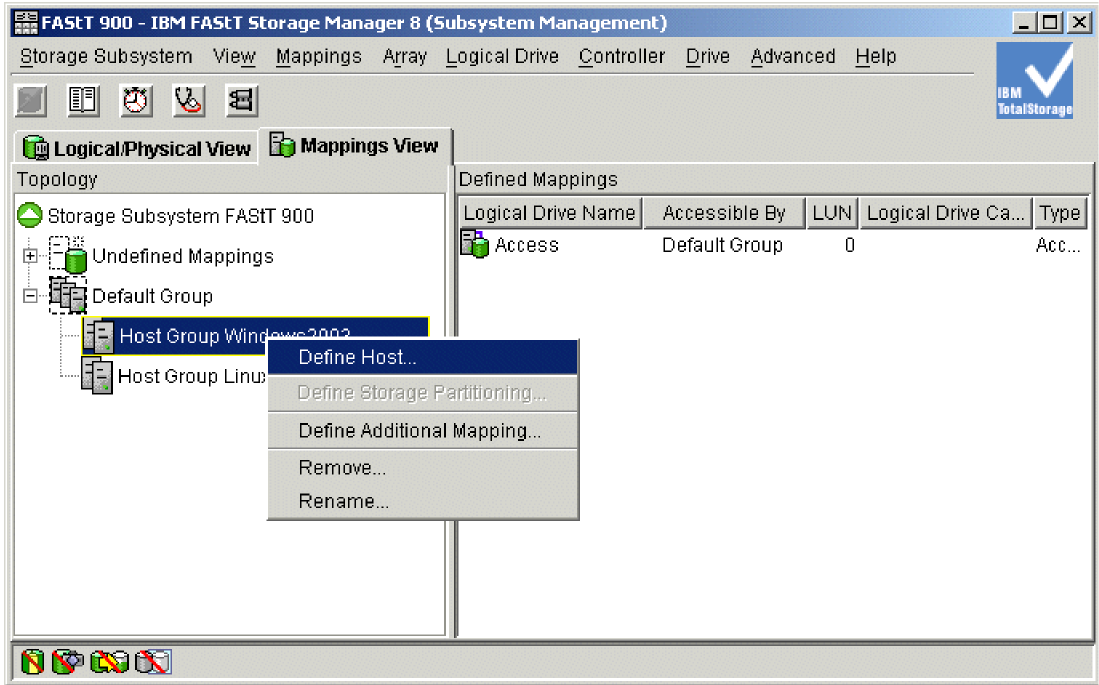

Figure 5-35 Selecting Define Host

6. On the Define Host dialog (Figure 5-36), 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.

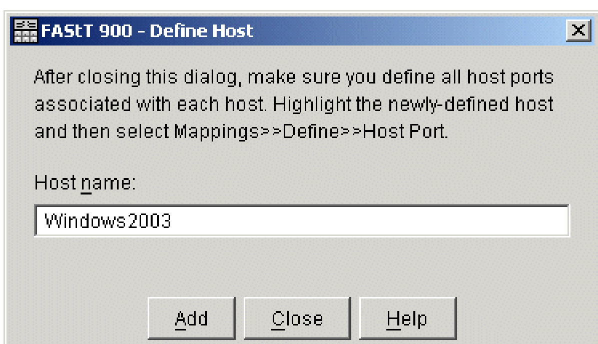

Figure 5-36 Define Host

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.

7. 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 do not, 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 5-37.

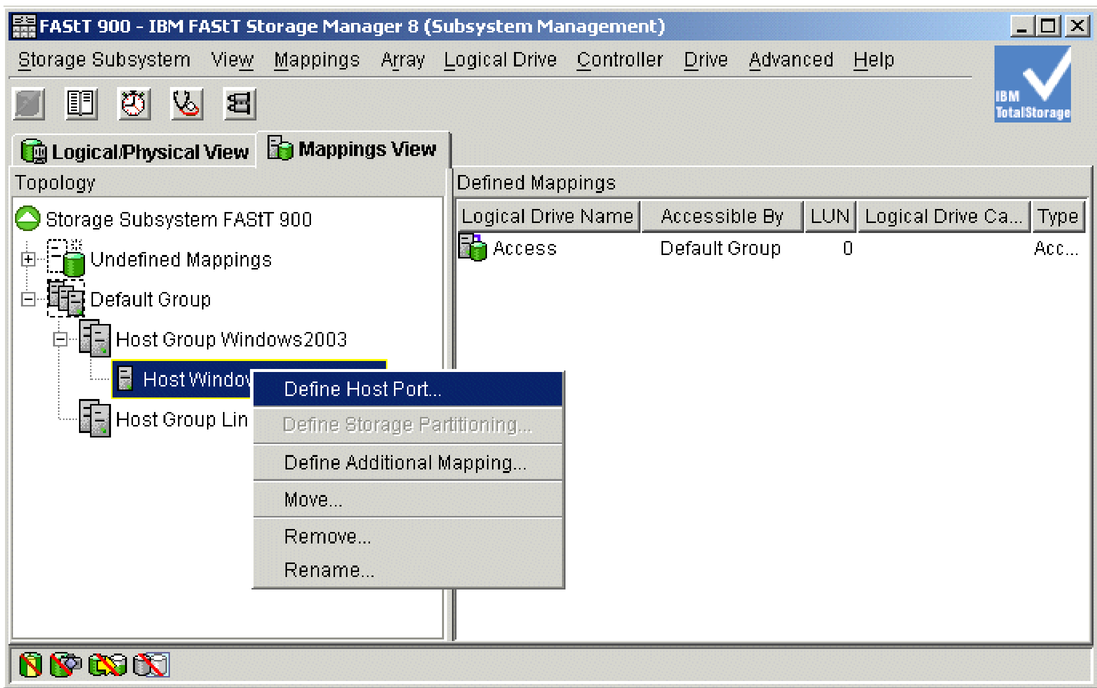

Figure 5-37 Choosing Define Host Port

8. In the Define Host Port dialog (Figure 5-38 on page 155), 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 5-1 on page 141.

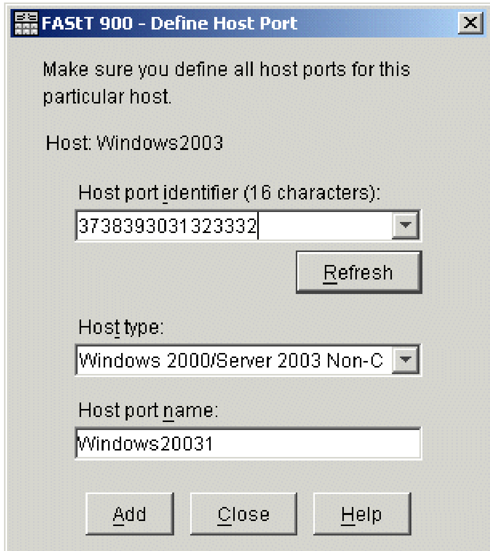

Figure 5-38 Enter the host port information

9. Define the mapping for each logical drive created in “Creating arrays and logical drives” on page 144. 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 5-39, highlight the host group to which you want to map a new logical drive. Right-click and choose **Define Additional Mapping**.

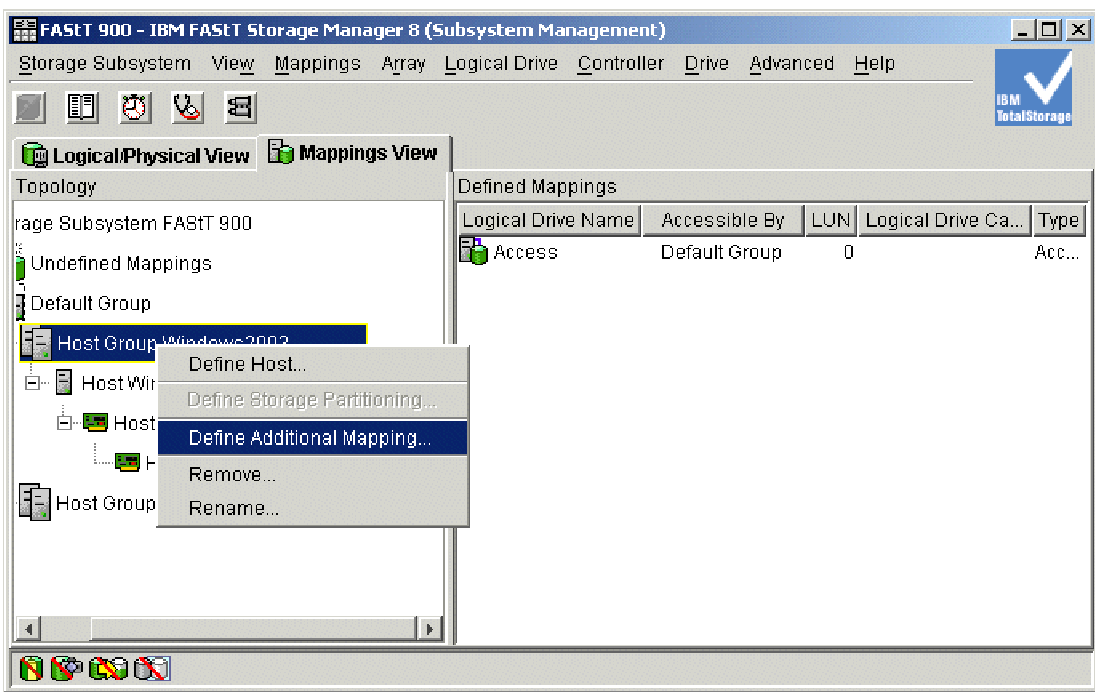

Figure 5-39 Define Additional Mapping

10. 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 5-40), according to the mapping table (Table 5-1 on page 141).

   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.

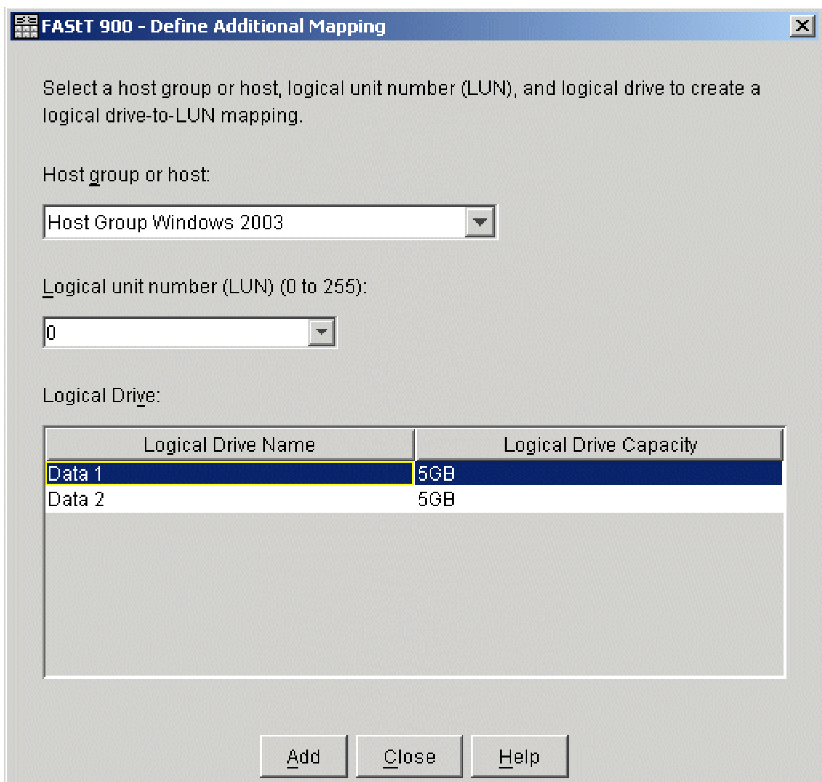

Figure 5-40 Define Additional Mapping

Note: If you create a new mapping or change an existing mapping of a logical drive, the change is effective 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.

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 requires that a new configuration be 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 mounting them.

If you attached a Linux or AIX system to the FAStT Storage Server, you need to delete the mapping of the access LUN. 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.

So far we updated the storage subsystem to the latest level, configured logical drives, and set up storage partitioning.

The next step is to define the alerting methods in case of failure. The final step is to document and save the configuration as it is now.

5.4.6 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 you of 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 5-41). 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.

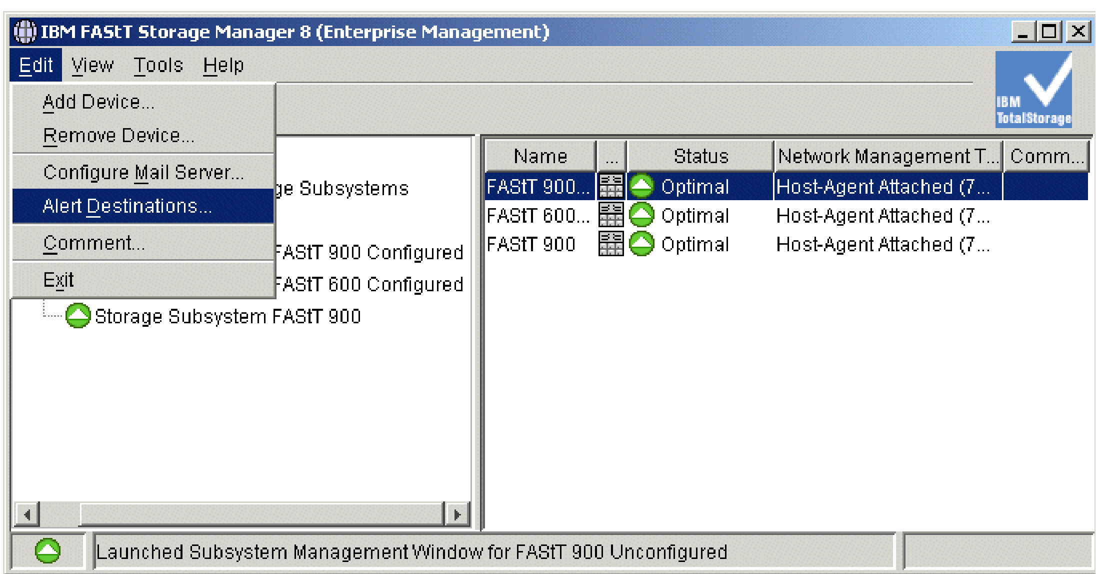

Figure 5-41 Defining Alert Destinations

An icon in the lower left corner of the Enterprise Management window indicates that the Event Monitor is running on this host.

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 5-42.

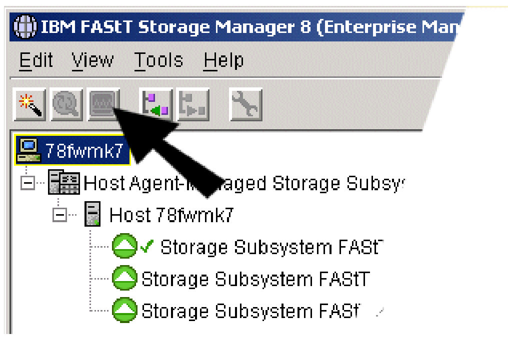

Figure 5-42 Unsynchronized Event Monitor icon

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, 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.

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.

IBM FAStT service alert

FAStT service alert is a feature of the IBM TotalStorage FAStT Storage Manager that monitors system health and automatically notifies the IBM Support Center when problems occur. Service alert sends an e-mail to a call management center that identifies your system and captures any error information that can identify the problem. The IBM Support Center analyzes the contents of the e-mail alert and contacts you with the appropriate service action.

See Appendix B, “IBM FAStT Service Alert” on page 391, for more information on how to set up Service Alert.

5.5 Saving the subsystem profile

Configuring a FAStT Storage Server is a complex task and it is therefore essential to save the configuration. The information can be saved in a single location (file) known as the subsystem profile. The profile stores 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 5-43.

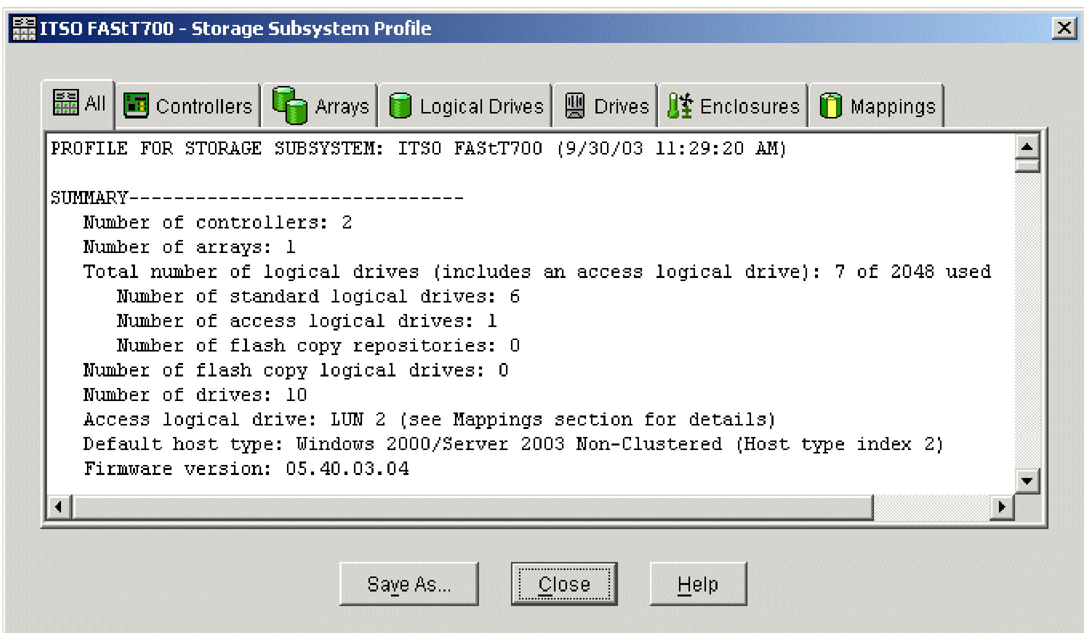

Figure 5-43 Storage Subsystem Profile

5.6 Advanced configuration

This section discusses how to access and configure the advanced management topics presented in “Advanced functions” on page 90.

5.6.1 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, for one to two weeks, 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 5-44) shows you the global cache settings. The values for start and stop flushing are discussed in “Cache parameters” on page 94, as well as the cache block size value. The changes occur immediately.

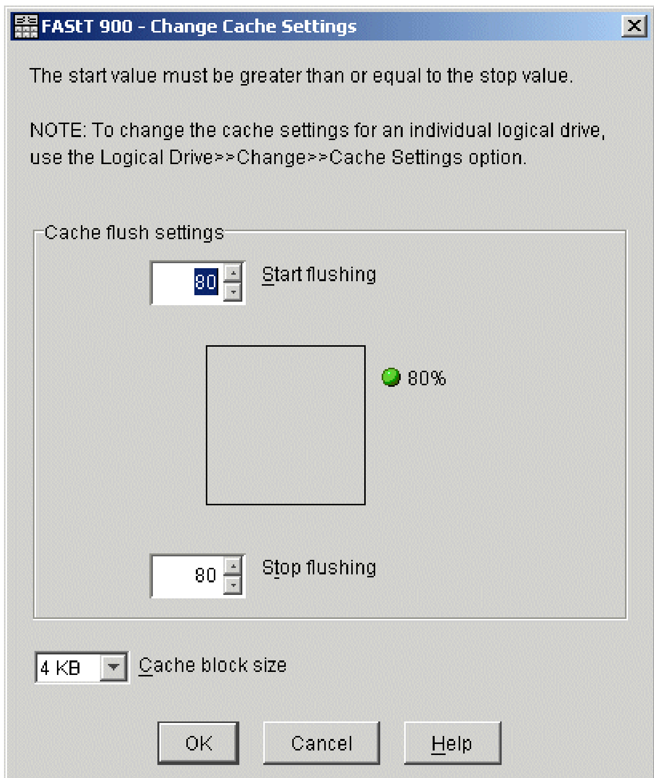

Figure 5-44 Global cache settings

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.

To change any settings, right-click a logical drive and choose **Change -> Cache Settings**. The Change Cache Settings dialog (Figure 5-45) allows you to change the parameters that we described.

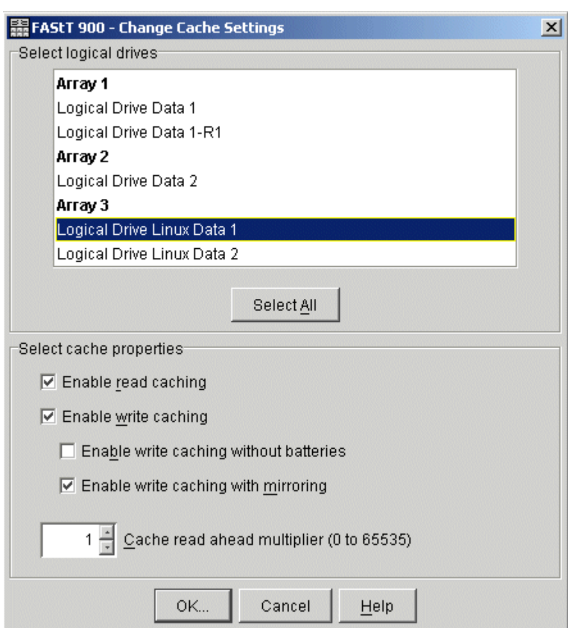

Figure 5-45 Logical drive cache settings

5.6.2 Modification priority

The modification priority defines how much processing time is used for operations modifying the logical drive relative to the system performance.

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 5-46).

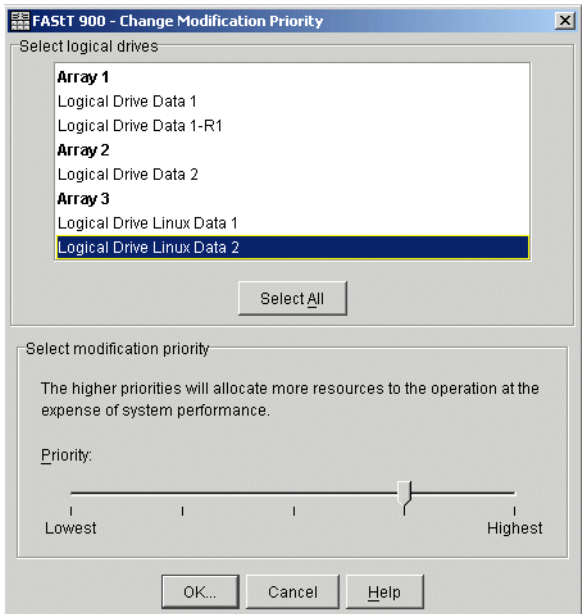

Figure 5-46 Modification priority for a logical drive

If a logical drive modification is in progress, a status bar appears at the bottom of the dialog.

5.6.3 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 defective host bus adapter or switch. Normally the logical drive's ownership is alternated between the two controllers while you define them.

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 awhile to finish.

5.6.4 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 "RAID levels" on page 78.

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 163.

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 select **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 select **Add free Capacity (Drives)**. On the Add Drives window (Figure 5-47), you need at least one unassigned drive that can be added to the array.

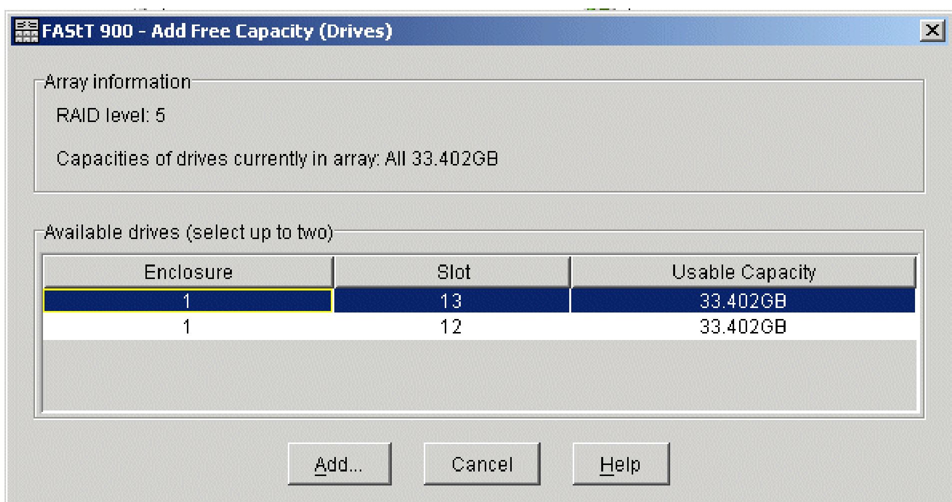

Figure 5-47 Adding new drives to an array

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. The array needs to be defragmented first to consolidate all free space nodes to one free space node for the array. Then all 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.

5.6.5 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 "Segment size" on page 80. 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

Note that Storage Manager 8.4 allows you to define a segment size of 512KB; however, this can be done from the command line interface only.

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 starts, 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

The dynamic logical drive expansion feature was introduced with Storage Manager Version 8. 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 “Arrays and logical drives” on page 75.

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 does not 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 5-48), 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 163, but you can only add two drives at a time here, because you are combining the two operations.

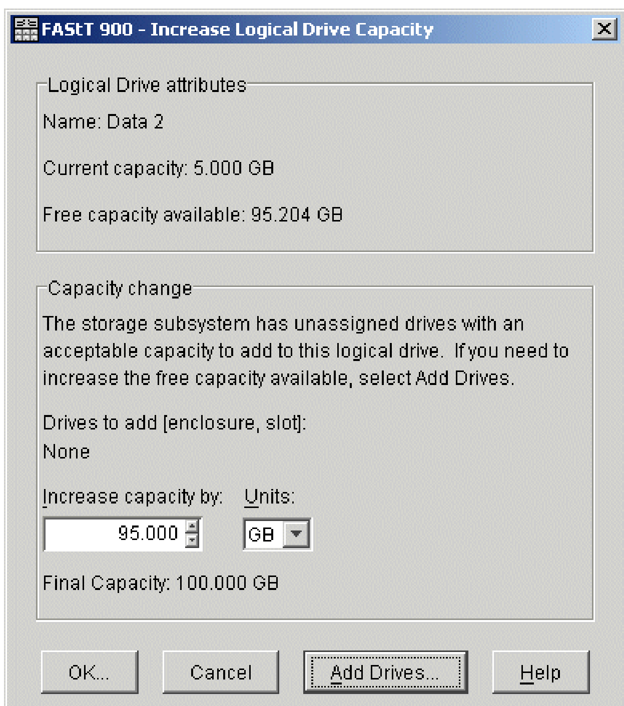

Figure 5-48 Dynamic logical drive expansion

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.

5.6.6 Media scanning

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 5-49), you can also define the duration of one complete scan of all selected logical drives.

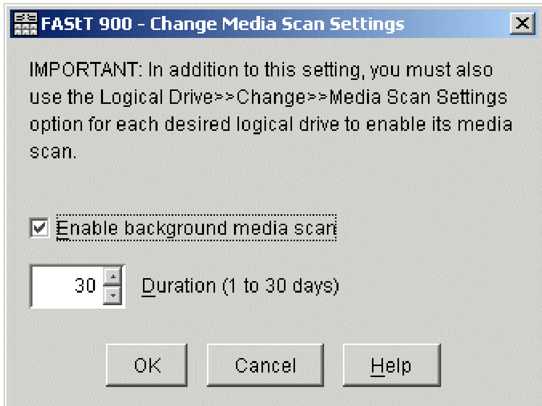

Figure 5-49 Enabling media scan

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 5-50), select the logical drives to include in the media scan.

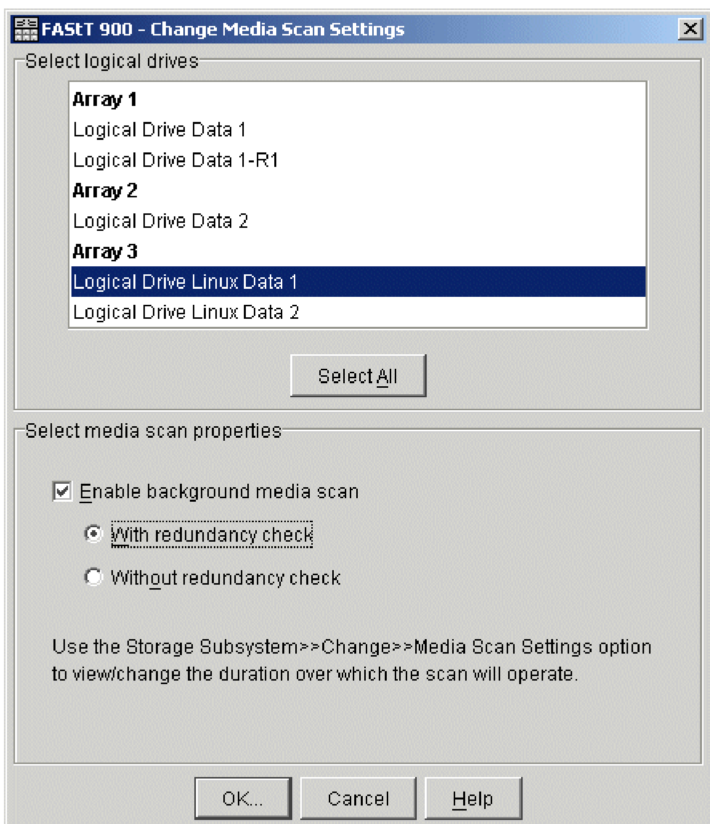

Figure 5-50 Enabling a logical drive media scan

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.

5.6.7 Component properties

All the information received in the profile (see “Starting the FAStT Storage Manager client” on page 137) is gathered from the various components of the storage subsystems as controllers, disk drives, and enclosures. The information collected is also available for each component. This may be useful to track or locate a failure in a component.

To see all of the properties of an enclosure, click the components icon (indicated by the arrow in Figure 5-51).

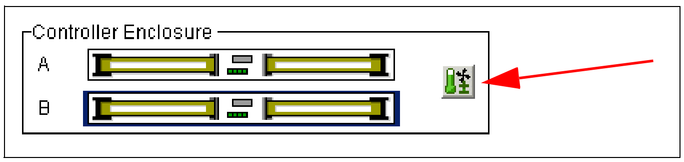

Figure 5-51 Components icon

You can also highlight the component you want to examine, right-click, and choose **Properties**. The properties appear for this component, for example, for one of the controllers (see Figure 5-52).

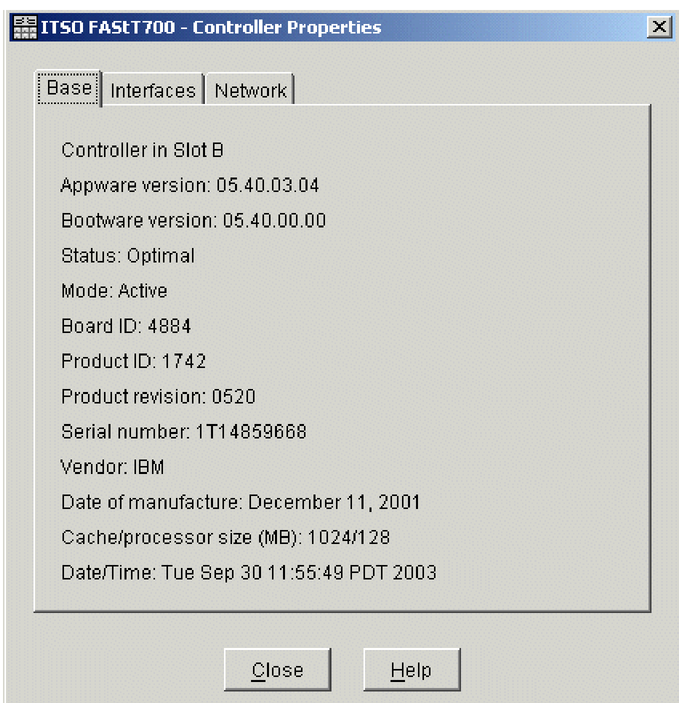

Figure 5-52 Properties of a controller

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 5-53 on page 168) 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 5-53 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 5-54.

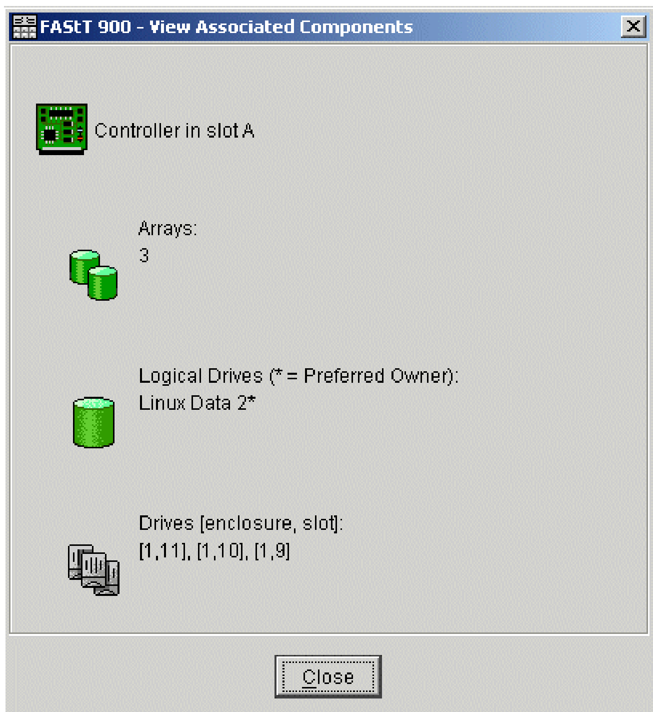

Figure 5-54 Associated components of Controller A

6

Troubleshooting and diagnostics

In this chapter we discuss IBM FAStT MSJ and Storage Manger 8 in conjunction with troubleshooting and diagnostics on FAStT600/700/900. We explain, based on examples, how to check the Major Event Log (MEL) and subsequent actions to take for various operating systems.

6.1 Diagnostics tools

This section discusses some of the diagnostic tools available for the FAStT.

6.1.1 FAStT MSJ

We have introduced FAStT MSJ in “FAStT MSJ” on page 128. The FAStT Management Suite Java GUI can be used to manage the host bus adapters in servers running the FAStT MSJ agent. This section gives an overview on how to use the tool to make changes to the configuration of host bus adapters, perform diagnostic tests, and check the performance of the host bus adapters.

Setting the FAStT MSJ client

When you launch the FAStT MSJ GUI client, the IBM FAStT Management Suite Java - HBA View window, shown in Figure 6-1, opens.

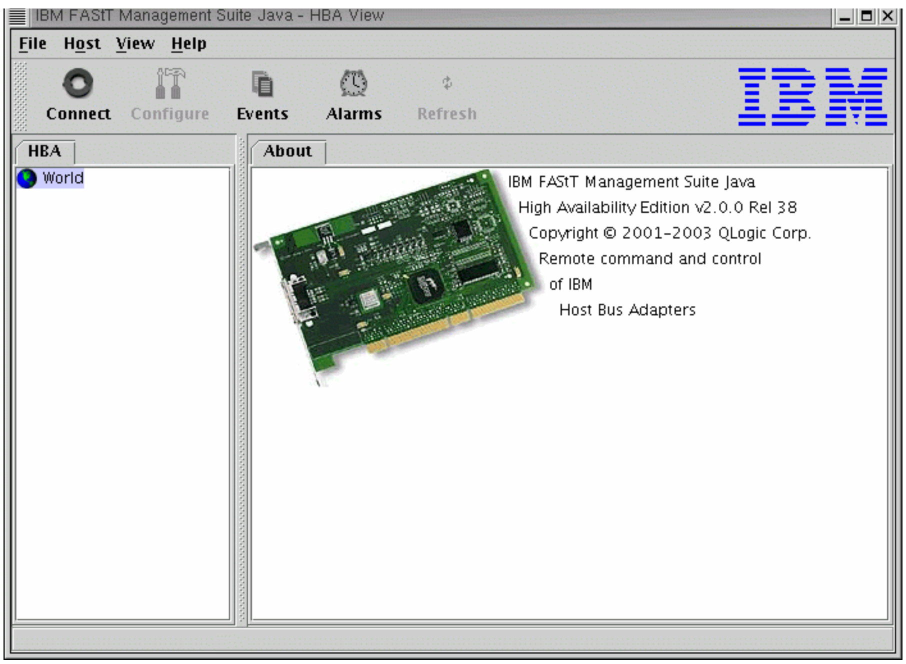

Figure 6-1   FAStT HBA view

To connect to a specific host, click the **Connect** icon; the Connect to Host window opens as shown in Figure 6-2 on page 171.

Enter the name or the IP address of the server or select from the pull-down menu to which server you want to connect, and click the **Connect** push-button.

At that point you are returned to the HBA View window and the host you specified in the previous dialog is now displayed in the left-hand side pane (known as the HBA tree panel). The host bus adapters installed in the server appear below the host name (Figure 6-3 on page 171).

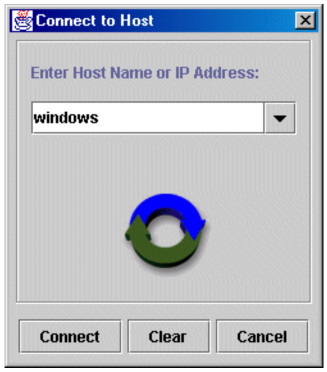

Figure 6-2 FAStT MSJ connect to host

When you click the host name, two tabs display in the right-hand pane (known as the Tab panel; see Figure 6-3). They are:

- Information: Contains basic information about the currently connected server, agent version running on the connected host, and the OS version.
- Security: Contains security settings for the connected agent. It lets you set host security and application security.
  - Host security defines the authorized user with administrator or root privileges.
  - Application security specifies the password for changing settings on the HBAs (for firmware or NVSRAM download or BIOS settings); the default password for the HBAs is *config*.

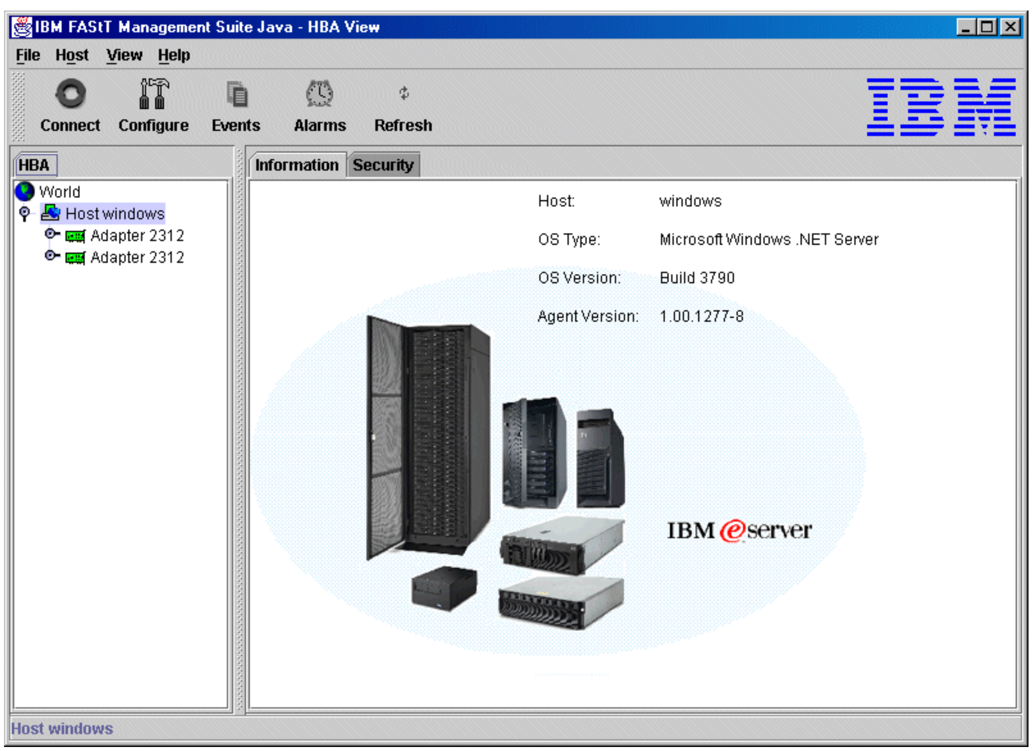

Figure 6-3 FAStT MSJ Host view

Viewing event and alarm logs

FAStT MSJ records an extensive set of information to the event and alarm logs. The logs are saved as text files (alarms.txt and events.txt) in the folder where FAStT MSJ is installed. FAStT MSJ can parse and view these logs in a window. To view these logs, click **Event Log** or **Alarm Log** from the View menu, or click the appropriate button on the button bar.

Using the FAStT MSJ diagnostic tools

When you click one of the host bus adapters in the HBA tree panel, the tab panel displays seven tabs as shown in Figure 6-2 on page 171.

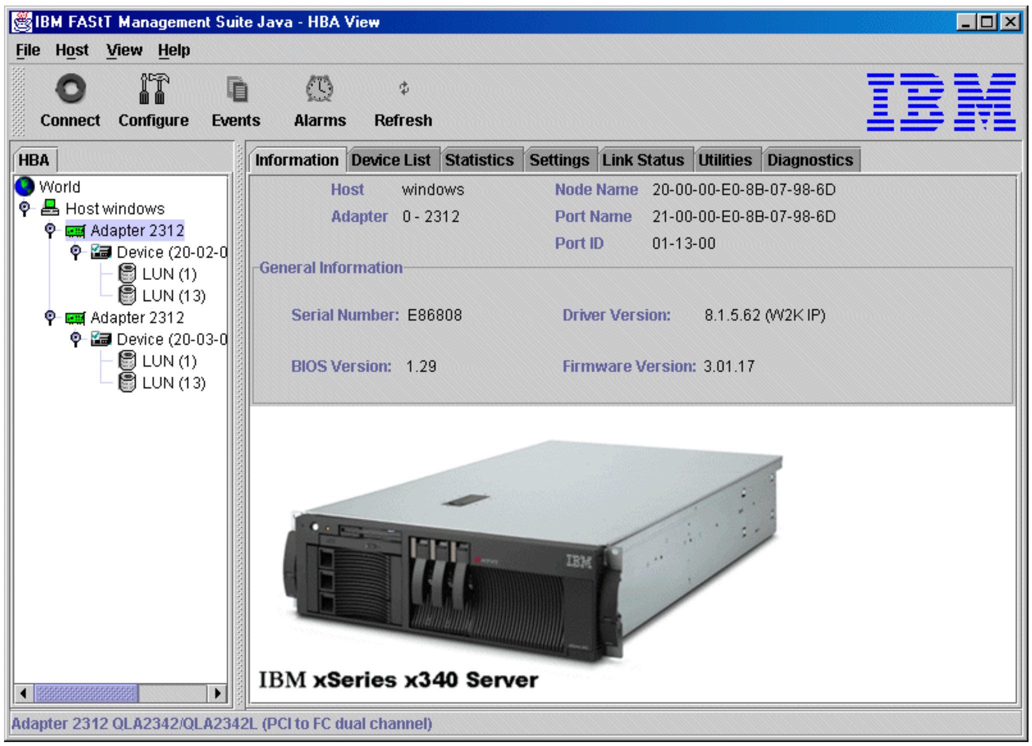

Figure 6-4 FAStT MSJ HBA view

The seven tabs are:

- Information: Displays general information about the server and host bus adapters, such as world-wide name, BIOS version, and driver version.
- 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 of time.
- NVSRAM Settings: Displays the current settings and allows you to make remote configuration changes to the NVSRAM of the adapters. All changes require a reboot of the server.
- Link Status: Displays link information for the devices attached to an adapter connected to a host.
- Utilities: Allows you to update the flash and NVSRAM remotely (not on Linux).
- Diagnostics: Allows you to run diagnostic tests remotely.

You can find detailed information on all these functions in the users guide that is bundled with FAStT MSJ download package.

We briefly introduce here some of the possibilities of the IBM FAStT MSJ.

Statistics

The Statistics panel displays the following information:

- Adapter Errors: The number of adapter errors reported by the adapter device driver (connection problem from or to switches or hubs).
- Device Errors: The number of device errors reported by the adapter device driver (I/O problems to FAStT, etc.); this usually gives the first hint about what path to the FAStT controller has a problem.
- Reset: The number of LIP resets reported by the adapter's driver. If you get increasing numbers there might be a communication problem between HBAs and storage.
- I/O Count: Total numbers of I/Os reported by the adapter's driver.
- IOPS (I/O per second): The current number of I/Os processed by the adapter.
- BPS (bytes per second): The current numbers of bytes processed by the adapter.

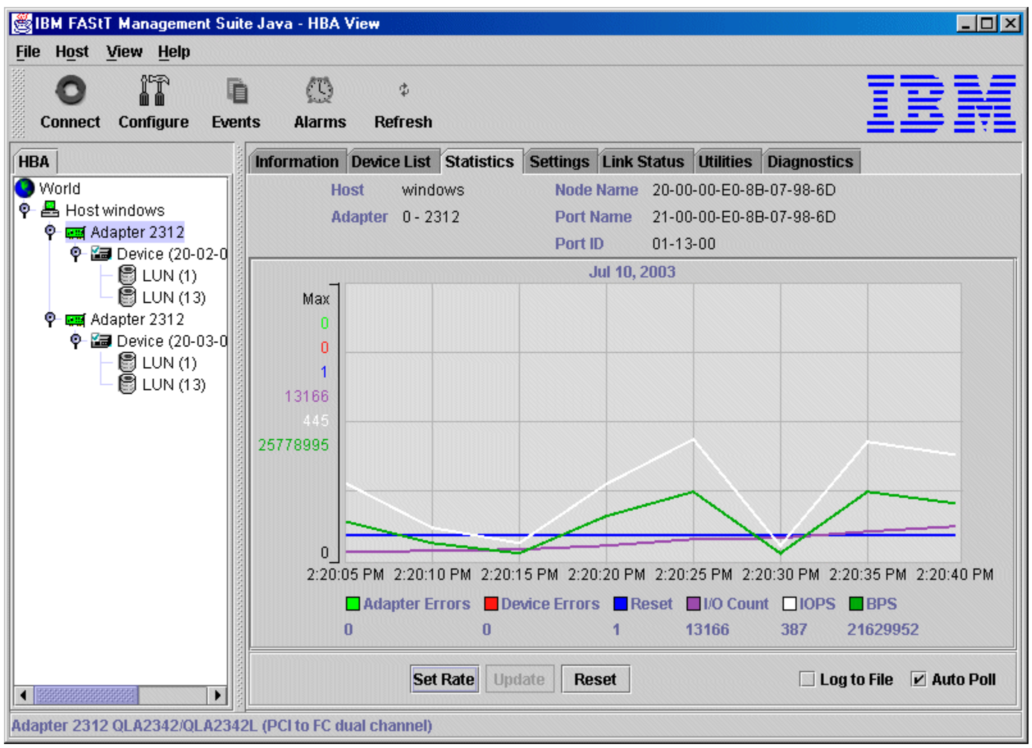

Figure 6-5 Statistics Window in IBM FAStT MSJ

Link Status

If you experience problems with connectivity or performance or you see entries from RDAC or the HBA driver in Windows, the Link Status tab (Figure 6-6 on page 174) is where you should start from to narrow down the device causing problems (faulty cable, gibbets, sets).

The following information can be retrieved from the Link Status window:

- Link Failure: The number of times the link failed. A link failure is a possible cause for a time-out (see Windows Event Log).
- Loss of Sync: The number of times the adapter had to re-synchronize the link.
- Loss Signal: The number of times the signal was lost (dropped and re-connected).

- Invalid CRC: The number of Cyclic Redundancy Check (CRC) errors that were detected by the device.

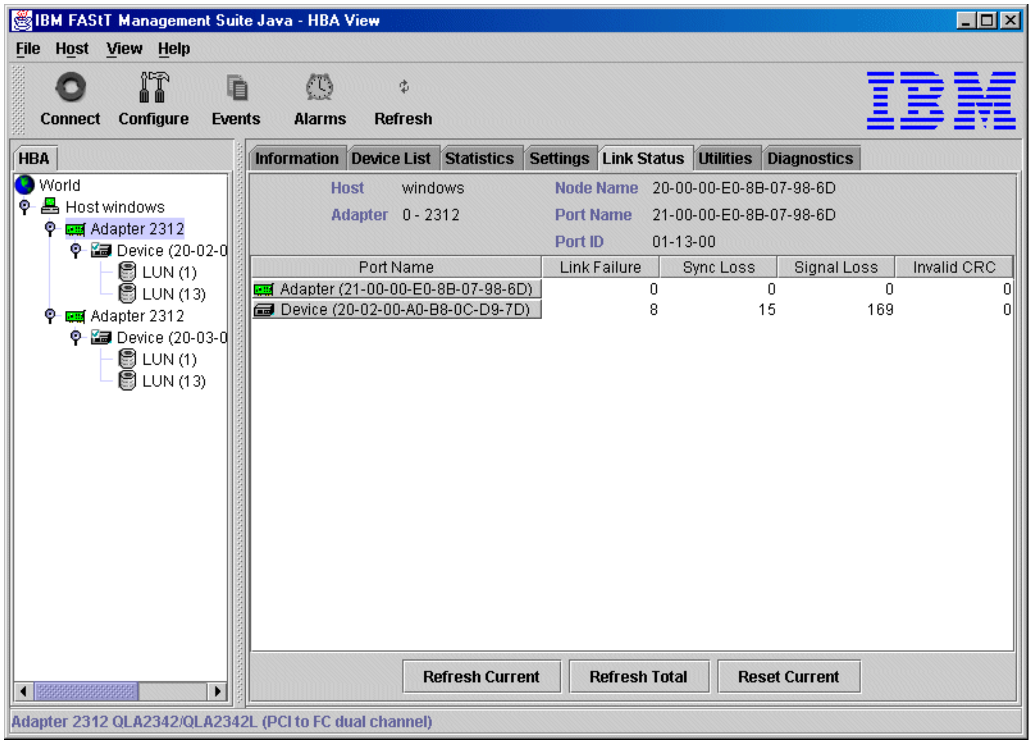

Figure 6-6 Link Status in IBM FAStT MSJ

Diagnostics

Using the Diagnostics panel (see Figure 6-7 on page 175) you can perform the loopback and read/write buffer tests.

- The loopback test is internal to the adapter. The test evaluates the Fibre Channel loop stability and error rate. The test transmits and receives (loops back) the specified data and checks for frame CRC, disparity, and length errors.
- The read/write buffer test sends data through the SCSI Write Buffer command to a target device, reads the data back through the SCSI Read Buffer command, and compares the data for errors. The test also compares the link status of the device before and after the read/write buffer test. If errors occur, the test indicates a broken or unreliable link between the adapter and the device.

The Diagnostics panel has three main parts:

- Identifying Information: Displays information about the adapter being tested
- Diagnostic Configuration Error: Contains testing options (like data patterns, number of tests, test increments)
- Loopback Test Results: Displays the results of a test showing whether the test passed or failed and error counters
  - For a loopback test, the test result includes the following information: Test Status, CRC Error, Disparity Error, and Frame Length Error.
  - For a read/write buffer test, the test result shows the following information: Loop ID/Status, Data Miscompare, Link Failure, Sync Loss, Signal Loss, and Invalid CRC.

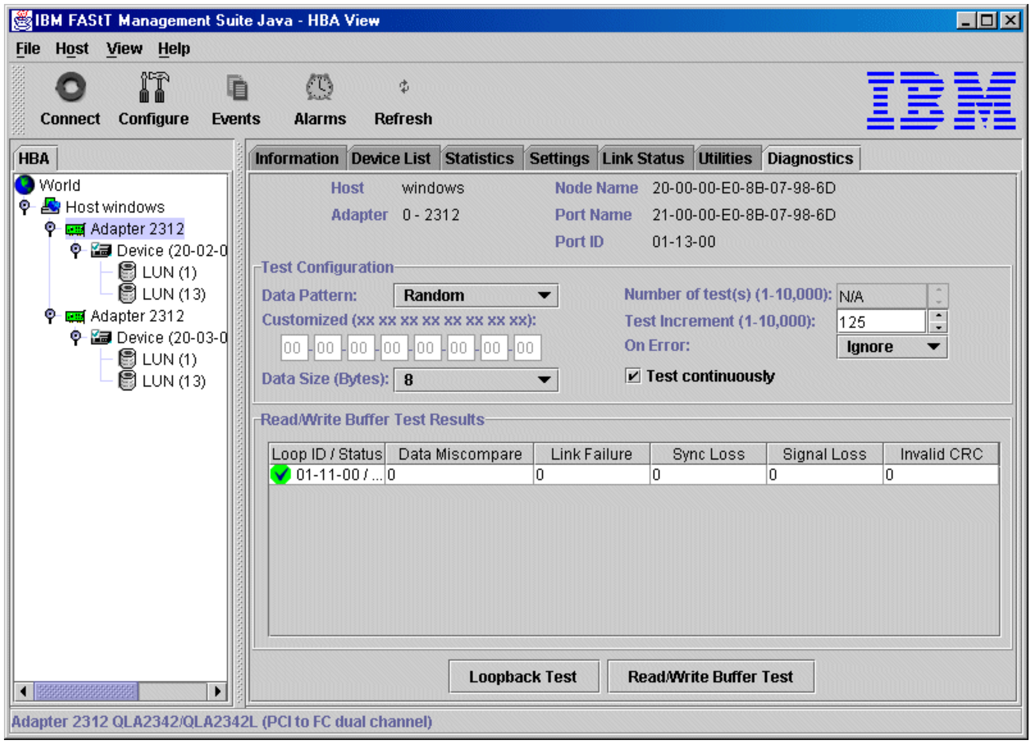

Figure 6-7 IBM FAStT MSJ diagnostics

6.1.2 FAStT and Storage Manager 8.4 error reporting and diagnostics

The FAStT Storage Server and Storage Manager offer several functions to log faults, notify users of faults, and guide them through the necessary recovery steps.

Controller diagnostics

The controller diagnostic feature allows a user to verify that a controller is functioning properly by performing various internal tests.

The diagnostics use a combination of three different tests: Read test, write test, and data loopback test. You should run all three tests at initial installation and whenever 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 will *not* be available for I/O. When diagnostics are completed, the controller should automatically allow data to be transferred to it. However, if there is a situation where data transfer is not re-enabled, select the controller and use the **Data Transfer -> Enable** option.

To run the diagnostics, highlight a controller from the Subsystem Management window. Then select either the **Controller -> Run Diagnostics** pull-down menu option, or **Run Diagnostics** from the right-mouse pop-up menu. The Run Diagnostics dialog (shown in Figure 6-8 on page 176) is displayed.

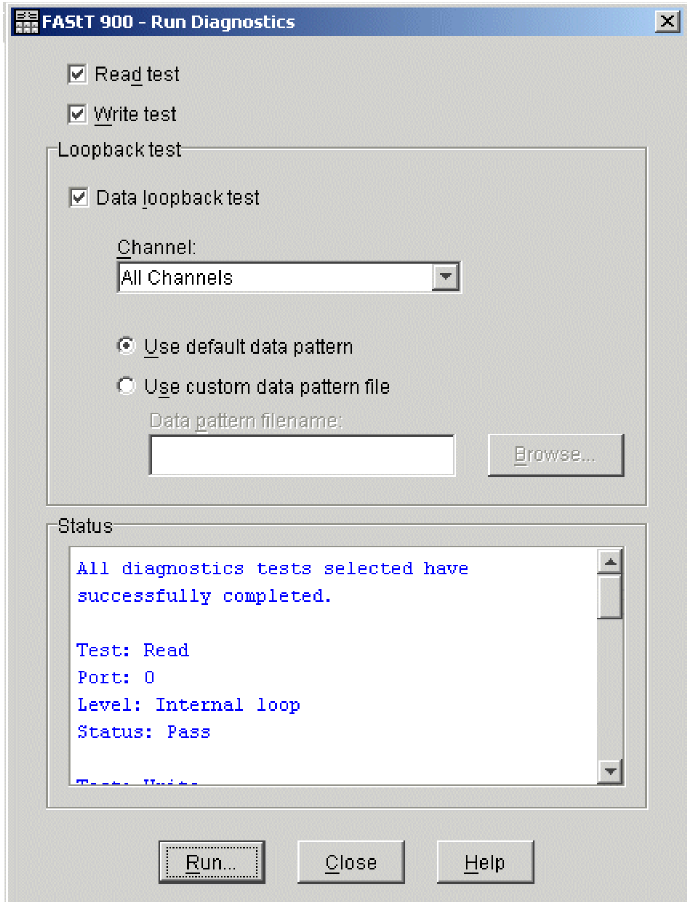

Figure 6-8 FAStT Controller DIagnostics

The checkboxes in the FAStT Run Diagnostics window are:

- 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: 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: This test can be run only on controllers that have Fibre Channel connections between the controller and the drives. The 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.

All test results are displayed in the diagnostics dialog status area.

Events are written to the Event Log when diagnostics are started, and when the tests have completed. These events will help you to evaluate whether diagnostics testing was successful or failed, and the reason for the failure. To view the Event Log, use the **View -> Event Log** pull-down menu option in the Subsystem Management window.

Major Event Log

The Major Event Log (MEL) is the primary source for troubleshooting the FAStT Storage Server. The MEL viewer can be used in two modes.

- View only critical events: For a quick overview of all events that might affect the operational status of the FAStT. For a full listing of these events see “Storage Manager Critical Events” on page 399.

  Storage Manager can forward all critical events through SNMP-trap or designated e-mail (see “Monitoring and alerting” on page 157).

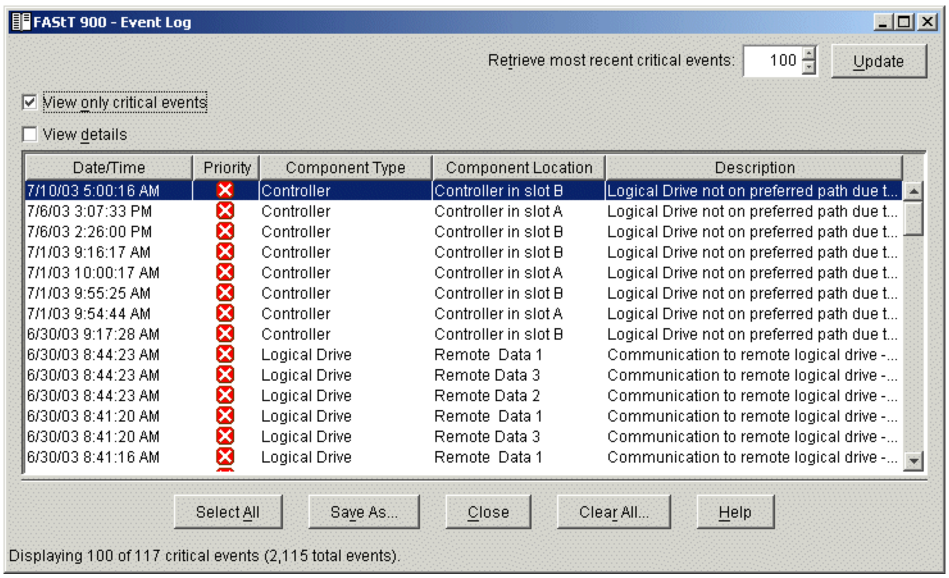

Figure 6-9 SM 8.4 Major Event Log

- View all events: Detailed information on all events logged by the FAStT controller. You can choose how many events you want to have listed. The maximum number is 8191.

Note that with Storage Manager 8.4, settings of the MEL viewer are persistent (see “MEL viewer settings” on page 20 for details).

If you troubleshoot your system you have to use the full Event Log since it will give you information about actions that took place before the actual critical event has happened.

Another reason to check the MEL is of preemptive nature. If you find, for instance, many destination driver errors to a specific drive, you should check the read link status diagnostics and possibly the operating system error logs for problems such as time-outs. An increased number of destination driver errors can indicate a faulty device; however, these logs are part of normal logging operation and should be examined in conjunction with other MEL entries. For example, if you encounter a problem where several attempts to rebuild a logical driver after drive failure never succeeds; the root causes of this problem are usually “bad blocks”, unreadable segments of the logical drive (not the physical HDD). In the event log these will be logged as follows.

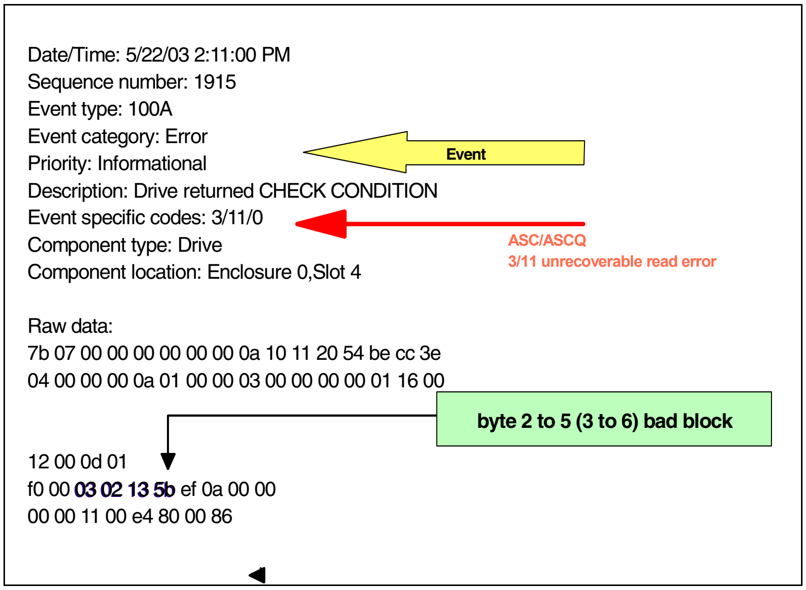

Figure 6-10 Failed rebuild logging

The actual drive taken out of service will be the drive that either the HSP or drive data are copied back to the parity data. You will find entries in the log similar to those in Figure 6-11. Even so the rebuild has failed the MEL will stat that the rebuild have finished and immediately after you will see that the drive is taken offline.

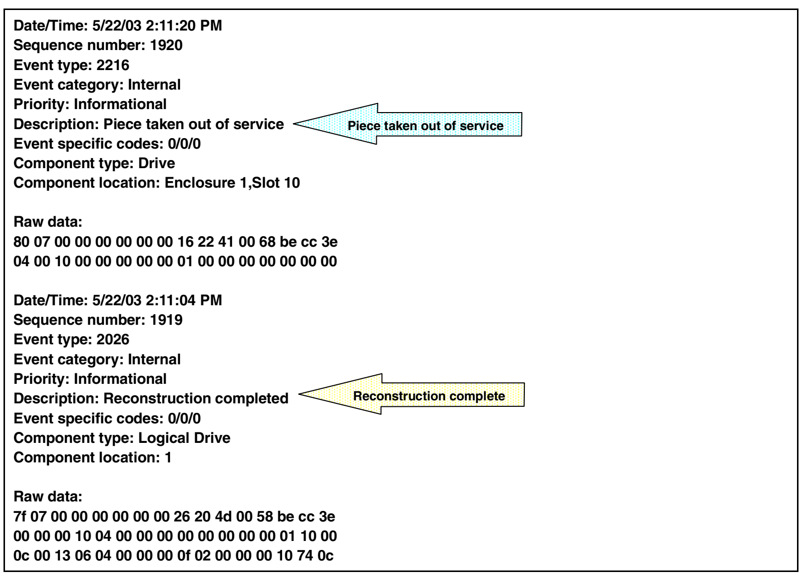

Figure 6-11 Drive taken out of service

There is no point replacing the drive as the bad blocks reside on a different drive. You have two possibilities:

- Back up your data, delete and recreate the logical drive, and enable media scan (preferred).
- Overwrite the bad block and start the rebuild again. However, the overwrite procedure has the drawback that the data in that block is lost and one cannot determine if there are more bad blocks on the logical drive. This procedure can only be done through the serial

interface connection or using a login client (this should only be done by technical support personnel).

Recovery Guru

Use Recovery Guru to help troubleshoot storage subsystem problems. The Recovery Guru will explain the cause of the problem. If necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

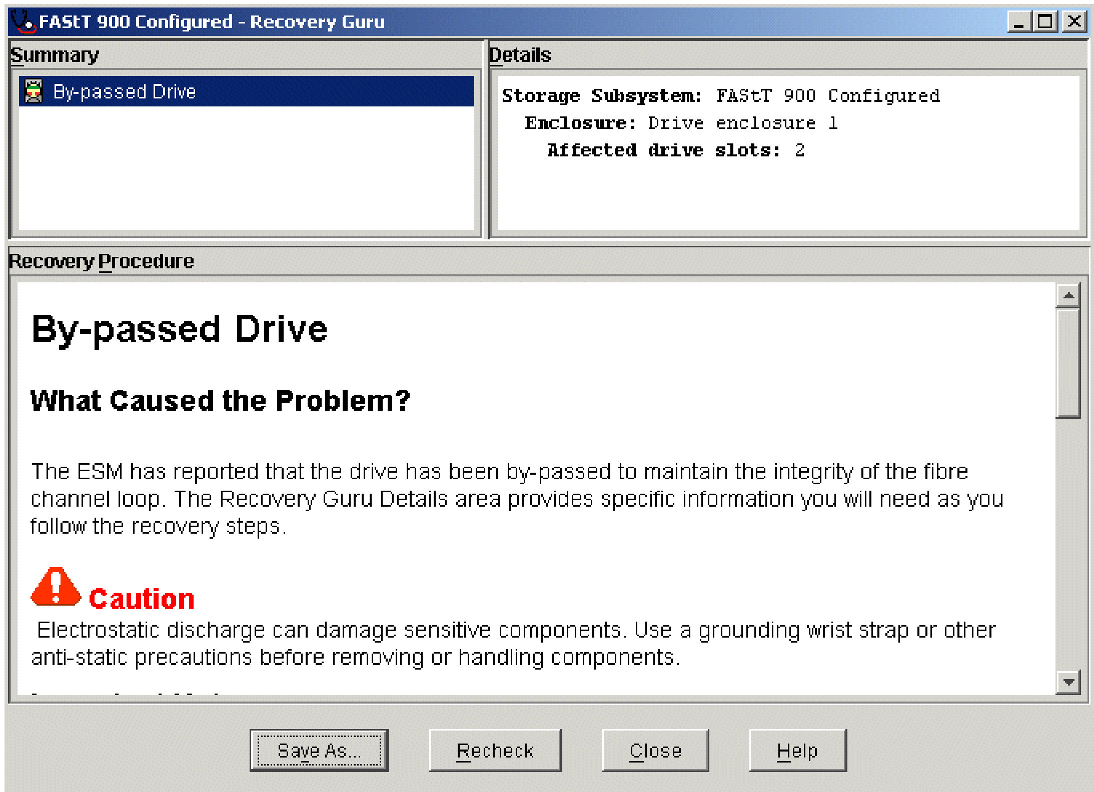

Figure 6-12 Storage Manager 8.4 Recovery Guru

You can find a list of Guru Alerts in the Appendix A, “FAStT Storage Manager 8.4 Recovery Guru Events” on page 21.

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

Read Link Status diagnostics

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. ITW errors can occur during normal operations but should not exceed 3–4 counts in 24 hours.

Usually RLS diagnostics should be run over a period of 24 hours with workload applied to the controller. Otherwise, with no I/O it is very unlikely that RLS will record anything.

From the Storage Subsystem menu right click **Storage Subsystem** select **Read Link Status Diagnostics**.

Figure 6-13 on page 182 show the RLS Diagnostics window. It lists all drives for each redundant drive's loop and the controllers. The channel numbers correspond with the drive-side mini-hubs.

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.
- ITW: The total number of Invalid Transmission Word errors detected on the Fibre Channel loop from the baseline time to the date and time. ITW may also be referred to as the "Received Bad Character Count".
- 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. Errors logged again a single drive point usually to a drive failure. If you see an increase in numbers against multiple drives across one channel the problem probably lays with cabling or connection (including mini-hubs or SFTs).

Errors logged against the ESM board can point to bad ESM boards or bad connectivity (cables, SFT, mini-hubs).

If you experience an increase on the controller blades there might be a problem at the mid plane or the controller itself.

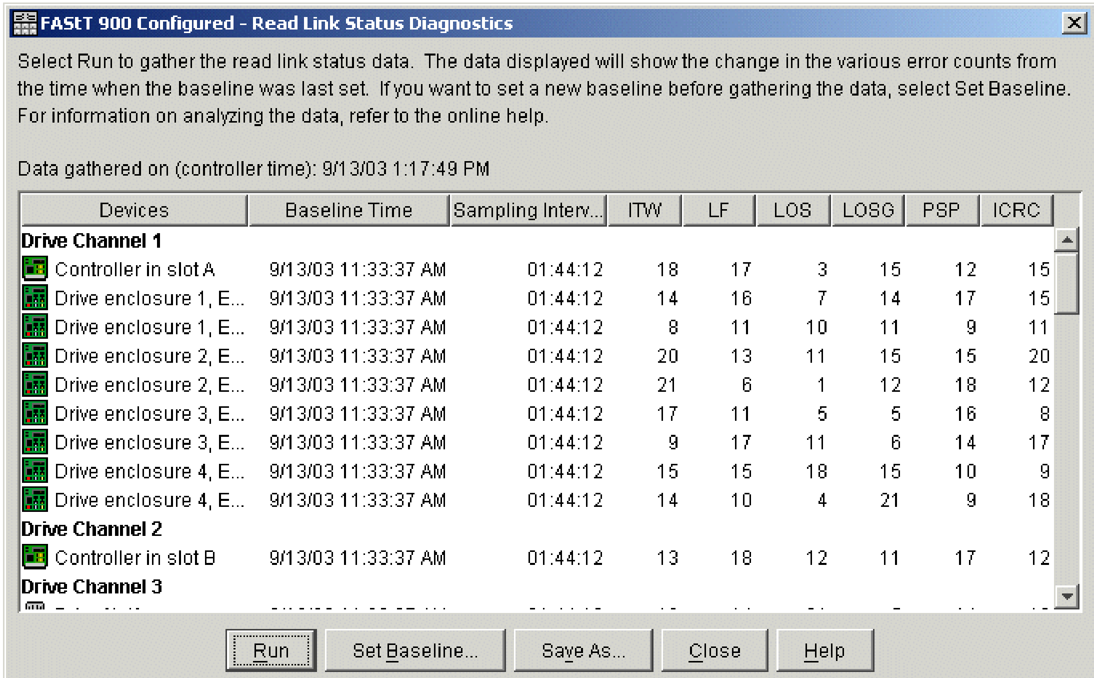

Figure 6-13 Storage Manager 8.4 Read Link Status Diagnostics

In our example in Figure 6-13 the baseline was set to September 13. When you open the RLSD window the most recent data is displayed. If you click **Run** the window will be updated (refreshed). In order to reset the counters, click **Set Baseline**. If you need to submit diagnostic data to the support staff, the file as a csv file. Figure 6-14 shows an example of RLSD file in .csv format.

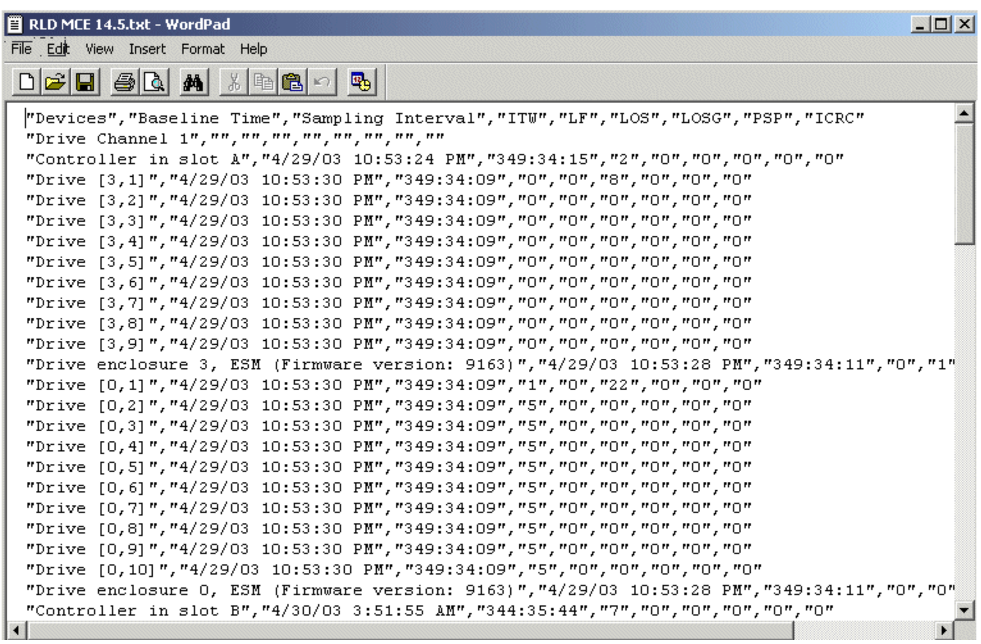

Figure 6-14 Example .csv file format

You can use the .csv format to import the data in a database for better readability. As an example, Figure 6-15 on page 183 shows the display, under Lotus Approach, of error information collected for a bad drive that generated errors across the drive loop. The event log shows that many destination drive errors were logged. Not shown in Figure 6-15 is that Drive Channel one was virtually error free. Since the drives have two ports, we can suspect that the port connected to Drive Channel 2 has a problem. In a fibre environment it is likely that the device that is downstream (AL-PA ID) from the error logging device is the faulty component. In the case below, however, the MEL pointed to drive 3,1 (which was ultimately replaced).

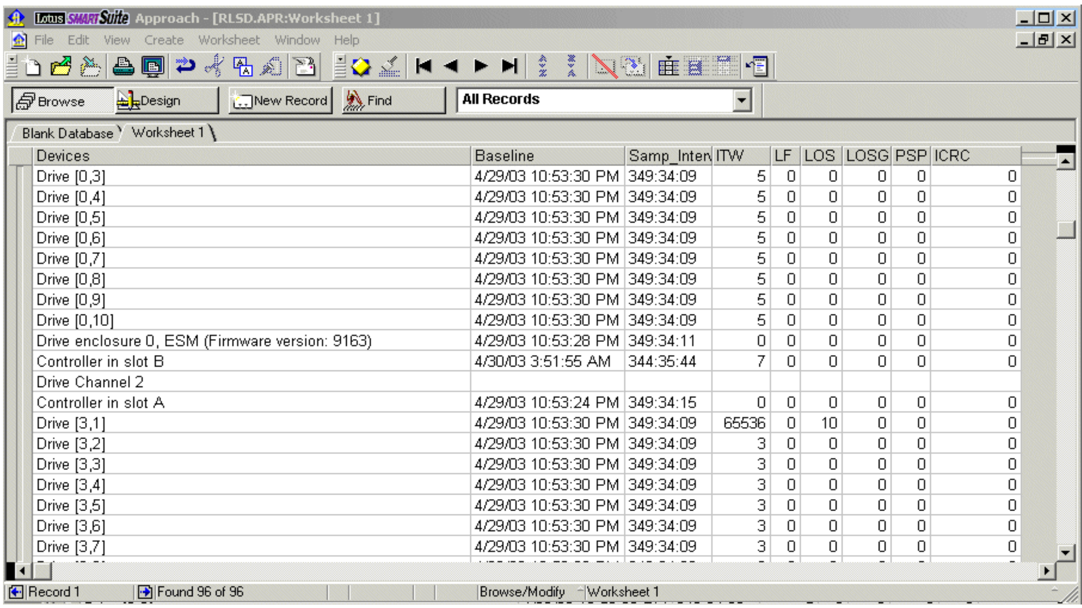

Devices	Baseline	Samp_Inten	ITW	LF	LOS	LOSG	PSP	ICRC
Drive [0,3]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,4]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,5]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,6]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,7]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,8]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,9]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive [0,10]	4/29/03 10:53:30 PM	349:34:09	5	0	0	0	0	0
Drive enclosure 0, ESM (Firmware version: 9163)	4/29/03 10:53:28 PM	349:34:11	0	0	0	0	0	0
Controller in slot B	4/30/03 3:51:55 AM	344:35:44	7	0	0	0	0	0
Drive Channel 2								
Controller in slot A	4/29/03 10:53:24 PM	349:34:15	0	0	0	0	0	0
Drive [3,1]	4/29/03 10:53:30 PM	349:34:09	65536	0	10	0	0	0
Drive [3,2]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0
Drive [3,3]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0
Drive [3,4]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0
Drive [3,5]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0
Drive [3,6]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0
Drive [3,7]	4/29/03 10:53:30 PM	349:34:09	3	0	0	0	0	0

Figure 6-15 RLSD Approach database

FAStT Storage Server diagnostics

The FAStT Storage Server and its EXPs have, depending on the model, server status and alert LED indicators. Table 6-1 shows you the meaning for the different LEDs.

Table 6-1 FAStT and EXP status LED

Problem indicator	Component	Possible cause	Possible solutions
Amber LED on	Drive	Drive failed	Replace the drive.
	Fan	Fan failed	Replace the fan.
	RAID Controller Fault LED	RAID controller failed	Replace the RAID Controller.
	Expansion port bypass LED	SFP port empty	It is normal.
		Fibre Channel cable is not attached to the expansion unit	It is normal.
		No incoming signal detected	Reattach the SFPs and Fibre Channel cables. Replace input and output SFPs or cables as necessary.

Problem indicator	Component	Possible cause	Possible solutions
	Front Panel	General system error	Indicates that a Fault LED somewhere on the storage server has turned on.
Amber LED is on and green LED is off	Power Supply	Power switch is turned off or AC power failure	Turn on all power-supply power switches.
Amber and green LED on	Power-Supply	Power supply failure	Replace the failed power-supply.
All green LEDs off	Every component	Storage subsystem power is off	Check power cord and PDU.
		AC power failure	Check main circuit breaker and AC outlet.
		Power supply failure	Replace power supply.
		Mid-plane failure	Replace the Mid-plane.
Amber LED flashing	Drive	Drive rebuild or identity is in process	It is normal.
One or more green LEDs off	Power Supply	Power cord unplugged or switches off	Check power cord and power supply switch.
	All Drive	Mid-plane failure	Replace the mid-plane.
	Front panel	Power Supply problem	Make sure power is good.
		Hardware Failure	If any other LEDs are on, replace mid-plane.
	Battery	Battery failure	Replace battery.
	Cache active	The cache is disabled, the cache has failed, battery failure	Use Storage Manager to enable the cache; replace the RAID controller; replace the battery.
	Host Loop	Host, managed hub or switch is off or failed	Check if hub or switch is on; replace attached devices that have failed.
		Fibre Channel cable failed	Check hub or switch; replace attached device that has failed.
		SFP failed	Ensure SFP is seated properly; replace SFP.
		RAID Controller has no power or failed	Ensure that the unit is powered on. Replace RAID controller.

Problem indicator	Component	Possible cause	Possible solutions
	Expansion loop	Drives are improperly installed or not installed	Ensure that the drives are properly installed.
		RAID controller has no power or has failed	Ensure that the unit is powered on. Replace the RAID controller.
		Expansion port device failed	Replace the drive; replace the expansion unit SFP or Fibre Channel cable.
Intermittent or sporadic power loss to the storage server	Some or all component	Defective AC power source or partially plugged-in power cord	Check AC power source. All power unit.
		Power supply failed	Check for Fault LED on the power supply, and replace bad one.
		Mid-plane failed	Replace mid-plane.
Unable to access drive	Drives and Fibre Channel loop	Fibre Channel cabling failed	Check Fibre Channel cables are undamaged and properly connected.
		RAID controller failed	Replace the RAID controller.
		SFP failed	Ensure SFP is seated properly; replace SFP.
Random errors	Subsystem	Mid-plane failed.	Replace mid-plane.

For details:

- For the FAStT 600 see “FAStT600 indicator lights” on page 32.
- For the FAStT 700/900, see “FAStT 900 indicator lights” on page 41 and “RAID controller” on page 29.

6.2 Operating systems

This section provides practical information for problem determination on specific OS platforms. The intent is to give you a starting point.

6.2.1 Windows 2000

For troubleshooting problems in Windows 2000/2003, look first in the event log. There you can find information logged by either the RDAC driver or the adapter (Qlogic) driver. For example, you could have the following:

- RDAC entries. In order to decode RDAC problems display the log in Word data format. The offset 0x010 displays the reason why RDAC logged an entry, in our case 0000126. For a description see Figure 6-16 on page 186. ID 18 events with qualifier value 126 is often seen in a clustering environment when the not owning note boots up and tries to access the drives but be rejected because of a SCSI reservation.

Status Code	Message
0x116	The maximum number of paths to a controller have been exceeded. RDAC currently supports only one path to a controller.
0x117	WWN mismatch on two sides of a LUN.
0x118	Unable to allocate an RDAC structure.
0x122	A controller failover occurred due to an I/O error. Examples of this include resets and LUNS becoming ready.
0x123	A path to a multi-path controller has failed.
0x124	A controller failover has failed. No retries will be performed.
0x125	Controller failover event due to all paths being removed.
0x126	Controller failover due to some LUNs being owned by an inaccessible alternate controller at startup.
0x201	SCSI recovered error on a Controller.
0x202	SCSI hardware error on a Controller.
0x203	SCSI unit attention on a Controller.
0x204	SCSI aborted command on a Controller.
0x205	SCSI media error on a Controller.

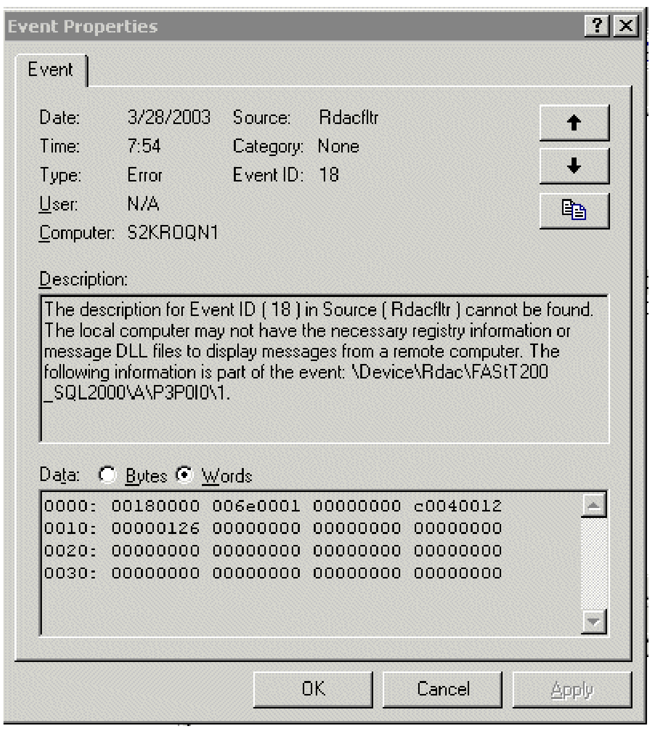

Figure 6-16 RDAC event log entry in Windows 2000

- The second example (Figure 6-17 on page 187) is a Qlogic 2300 driver entry. The relevant data for the user is the event ID number; however, there is no more additional information available and you need to consult other sources such as the MEL or the switch/hub event logs for further problem determination.

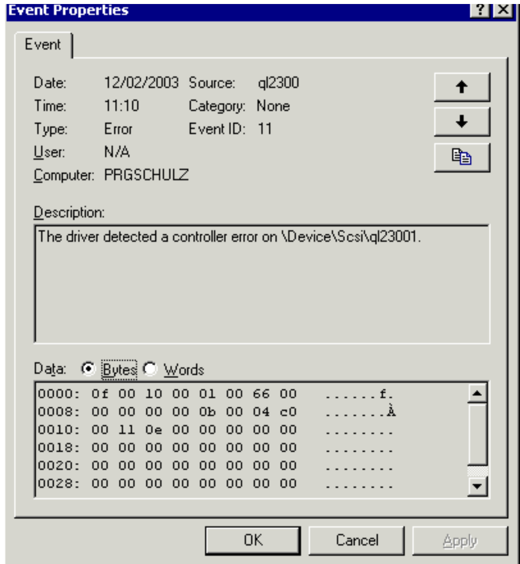

Figure 6-17 Qlogic Driver Entry in Windows 2000 Event Log

For more information on problem determination and troubleshooting refer to the *Hardware Maintenance Manual and Problem Determination Guide,* GC26-7528-03, available at:

ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers/hmpd3.pdf

6.2.2 Linux

In Linux there are two places you should check when running into problems.

- The message files in /var/log/message and /var/log/dmsg.
- The SCSI process queue. To check it run the following from the command line:

  `ls /proc/scsi/qla2300/`

  This command returns the list of the SCSI host bus adapters and their IDs (HBA IDs as seen by the OS; see Figure 6-18). Note that the driver name depends on your HBA version; if you are using qla2200 adapters, you will see need to substitute qla2300 with qla2200.

  Figure 6-18 shows the output from the **ls** command.

```
[root@linux root]# ls -al /proc/scsi/qla2300/
total 0
dr-xr-xr-x    2 root     root           0 Jul 11 12:41 .
dr-xr-xr-x    4 root     root           0 Jul 11 12:41 ..
-rw-r--r--    1 root     root           0 Jul 11 12:41 1
-rw-r--r--    1 root     root           0 Jul 11 12:41 2
crwxrwxrwx    1 root     root      254,   0 Jul 11 12:41 HbaApiNode
[root@linux root]#
```

Figure 6-18 Linux Red Hat ls command

  `less /proc/scsi/qla2300/[number of the SCIS HBA ID]`

  For example, the output list (Figure 6-19 on page 188).

If you are not sure whether the module is loaded at all, run **lsmod** from the command prompt. You get a list of all the modules loaded.

```
[root@linux root]# lsmod
Module                  Size  Used by    Tainted: P
ide-cd                 35296   0  (autoclean)
cdrom                  35520   0  (autoclean) [ide-cd]
soundcore               7940   0  (autoclean)
autofs                 13796   0  (autoclean) (unused)
e1000                  52740   0  (unused)
eepro100               21968   1
usb-ohci               23392   0  (unused)
usbcore                68864   1  [usb-ohci]
ext3                   73536   5
jbd                    55048   5  [ext3]
qla2300               252256   0
aic7xxx               127200   6
sd_mod                 13468   6
scsi_mod              124988   3  [qla2300 aic7xxx sd_mod]
[root@linux root]#
```

Figure 6-19 Output from lsmod

When Linux is loading all of the drivers that it uses, it will log some messages into a log file. This log file can be looked at with the **cat** command. The messages file is found in /var/log/messages; some of the same messages can be found by running the **dmesg** command form the command prompt.

To view the /var/log/messages file, use the command **cat /var/log/messages**.

To view the dmesg type **dmesg**.

```
Jul 10 13:59:45 linux kernel: Partition check:
Jul 10 13:59:45 linux kernel:  sda: sda1 sda2 sda3 sda4 < sda5 sda6 sda7 >
Jul 10 13:59:45 linux kernel: qla2x00_set_info starts at address = f8848060
Jul 10 13:59:45 linux kernel: qla2x00: Found  VID=1077 DID=2312 SSVID=1077 SSDID
=100
Jul 10 13:59:45 linux kernel: scsi(1): Found a QLA2312  @ bus 10, device 0x1, ir
q 29, iobase 0x3000
Jul 10 13:59:45 linux kernel: scsi(1): Allocated 4096 SRB(s).
Jul 10 13:59:45 linux kernel: scsi(1): Configure NVRAM parameters...
Jul 10 13:59:45 linux kernel: scsi(1): 64 Bit PCI Addressing Enabled.
Jul 10 13:59:45 linux kernel: scsi(1): Scatter/Gather entries= 896
Jul 10 13:59:45 linux kernel: scsi(1): Verifying loaded RISC code...
Jul 10 13:59:45 linux kernel: scsi(1): Verifying chip...
Jul 10 13:59:45 linux kernel: scsi(1): Waiting for LIP to complete...
Jul 10 13:59:45 linux kernel: scsi(1): Cable is unplugged...
Jul 10 13:59:45 linux kernel: qla020: ConfigRequired is set.
Jul 10 13:59:45 linux kernel: qla2x00: Found  VID=1077 DID=2312 SSVID=1077 SSDID
=100
Jul 10 13:59:45 linux kernel: scsi(2): Found a QLA2312  @ bus 10, device 0x1, ir
q 30, iobase 0x3100
```

Figure 6-20 Sample information from /var/log/messages

```
Attached scsi disk sda at scsi0, channel 0, id 13, lun 0
SCSI device sda: 71096640 512-byte hdwr sectors (36401 MB)
Partition check:
 sda: sda1 sda2 sda3 sda4 < sda5 sda6 sda7 >
qla2x00_set_info starts at address = f8848060
qla2x00: Found  VID=1077 DID=2312 SSVID=1077 SSDID=100
scsi(1): Found a QLA2312  @ bus 10, device 0x1, irq 29, iobase 0x3000
scsi(1): Allocated 4096 SRB(s).
scsi(1): Configure NVRAM parameters...
scsi(1): 64 Bit PCI Addressing Enabled.
scsi(1): Scatter/Gather entries= 896
scsi(1): Verifying loaded RISC code...
scsi(1): Verifying chip...
scsi(1): Waiting for LIP to complete...
scsi(1): Cable is unplugged...
qla020: ConfigRequired is set.
qla2x00: Found  VID=1077 DID=2312 SSVID=1077 SSDID=100
scsi(2): Found a QLA2312  @ bus 10, device 0x1, irq 30, iobase 0x3100
```

Figure 6-21 Sample dmesg output

6.2.3 AIX

Table 6-2 shows a list of possible disk array errors that could be reported in the AIX error log. You can view the AIX error log by running the **errpt -a** command. You can also check your Storage Manager Major Event Log (MEL) to find out whether there is any correlation between the host, SAN, and FAStT Storage Subsystem.

You might need to validate your configuration or replace defective hardware to correct the situation.

Table 6-2 AIX error messages

AIX category	AIX error description	Explanation
FCP_ARRAY_ERR1	ARRAY OPERATION ERROR	A permanent hardware error involving the disk array media.
FCP_ARRAY_ERR2	ARRAY OPERATION ERRO	A permanent hardware error.
FCP_ARRAY_ERR3	ARRAY OPERATION ERROR	A permanent error detected by the array adapter.
FCP_ARRAY_ERR4	ARRAY OPERATION ERROR	A temporary error within the array communications adapter, and so on.
FCP_ARRAY_ERR5	UNDETERMINED ERROR	
FCP_ARRAY_ERR6	SUBSYSTEM COMPONENT FAILURE	A degradation condition has occurred other than a disk drive.

AIX category	AIX error description	Explanation
FCP_ARRAY_ERR7	CONTROLLER HEALTH CHECK FAILURE	A health check on the passive controller has failed.
FCP_ARRAY_ERR8	ARRAY CONTROLLER SWITCH	
FCP_ARRAY_ERR9	ARRAY CONTROLLER SWITCH FAILURE	An array controller switch has failed.
FCP_ARRAY_ERR10	ARRAY CONFIGURATION CHANGED	A logical unit has been moved from one controller to the other (most likely by the action of an alternate host).
FCP_ARRAY_ERR11	IMPROPER DRIVE TYPE FOR DUAL ACTIVE MODE	This error should not be possible on the 2102 array and exists for history reasons only. It may be reused for a different error in the future.
FCP_ARRAY_ERR12	POLLED AEN FAILURE	An automatic error notification has failed.
FCP_ARRAY_ERR13	ARRAY INTER-CONTROLLER COMMUNICATION FAILURE	
FCP_ARRAY_ERR14	ARRAY DRIVE FAILURE	
FCP_ARRAY_ERR15	CACHE BATTERY LOW/DATA LOSS POSSIBLE	If a controller card is replaced it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the error log. Do not replace the controller.

AIX category	AIX error description	Explanation
FCP_ARRAY_ERR16	CACHE BATTERY CHARGE BELOW 87.5%	If a controller card is replaced it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the error log. Do not replace the controller.
FCP_ARRAY_ERR17	WORLDWIDE NAME CHANGED	A controller has changed world-wide names (most likely either it was replaced without placing it in the reset state first or the cabling was changed so that a different controller with the same SCSI ID is on the loop).
FCP_ARRAY_ERR18	RESERVATION CONFLICT	An operation failed because the disk array volume (LUN) is reserved by another host.

6.2.4 Sun Solaris

For troubleshooting SUN/Solaris systems look at /var/adm/messages.

During boot, the driver will log whether it is loaded correctly and how many LUNs are connected to the adapter. Look further in /kernel/drv/.

Examine the .conf file for your corresponding adapter, for example, Jinc146.conf or fcaw.conf and the sd.conf for attached disk listing.

1. Run the following command to verify that your host can see your external FAStT storage: **`#cd /etc/raid/bin/SMdevices`**. You must see two access LUNS, and all configured storage (logical drives). If any of these are missing, run the following command:

   ```
   #cd /etc/raid/bin/hot_add
   ```

2. Run the following command to cross-verify between SMutils and what the host sees:

   ```
   #format
   ```

You must see two access LUNS, and all configured storage (logical drives). If you still do not see all your storage check the following:

- Recheck your zones.
- Make sure the zones have been enabled.
- Recheck the targets and WWPNs.
- Check to see that the FAStT Controllers and expansion drawers are powered on.
- Check to see that the switches are powered on.

The latest driver and links are available form the support Web site.

6.2.5 Netware

With NetWare 6.0 the system log screen was introduced. That screen will log all events to a file and display them to the server console.

For initial troubleshooting check the screen for any output.

In order to confirm that the server can see all attached LUNs at the server console run the following commands:

```
Scan all luns
```

And:

```
List all devices
```

That will list all accessible devices seen by the server.

> **Note:** *List All Devices* will show all LUNs to all paths available. In an multi HBA adapter setup you should see all LUNs listed twice. If not there is a problem with the path on one adapter or wrong partitioning on the controller side, install IBM FAStT MSJ for an easier way of troubleshooting the problem.

.

7

FlashCopy

We have introduced FashCopy in “FlashCopy overview” on page 100; this chapter covers the FlashCopy feature in details, and looks at 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 FlashCopy features either through the Storage Manager client GUI or using CLI and scripts.

7.1 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 later.

For a general overview of the FlashCopy functions, refer to “FlashCopy overview” on page 100; this section describes the various parameters for FlashCopy and what they mean.

7.1.1 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 “Script Editor and command line interface” on page 89 for more information.

Note the following points:

- Refer to Appendix A, “Additional host-specific instructions for FlashCopy logical drives” on page 351. 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 are 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 351, 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 "Disabling and recreating a FlashCopy logical drive" on page 195.

7.1.2 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 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 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.

7.1.3 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 increase 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 drives are no longer accessible. Missing logical drives, in most cases, are recoverable. See “Missing logical drives” on page 179.

FlashCopy logical drive maintenance

The default FlashCopy repository logical drive capacity is set to 20 percent 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 percent.

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 “Estimating FlashCopy repository life” on page 200.

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 opens (Figure 7-1).

Figure 7-1 FlashCopy Repository Logical Drive Properties

3. Select the **Capacity** tab to view the currently defined settings (Figure 7-2).

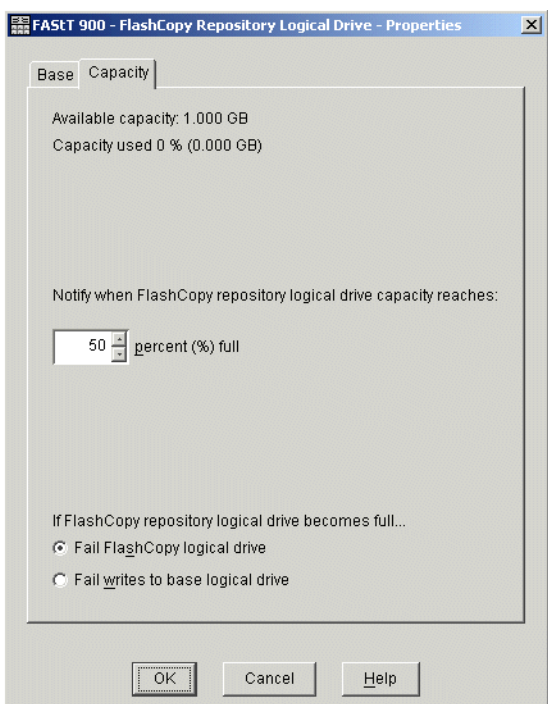

Figure 7-2 Repository capacity settings

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 “Increasing the capacity of a FlashCopy repository logical drive” on page 201.

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.

7.1.4 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 percent 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 “Estimating FlashCopy repository life” on page 200.
- 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 percent 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.

   ```
   [5 x 1024(K) x 1024(M) x 1024(G)]
   ```

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.

   ```
   [2,684,354.56 / 1024(K)]
   ```

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 percent 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 7-2 on page 197).

> **Note:** The percentage (%) full 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.)

7.1.5 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 percent 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, and 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 percent full during the FlashCopy's expected lifetime, then you should increase the FlashCopy repository logical drive's capacity. For more information, see "Increasing the capacity of a FlashCopy repository logical drive" on page 201.

7.1.6 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 164 for more information.

7.2 FlashCopy: Step-by-step

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
- Disabling a FlashCopy drive
- Re-creating a FlashCopy drive
- Resizing a FlashCopy repository drive
- Deleting a FlashCopy drive

7.2.1 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 select **Premium Features -> List** (see Figure 7-3).

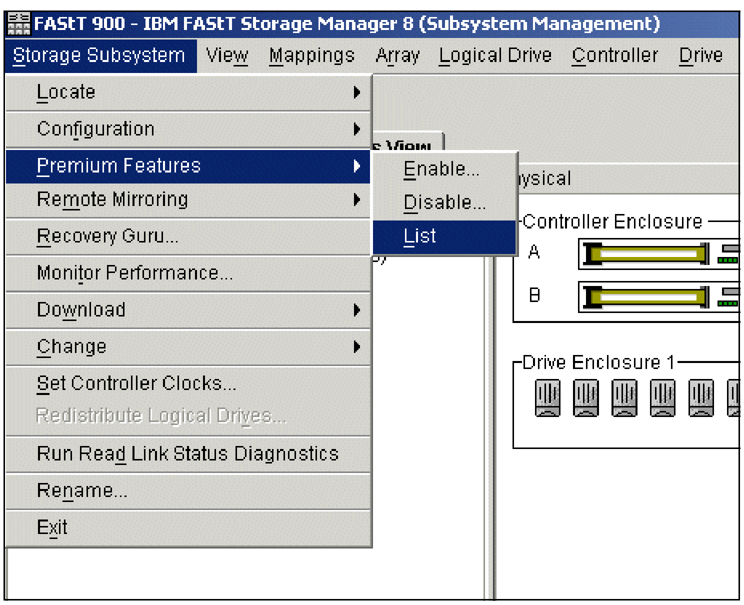

Figure 7-3 Listing premium features

   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 7-4 on page 203.

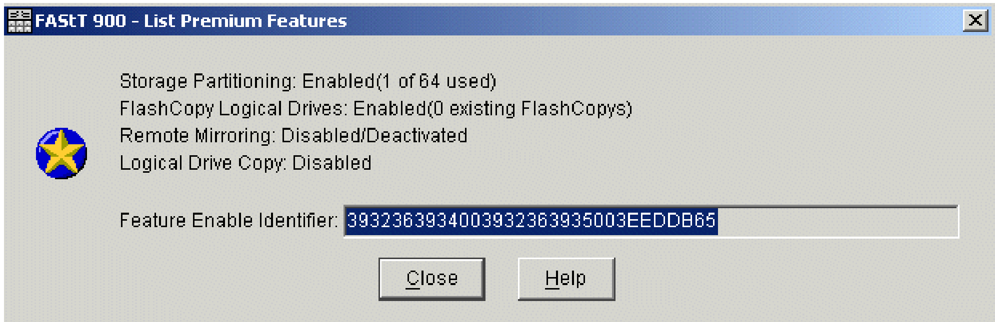

Figure 7-4 Premium features - FC Enabled

Figure 7-5 demonstrates an example where the FlashCopy copy feature is not enabled.

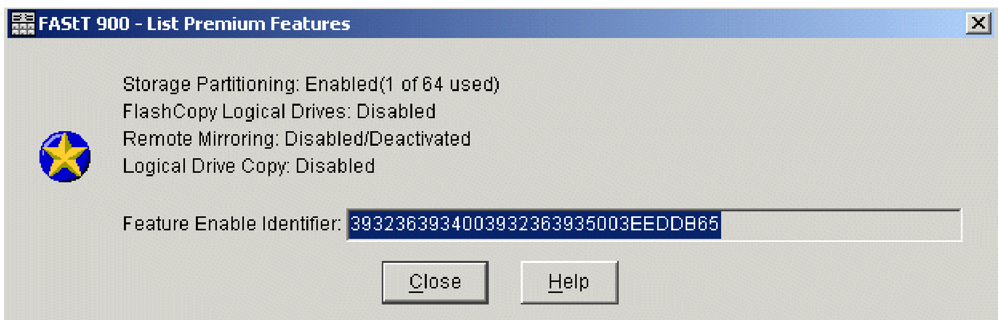

Figure 7-5 Premium features - FC disabled

If you identify that the FlashCopy feature is not enabled, see "Enabling the FlashCopy premium feature" on page 204 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 Subsystem Management window, as shown in Figure 7-6.

Figure 7-6 FlashCopy feature icon enabled

The example in Figure 7-7 shows a disabled FlashCopy feature.

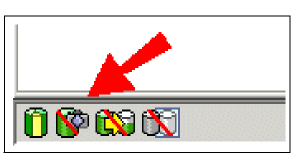

Figure 7-7 FlashCopy feature icon disabled

7.2.2 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 (Figure 7-8).

> **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**.

Figure 7-8 Select Feature Key File

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**.

7.2.3 Creating a FlashCopy drive

This section takes you through the procedure to create a FlashCopy drive from a base volume.

> **Important:** The following is important:
>
> - Refer to Appendix A, "Additional host-specific instructions for FlashCopy logical drives" on page 351, for additional instructions on the operating system-specific procedures. Failure to complete the additional steps may result in an inaccurate point-in-time image.f 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 a quiesced state for the duration of the FlashCopy creation.

b. It is also important to back up application recovery files, such as rollback 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 7-9.

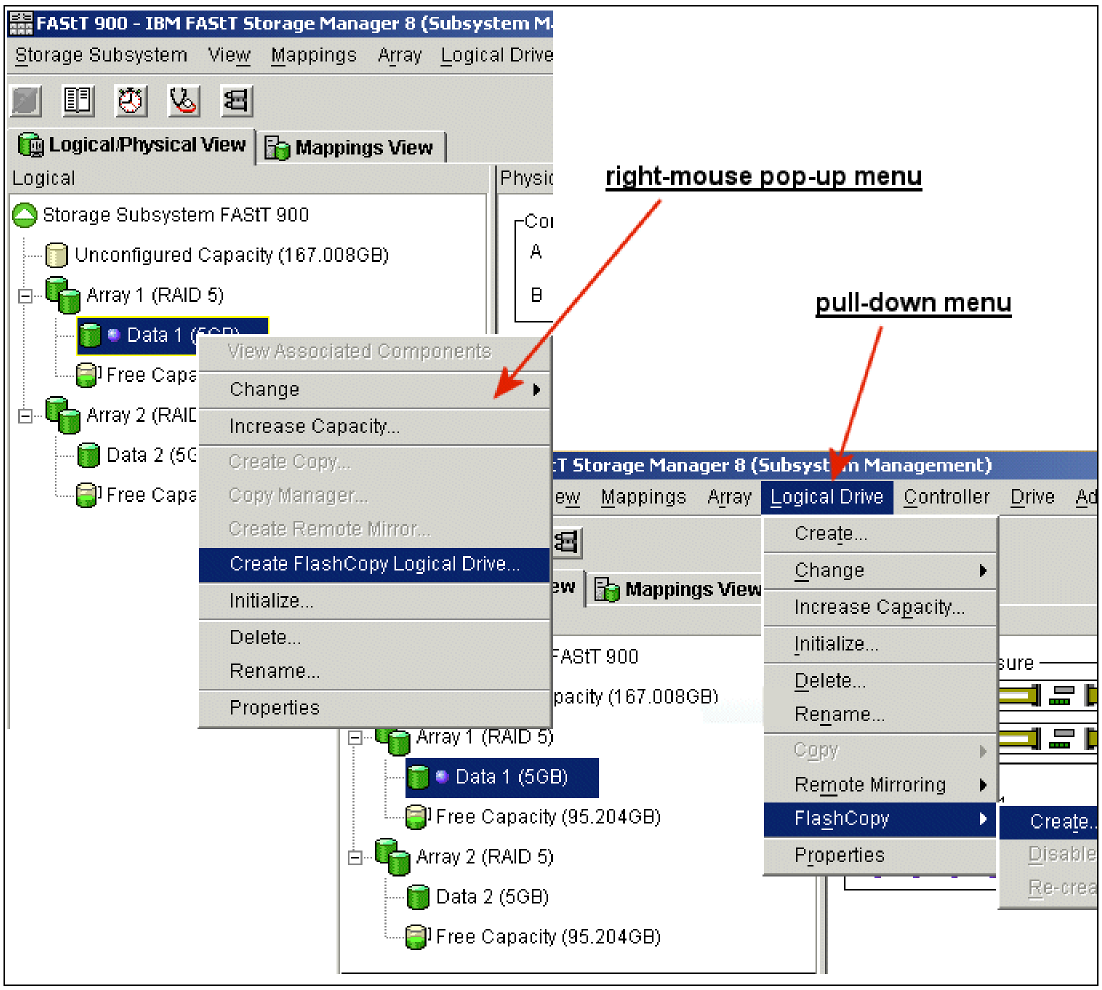

Figure 7-9 Create FlashCopy menu options

The Create FlashCopy Logical Drive Wizard begins as shown in Figure 7-10 on page 206.

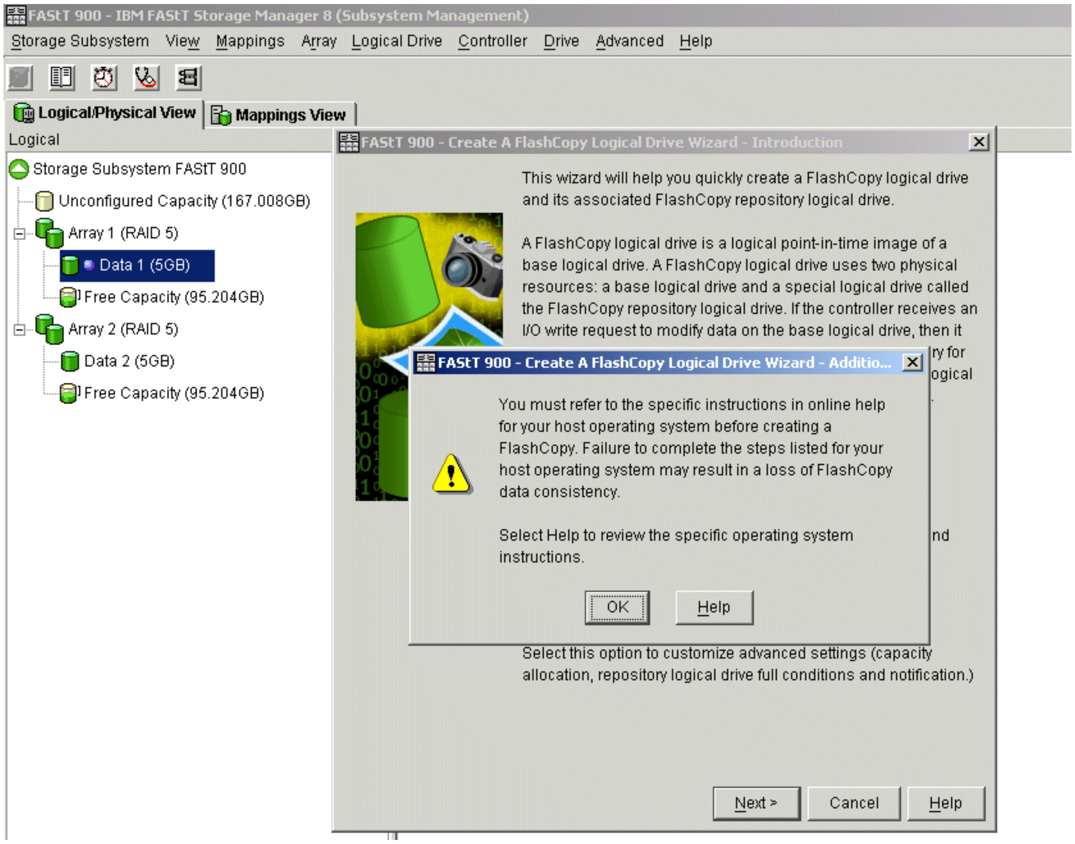

Figure 7-10 Create FlashCopy Logical Drive Wizard startup

> **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 7-10. Click **OK** to proceed to the wizard introduction window.

      Follow the instructions on each wizard panel, and click the Next button when you are 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.

      The Introduction window (Figure 7-11 on page 207) 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.

b. If you choose the simple path, as displayed in Figure 7-11, then the following steps are followed:

   i. The Specify Names panel is displayed. On the Specify Names panel (Figure 7-12 on page 208), you define the FlashCopy logical drive name and the name of its associated FlashCopy repository logical drive.

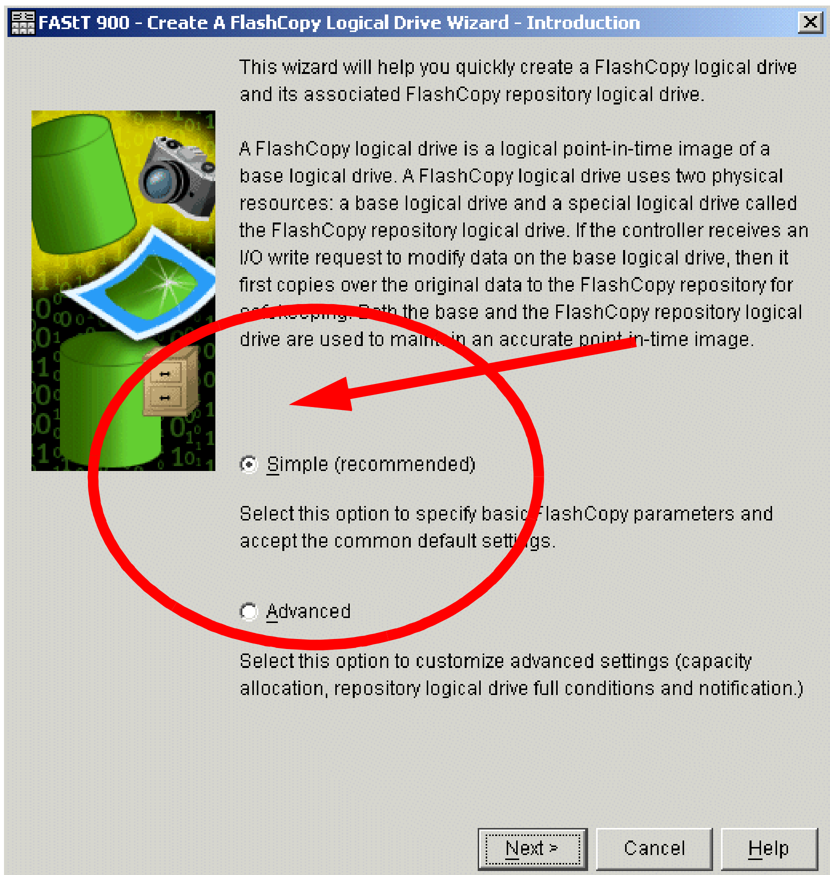

Figure 7-11 Create FC Logical Drive Wizard Introduction - Simple

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 fourth volume. For example, if you are creating the first FlashCopy logical drive for a base logical drive called Data 1, then the default FlashCopy logical drive name is Data 1-1, and the associated FlashCopy repository logical drive default name is Data 1-R1. The default name of the next FlashCopy logical drive you create based on Data 1 is Data 1-2, with the corresponding FlashCopy repository logical drive named Data 1-R2 by default.

Change the default names if required.

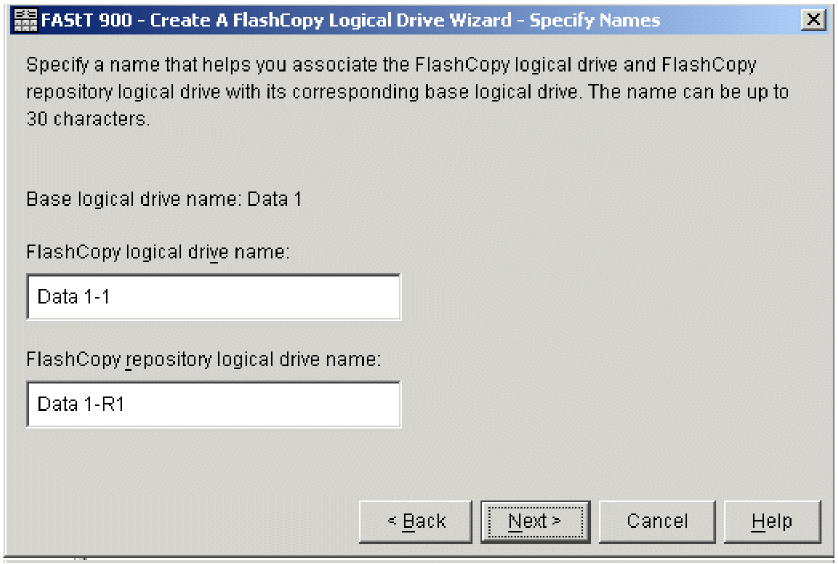

Figure 7-12 Create FC Logical Drive Wizard - Specify Names

Tips: Read over the following 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 Data 1 *DataVolMay28*, and do not use the software-supplied sequence number of 1. Then the default name for the next FlashCopy of Accounting is still *Data 1-2*.
- 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 continue.

ii. The next screen displayed is the Specify Repository Logical Drive Capacity. On the Specify FlashCopy Repository Logical Drive's Capacity window (Figure 7-13 on page 209), set the repository drive capacity as a percentage of the base logical drive's capacity. This defaults to 20 percent.

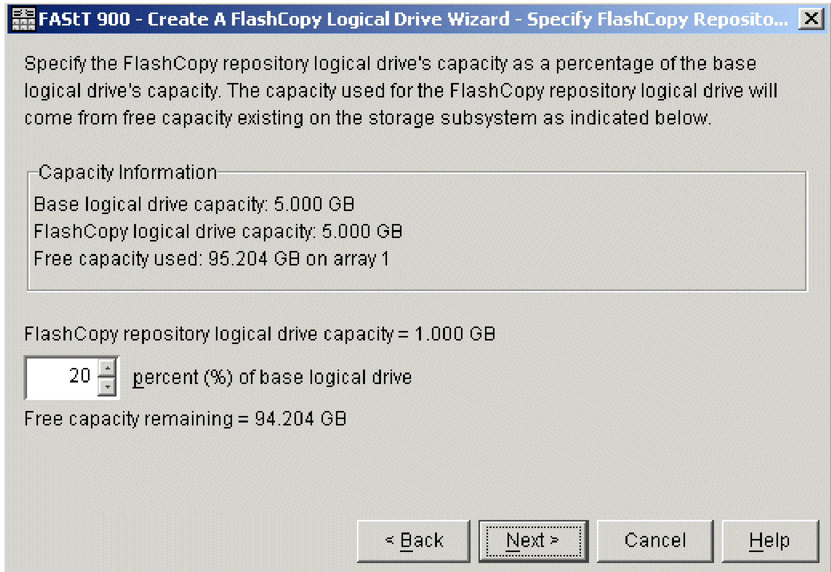

Figure 7-13 Specify repository logical drive capacity

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. Click **Next** to continue.

> **Important:** In most situations, the 20 percent default value should be ample capacity for your FlashCopy repository logical drive. For information on determining the size, see “Estimating FlashCopy repository logical drive capacity” on page 198 and “Estimating FlashCopy repository life” on page 200.

iii. The Create FlashCopy Wizard Preview screen appears (Figure 7-14 on page 210). It displays components associated with the FlashCopy. Review the information and click **Finish**.

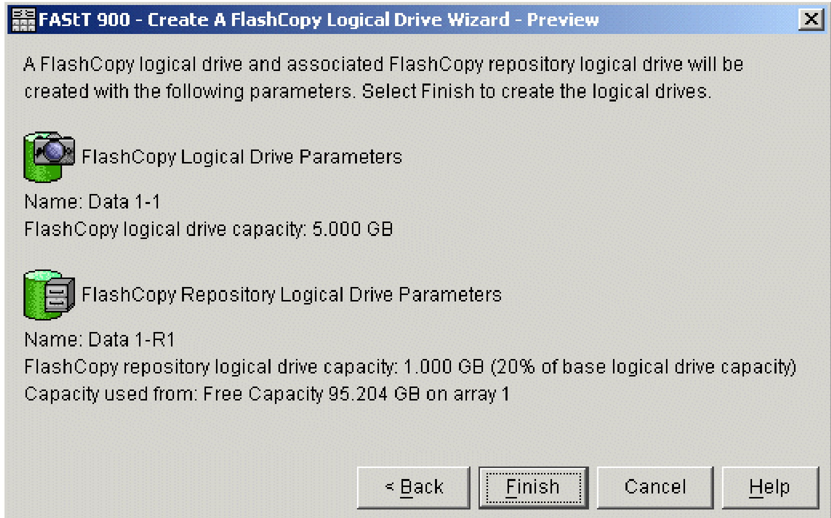

Figure 7-14 Create FlashCopy Logical Drive Wizard - Preview

iv. The Wizard then displays a complete message (Figure 7-15). Click **OK** to continue.

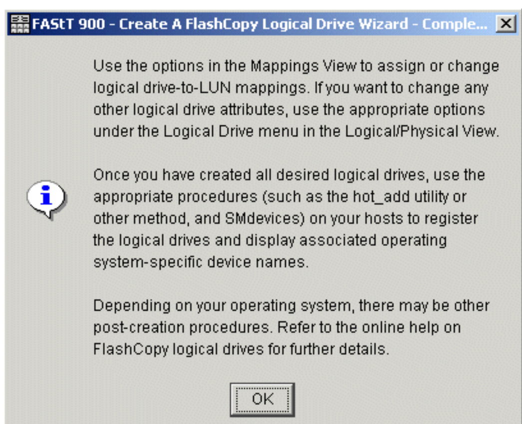

Figure 7-15 Create FlashCopy Logical Drive - Complete

v. You can now view the newly created FlashCopy logical drive and its associated repository logical drive from the Subsystem Management window. See Figure 7-16 on page 211.

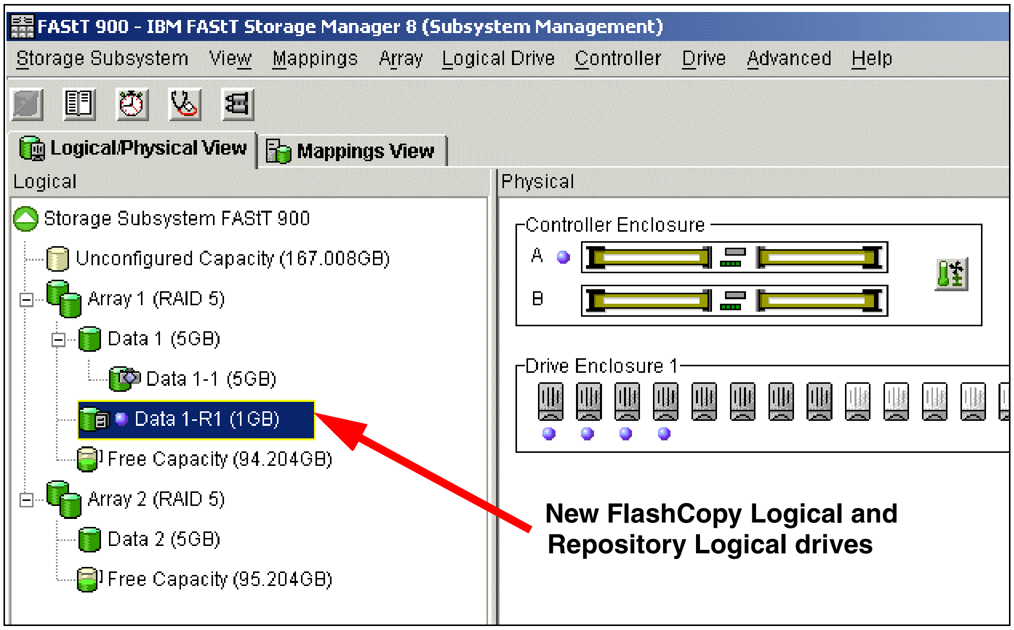

Figure 7-16 New FlashCopy volume in subsystem management

c. If you chose the advanced path, select **Advanced** from the Create FC Logical Drive Wizard Introduction window (Figure 7-17) and click **Next.**

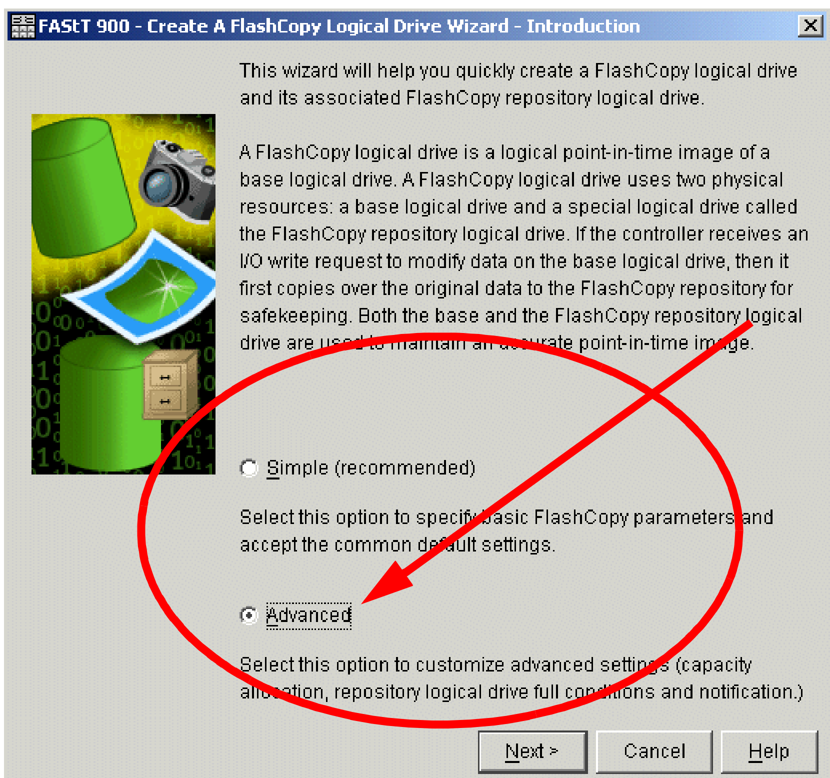

Figure 7-17 Create FlashCopy logical drive - Advanced

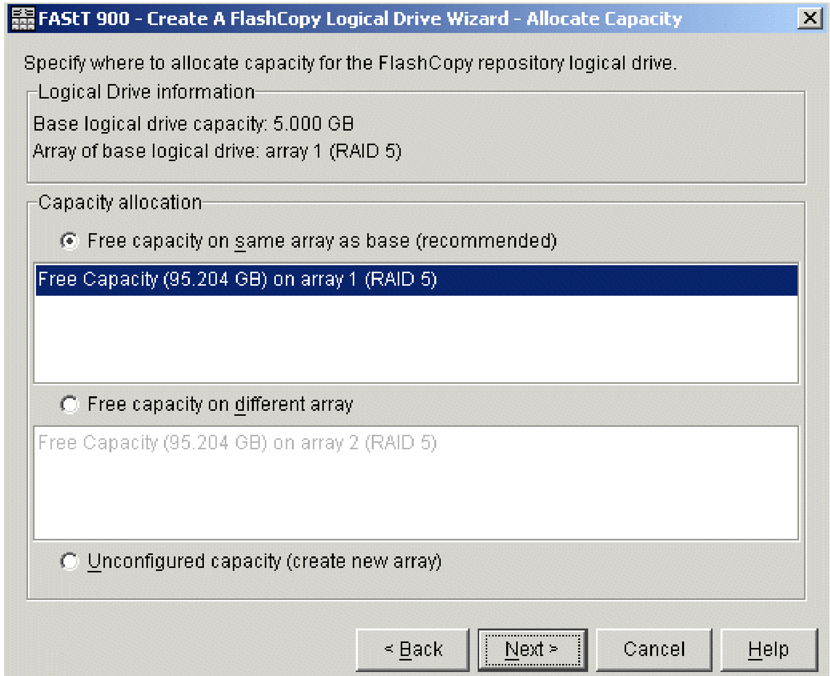

Figure 7-18 Create FlashCopy Logical Drive Wizard - Allocate Capacity

The Allocate Capacity window (Figure 7-18) appears on which you can choose free capacity on the *same array* as the base, free capacity on a *different array*, or unconfigured capacity to create a *new array*.

- Free Capacity on the same array

  If you select Free Capacity on the same array as the base logical drive and click **Next**, The Specify Names panel opens as seen previously in Figure 7-12 on page 208. Check or specify new names for the associated volumes then click **Next**. Highlight the free capacity on the same array and click **Next** to continue. The Specify Repository Logical Drive Capacity window appears as seen in Figure 7-13 on page 209. Make any modifications and click **Next**. The Specify Logical Drive Parameters window (Figure 7-19 on page 213) opens. Define the FlashCopy logical drive-to-LUN mapping parameter and the FlashCopy repository logical drive full conditions. Click the **Next** button.

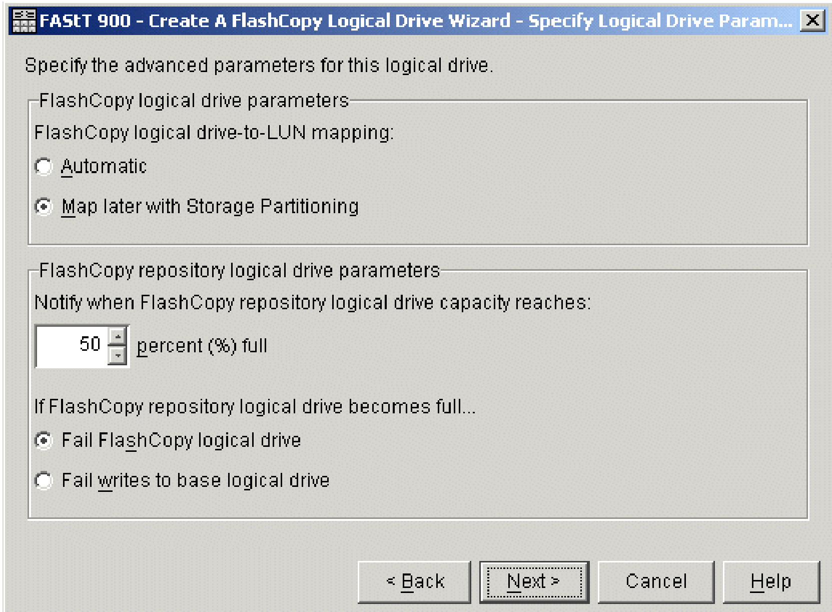

Figure 7-19 Create FlashCopy logical drive - Specify Parameters

The Preview window (Figure 7-20) will open and show 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. Note there is more information regarding logical drive parameters than was displayed in the Preview window when the simple path was used.

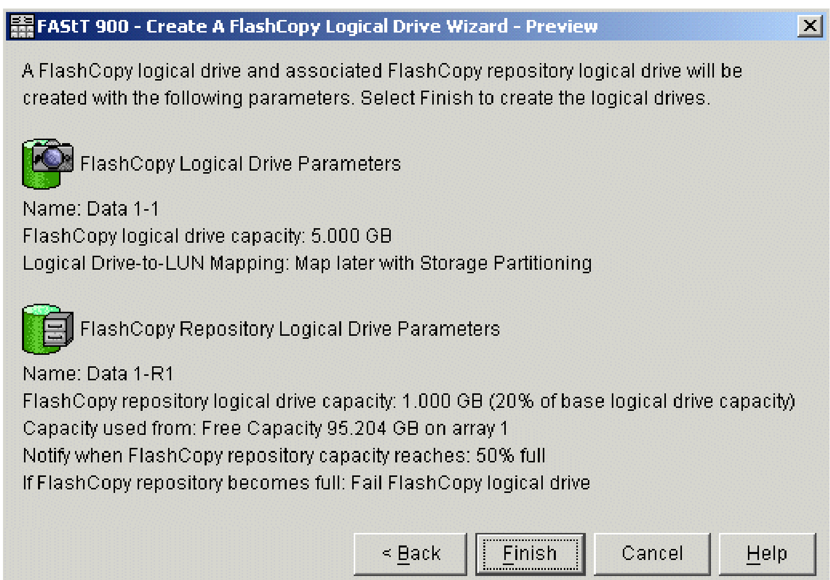

Figure 7-20 Create FlashCopy Logical Drive Wizard - Preview

Click **Finish**, and the Create FlashCopy Wizard Complete message box is displayed, as seen in Figure 7-15 on page 210. Read the notice and click **OK** to continue.

- Free capacity on a different array

  On the Allocate Capacity window (Figure 7-18 on page 212) you may select **Free capacity on different array** capacity and click **Next**. Click **OK** at the startup message, select **Advanced** at the Introduction window, and click **Next**.

  The Specify names window appears as per Figure 7-12 on page 208. Choose names for the FlashCopy repository and logical drives or accept the defaults and select.

  The Select Free Capacity window appears (Figure 7-21). In this window you will select the **Free capacity on a different array** button, then highlight from the available arrays where you want the repository logical drive to be located. Click **Next** to proceed.

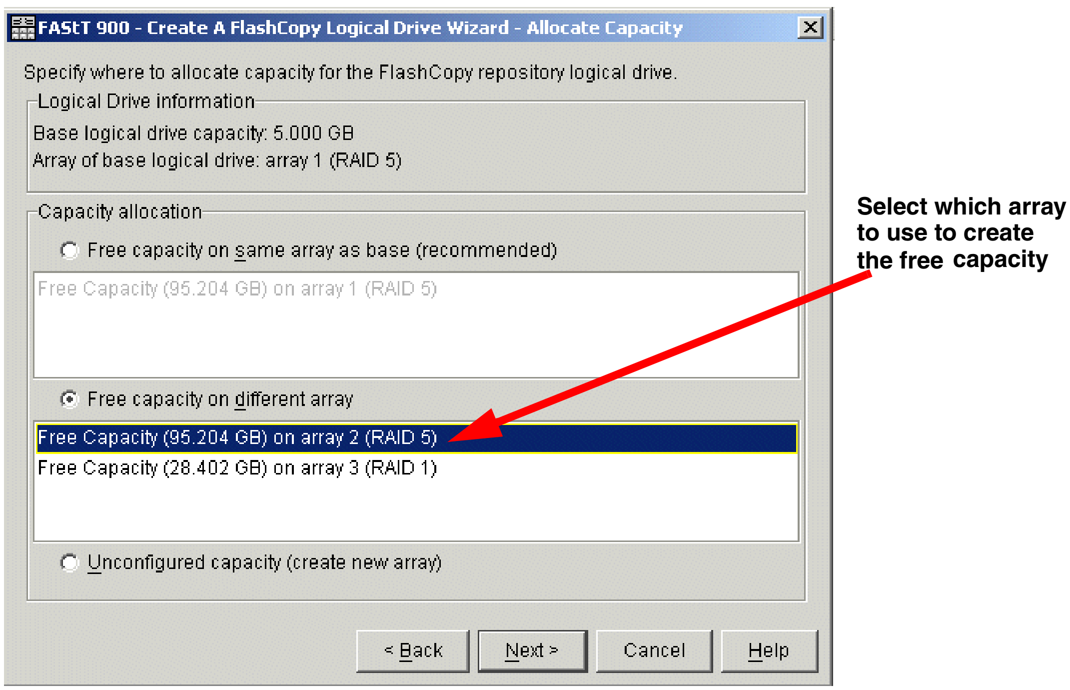

Figure 7-21 Create FlashCopy Logical Drive Wizard - Allocate capacity - Different array

The next window is the Specify FlashCopy Repository Logical Drive Capacity as was seen in Figure 7-13 on page 209. Choose the percentage of Base Logical as the size for the repository logical drive or accept the default of 20 percent. Click **Next** to proceed to the Specify FlashCopy Repository Logical Drive Parameters window as was seen in Figure 7-19 on page 213. Select the appropriate parameters, click **Next** to proceed to the Preview window as per Figure 7-20 on page 213. Click **Finish** and then **OK** at the Complete message.

Now when you view the FlashCopy Repository Logical Drives you have just created in the Subsystem Management window, you will observe it is in a different array to the base logical drive (Figure 7-22 on page 215).

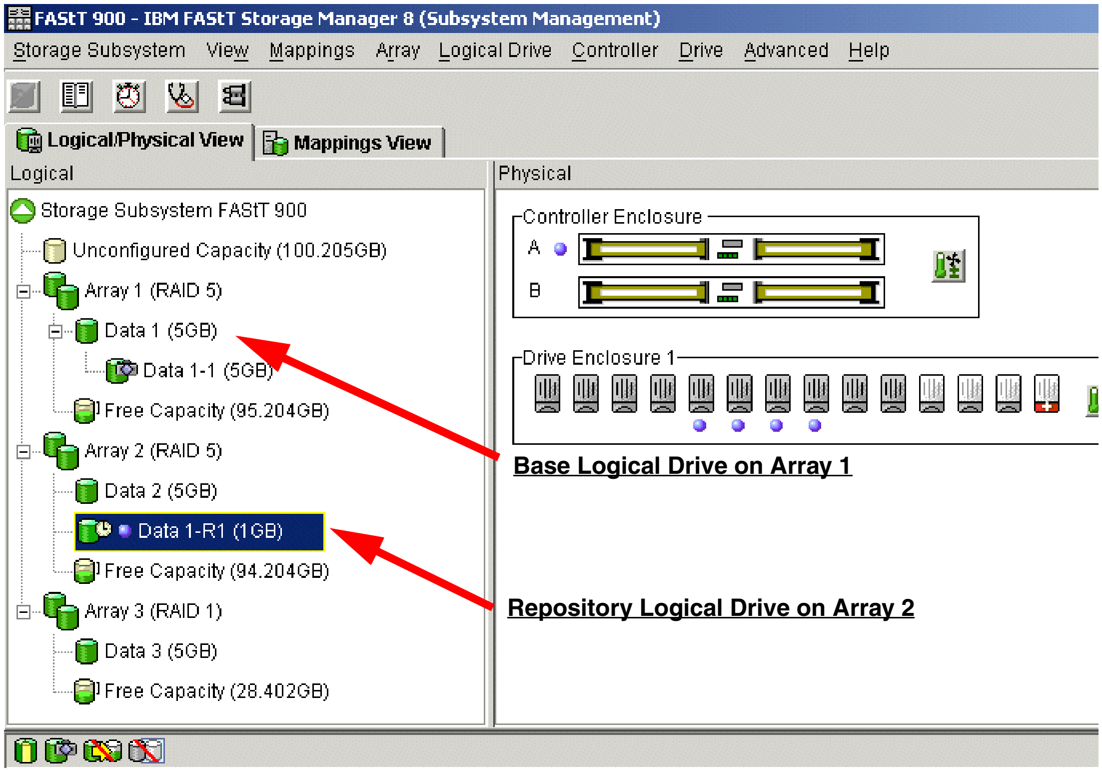

Figure 7-22 Repository logical drive on different array

- Unconfigured capacity

  Right-click the base logical drive from the Subsystem Management window, and select **Create a FlashCopy Logical Drive**, then OK the message and select **Advanced** from the Introduction window. The Allocate Capacity window, (Figure 7-23). Select **Unconfigured Capacity** and click **Next**.

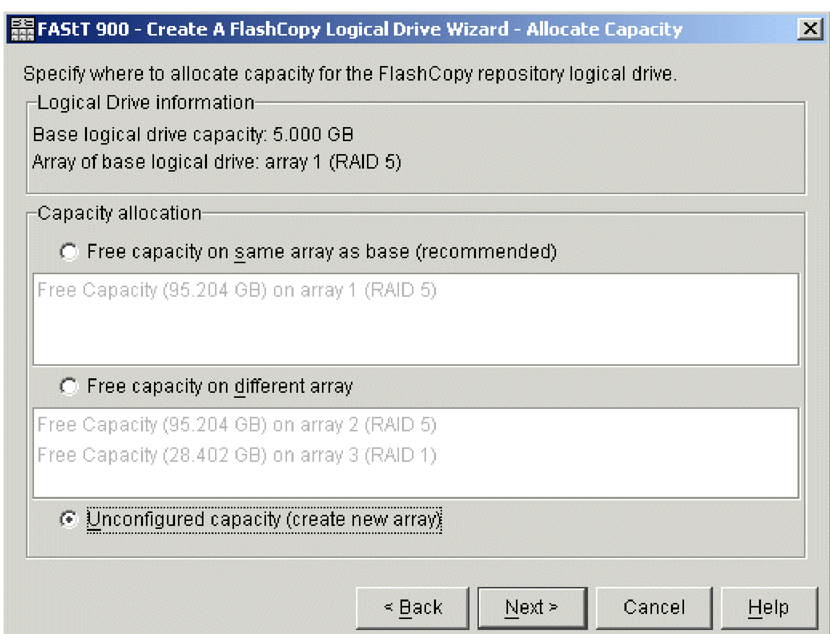

Figure 7-23 Create FlashCopy logical drive - Allocate capacity - Unconfigured capacity

The Create FlashCopy Wizard will then open the Specify Array Parameters window (Figure 7-24). Specify a new array where the repository will reside. Specify the RAID level of the array that meets the FlashCopy repository logical drive data storage and protection requirements. Click **Next** to continue.

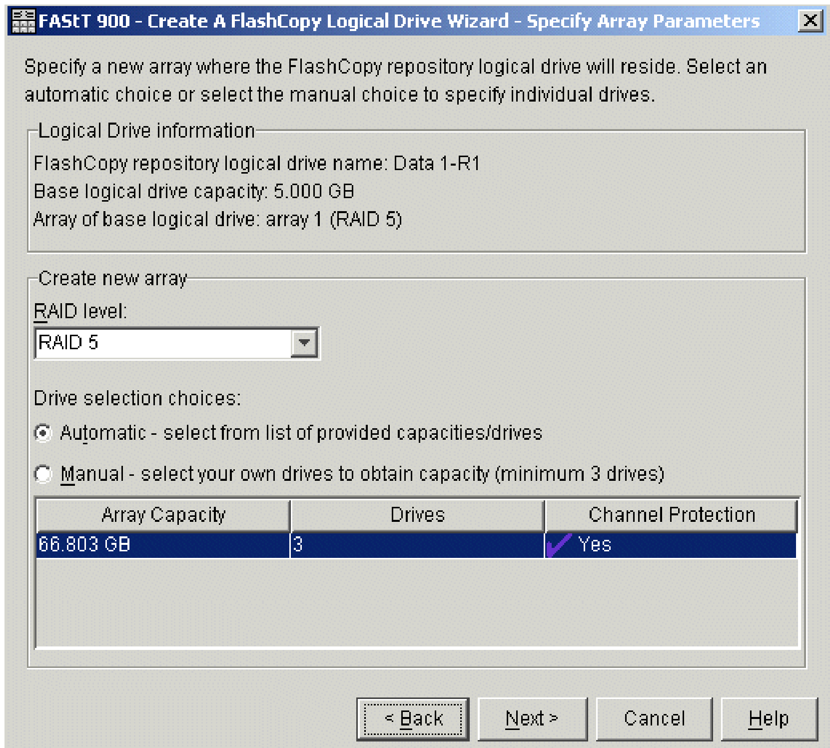

Figure 7-24 Create FlashCopy Logical Drive Wizard - Specify Array Parameters

The Wizard will take you to the Specify Repository Logical Drive Capacity window (Figure 7-25).

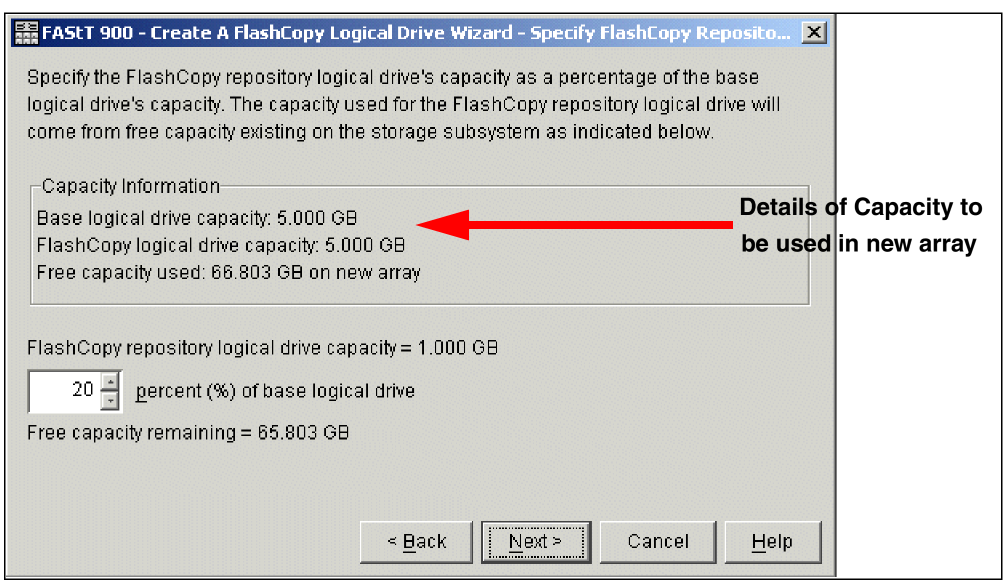

Figure 7-25 Create FlashCopy Logical Drive - Specify Repository Capacity - New Array

The Wizard then proceeds through the Specify Logical Drive Parameters window (Figure 7-19 on page 213). The Preview window and Complete messages are displayed, before returning to the Subsystem Management window.

You can observe the Repository Logical Drive just created in a new array (Figure 7-26).

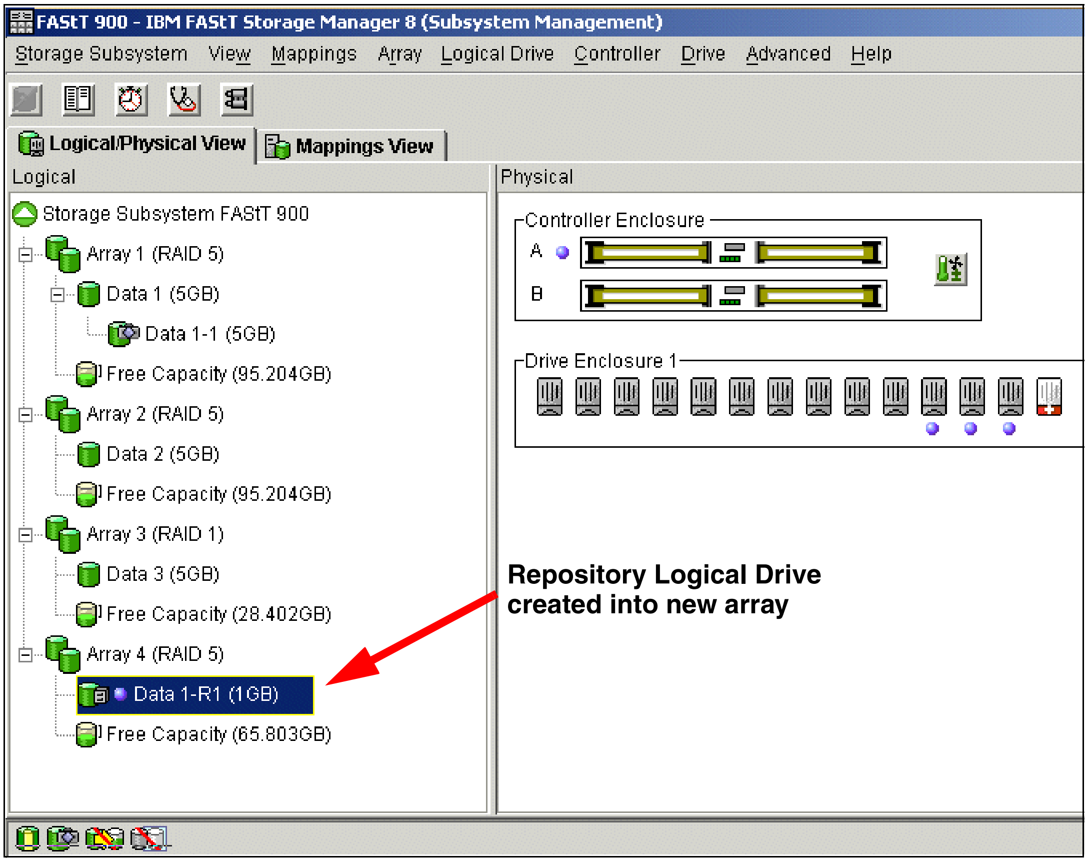

Figure 7-26 Repository Logical Drive on new array

4. Restart the host application.

   After you have created one or more FlashCopy logical drives, mount the base logical drive and restart the host application using that base logical drive.

7.2.4 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 351.

To map the FlashCopy logical drive to a host, follow these instructions:

1. Open the Mappings View of the Subsystem Management window. The newly created FlashCopy logical drive appears in the undefined mapping section (Figure 7-27 on page 218).

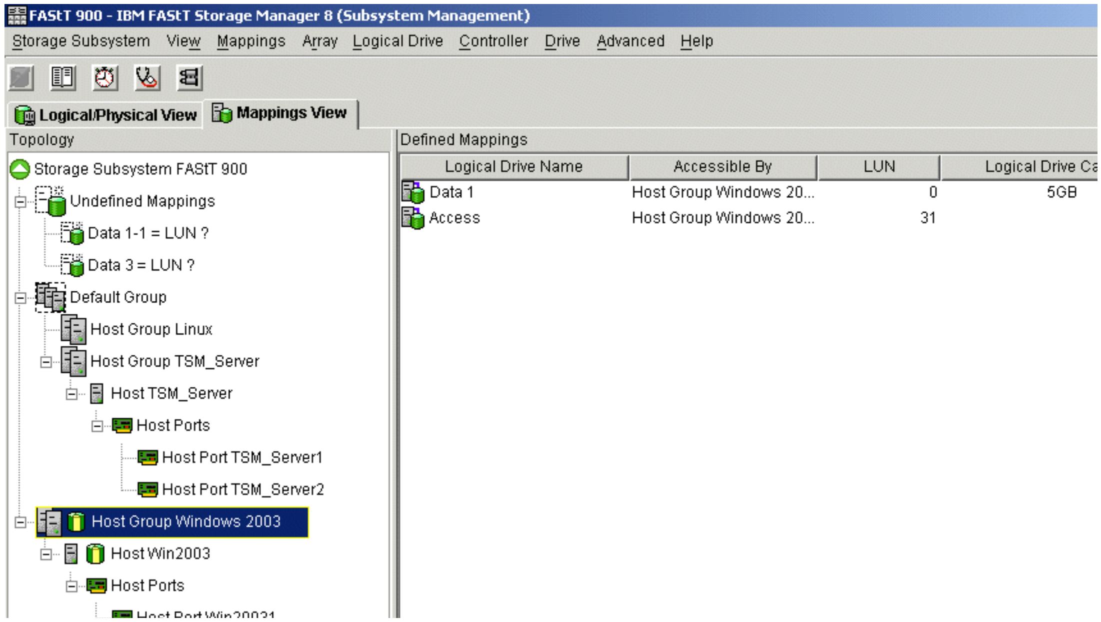

Figure 7-27 Subsystem Management window

2. Select the host or host group you want to map the drive to.
3. Right-click and select **Define Additional Mapping**, as shown in Figure 7-28.

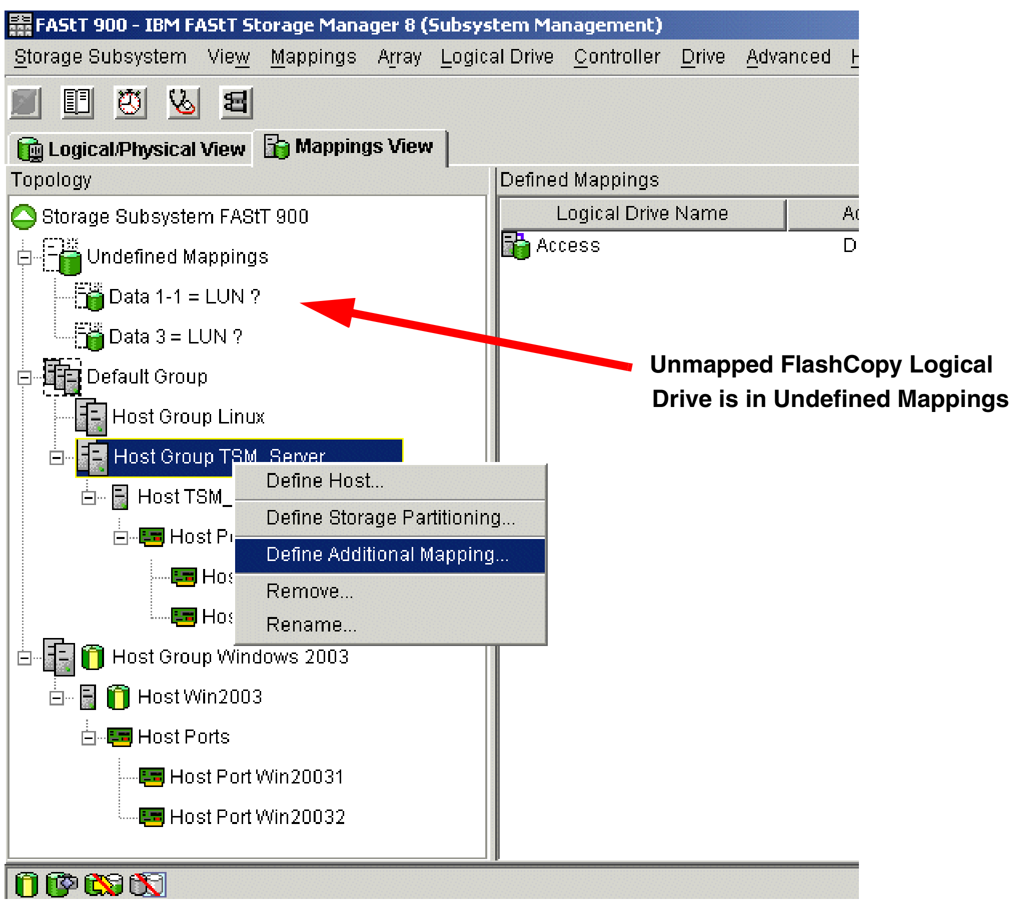

Figure 7-28 Define Additional Mapping from Subsystem Management

4. The Define Additional Mapping window (Figure 7-29) opens. Follow these steps:
   a. Select the FlashCopy drive.
   b. Select the host or host group.
   c. Set the LUN number.
   d. Click **Add**.

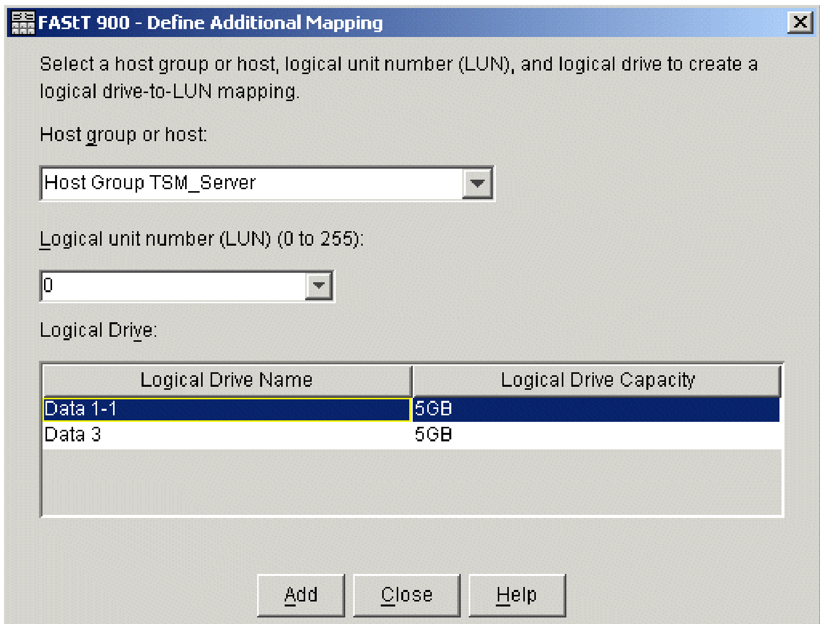

Figure 7-29 Define Additional Mapping

In the examples shown in Figure 7-28 on page 218 and Figure 7-29, the FlashCopy logical disk “Data 1-1” is made available to the TSM_Server host group. This enables the TSM_Server to access and mount the disk as part of its own file system.

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 7-30 on page 220.

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 “Disabling and recreating a FlashCopy logical drive” on page 195.

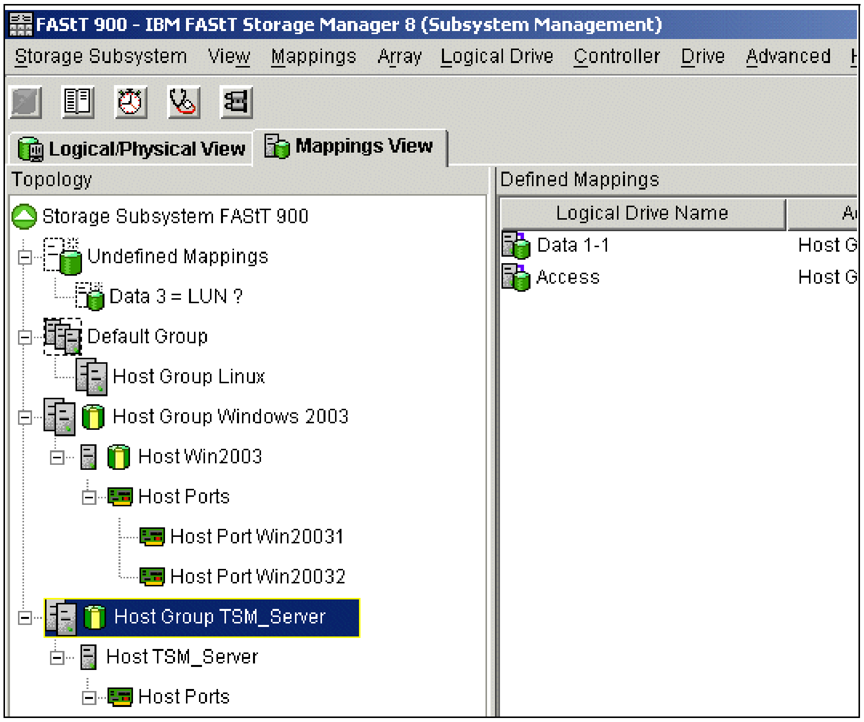

Figure 7-30 Mapped FlashCopy Logical Drive

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 351.

7.2.5 Viewing the FlashCopy drive status

The status of the FlashCopy logical drive can be determined by viewing the icons that will change depending on 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 7-31 on page 221.

FlashCopy status representation in logical view		
Logical drive status icon	**Mirror status**	**Icon**
FlashCopy	Optimal	
	Disabled	
	Failed	
	Offline	
Repository	Optimal	
	Degraded	
	Failed	
	Full	
	Offline	
	Warning	
Free capacity	Spare capacity	

Figure 7-31 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 7-32.

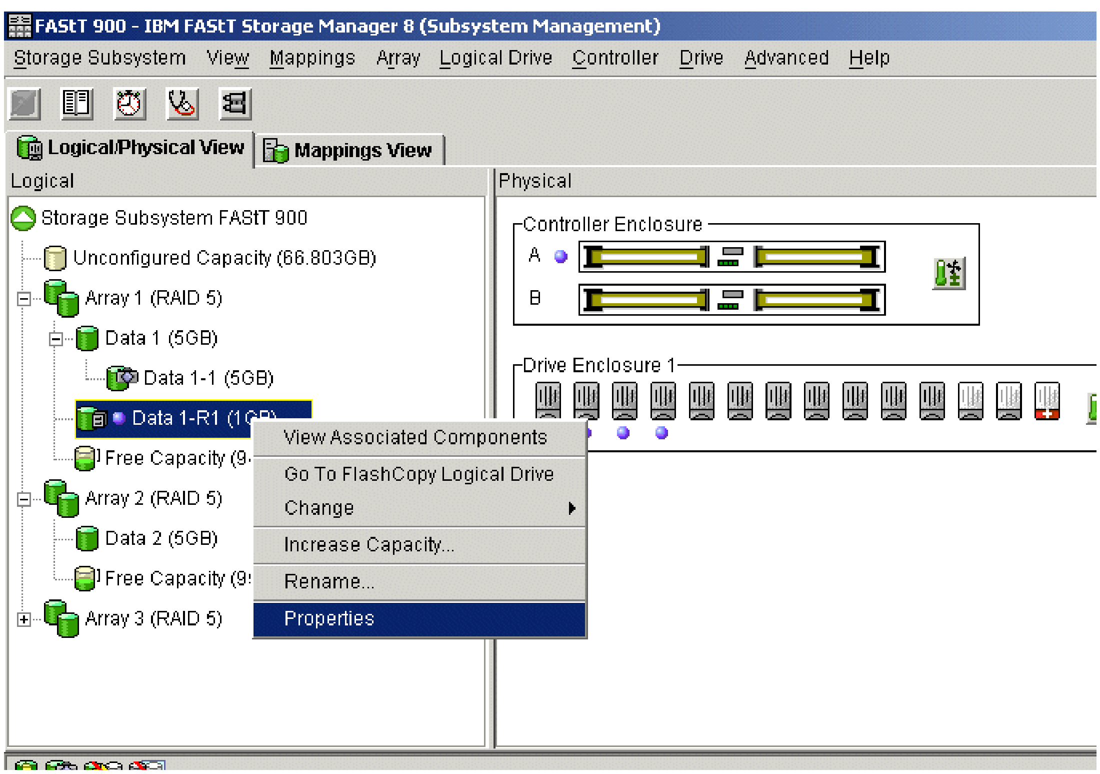

Figure 7-32 Repository Logical Drive properties

The FlashCopy Repository Logical Drive - Properties window opens as shown in Figure 7-33.

Figure 7-33 FlashCopy Repository Logical Drive Properties - Base tab

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 7-34 on page 223) 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 percent of the FlashCopy repository logical drives maximum capacity.

     Use the percent (%) full box to define the percentage at which a warning is issued.

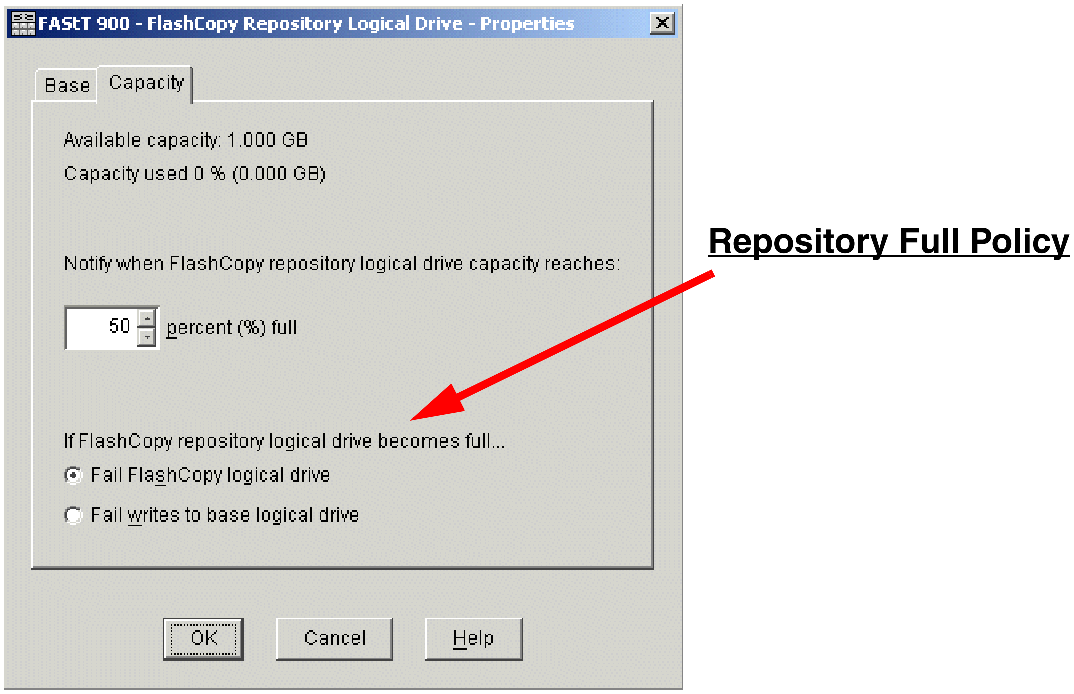

Figure 7-34 FlashCopy Repository Logical Drive Properties - Capacity tab

– 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 7-34) 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 7-34), 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 “Increasing the capacity of a FlashCopy repository logical drive” on page 201.

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 (Figure 7-35 on page 224).

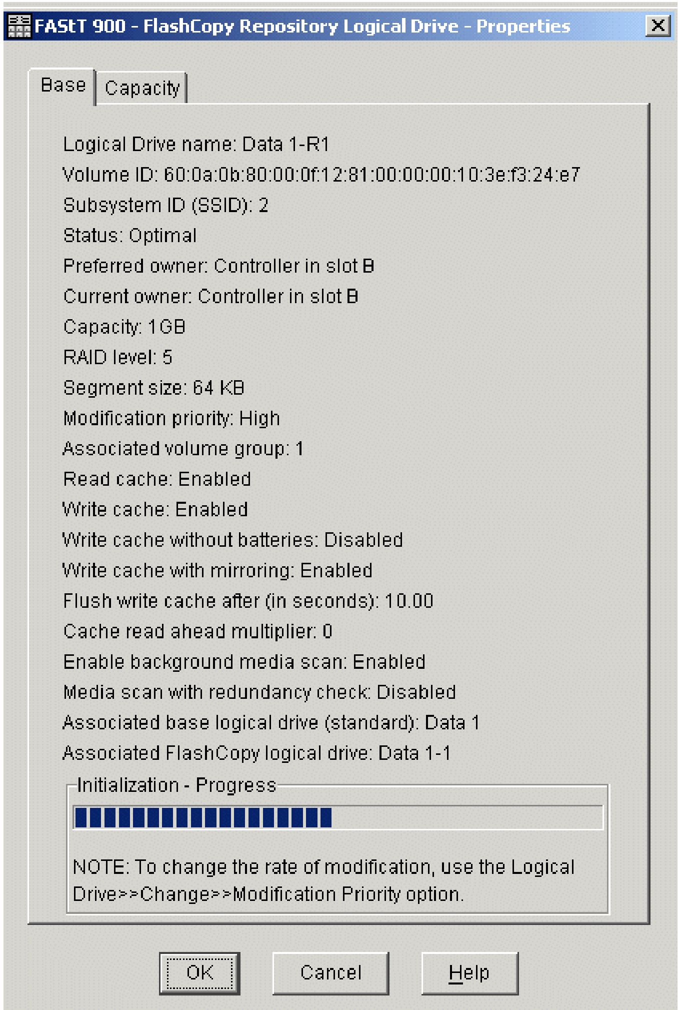

Figure 7-35 FlashCopy Repository Logical Drive Properties - Progress

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.

7.2.6 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 7-36.

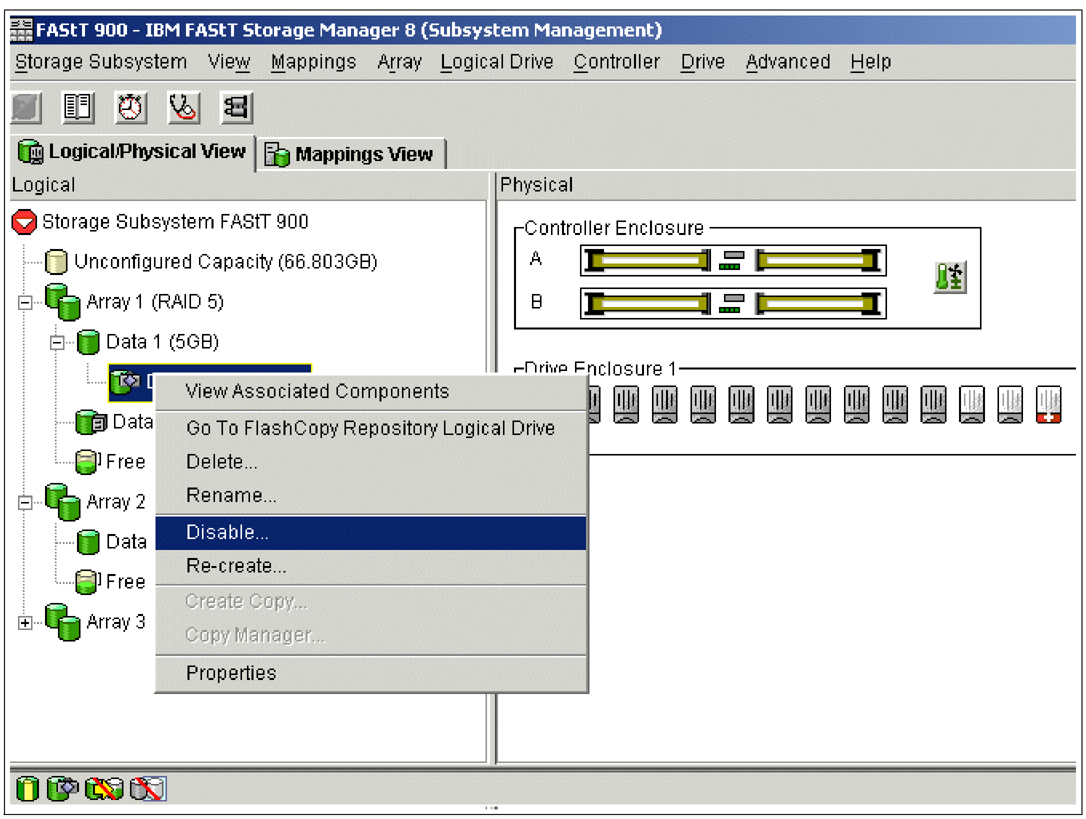

Figure 7-36 FlashCopy Logical Drive - Disable

2. The Disable FlashCopy Logical Drive confirmation window (Figure 7-37 on page 226) opens. On this window, type `yes` and click **OK** to begin the disable operation.

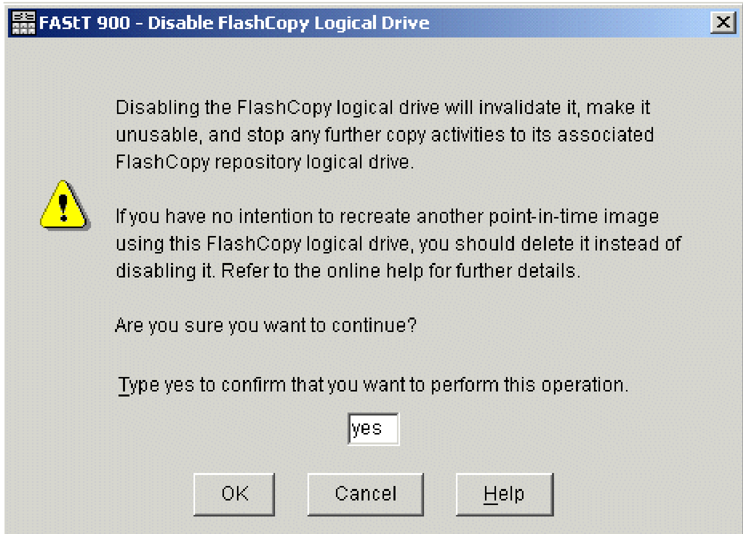

Figure 7-37 Disable FlashCopy confirmation

The FlashCopy icon in the Physical/Logical View now appears as disabled, as shown in Figure 7-38.

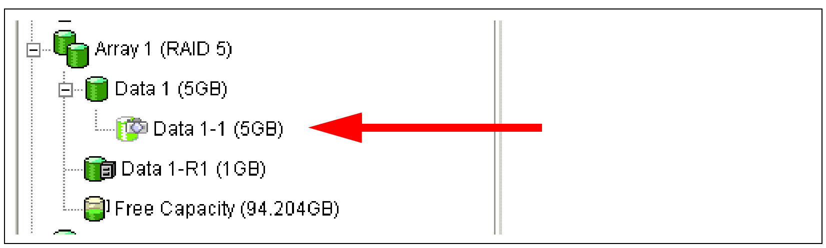

Figure 7-38 Icon showing the disabled FlashCopy logical drive

7.2.7 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 Logical Drive 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 Logical Drive 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 7-39.

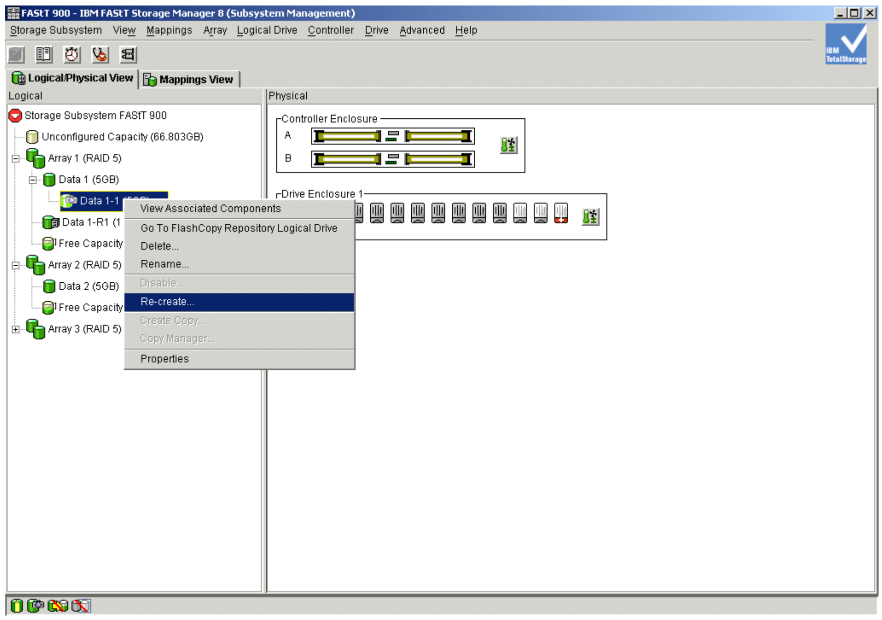

Figure 7-39 FlashCopy Logical Drive - Delete

2. The Re-create FlashCopy Logical Drive dialog opens (Figure 7-40). 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.

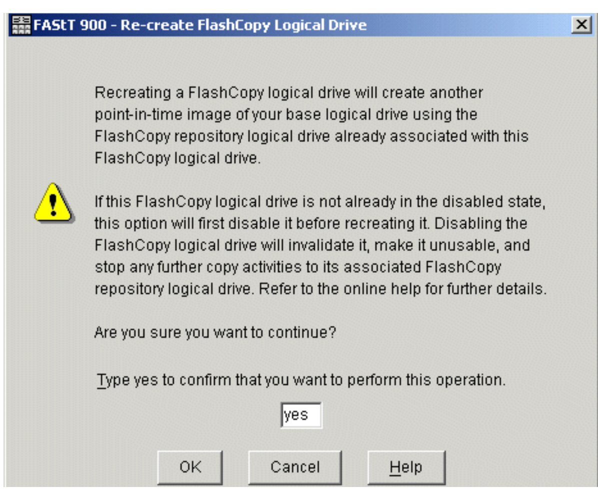

Figure 7-40 Re-creating a FlashCopy Logical Drive - Confirmation

Important: The following is 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, and creates a new point-in time copy

7.2.8 Resizing a FlashCopy Repository Logical 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 Logical 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** (Figure 7-41 on page 229).

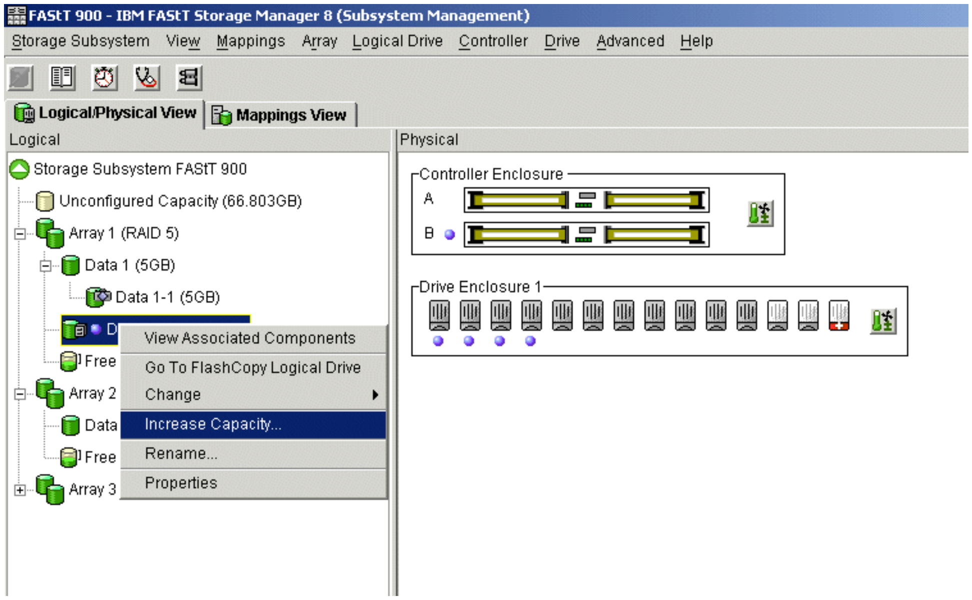

Figure 7-41 FlashCopy Repository Logical Drive - Increase Capacity

> **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 zero. 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 (Figure 7-42 on page 230).

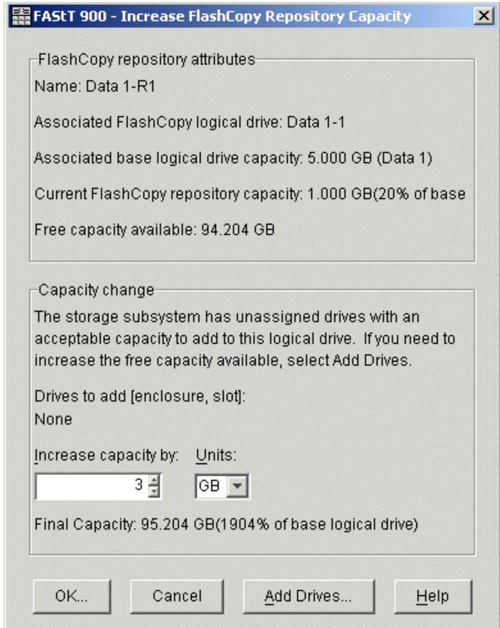

Figure 7-42 Increase FlashCopy Repository Capacity

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 (Figure 7-43).

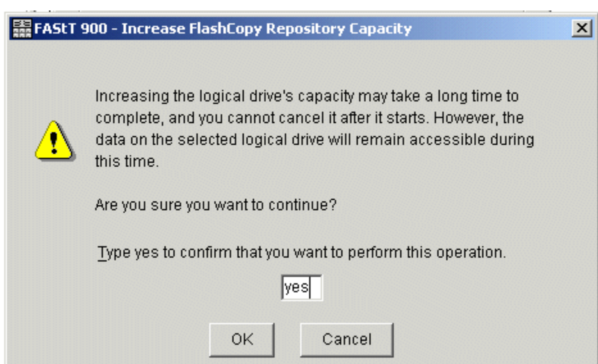

Figure 7-43 Increase FlashCopy Repository Capacity - Confirmation

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. At the Increase FlashCopy Repository Capacity window (Figure 7-42 on page 230), and select **Add Drives**.

The Increase Repository Capacity - Add Drives window opens. Enclosure, slot, and usable capacity details for the available free drives are displayed (Figure 7-44).

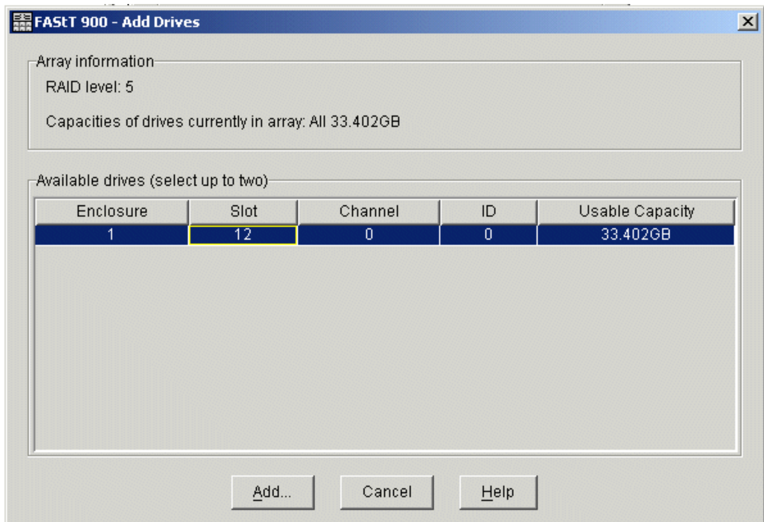

Figure 7-44 Add Drives

> **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 Drives window is closed. Check the “Drives to add [enclosure, slot]” area to ensure the correct drives are added (Figure 7-45 on page 232).

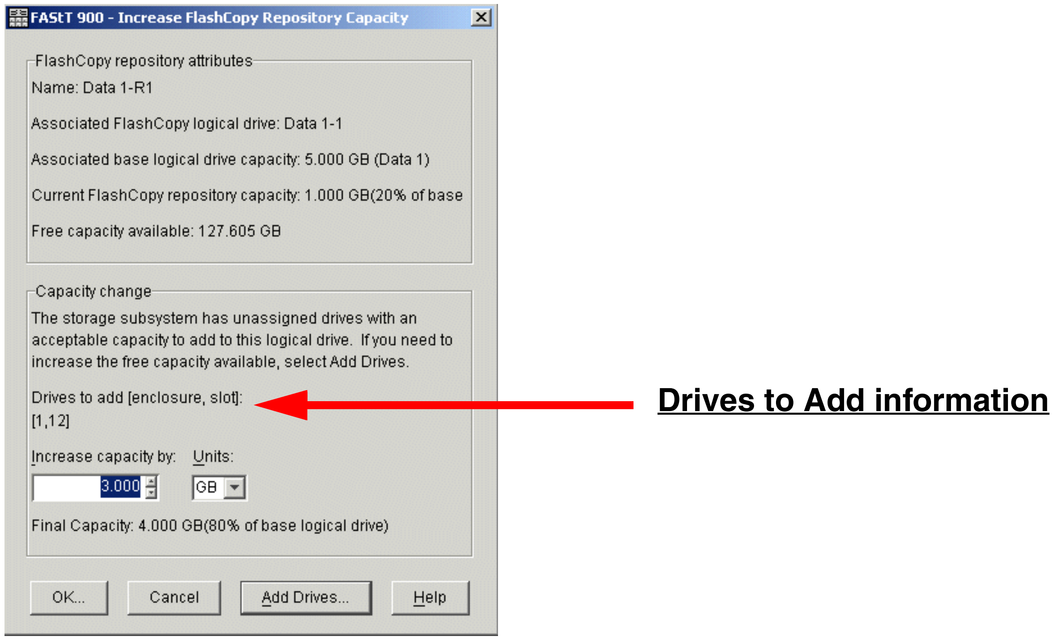

Figure 7-45 Increase FlashCopy Repository Capacity - Added drives

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 (Figure 7-46).

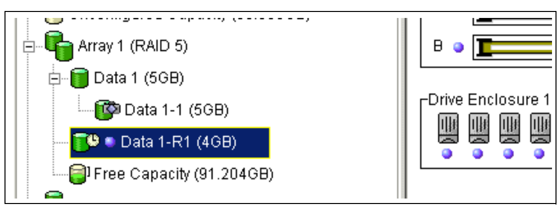

Figure 7-46 FlashCopy Repository Logical Drive with increased capacity

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 of 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.

7.2.9 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:** Read over the following points:
>
> - 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 Logical 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 7-47.

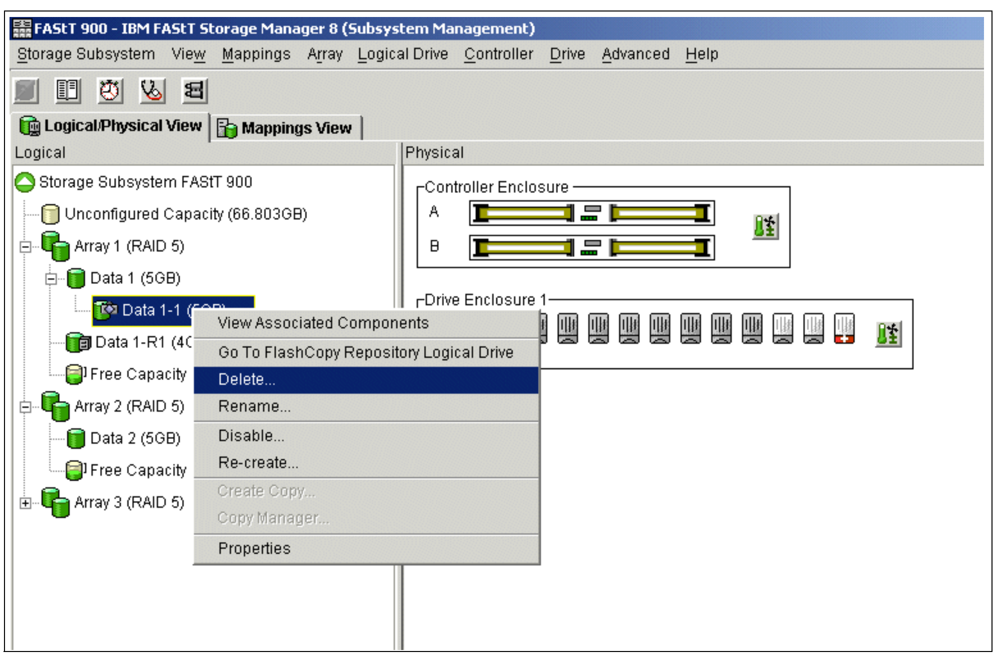

Figure 7-47 Deleting the FlashCopy Logical Drive

3. The Delete FlashCopy Logical Drive window opens as shown in Figure 7-48. Highlight the FlashCopy Logical Drive(s) you wish to delete and click **OK**. Be careful not to delete the wrong logical drive.

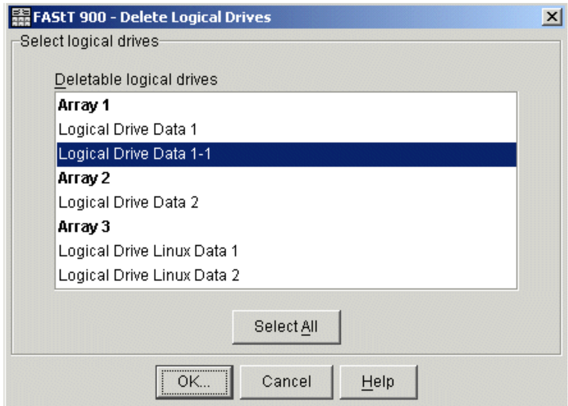

Figure 7-48 Delete FlashCopy Logical Drive dialog

4. The Confirm Delete Logical Drive(s) window opens (Figure 7-49).

   Read the warning message, type `yes`, and click **OK**.

   The FlashCopy logical drive and FlashCopy repository logical drive are deleted. Any data on them is destroyed.

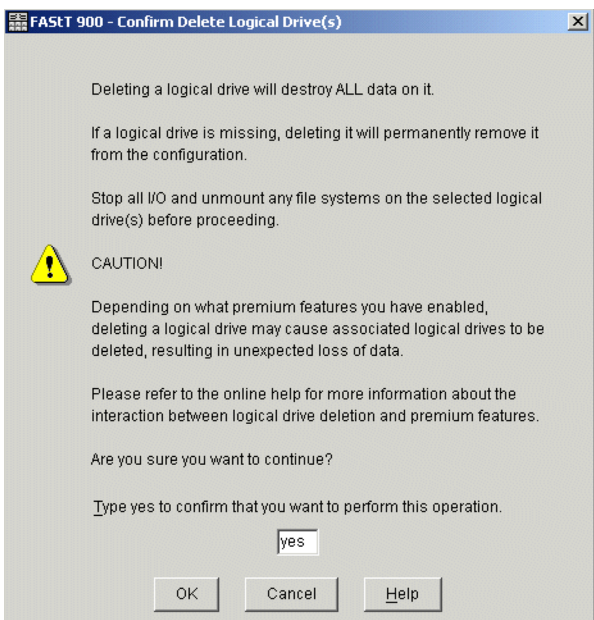

Figure 7-49 Confirm Delete Logical Drive(s)

7.3 Command Line Interface for FlashCopy

The Command Line Interface was introduced in “Script Editor and command line interface” on page 89.

This section discusses those commands that are specific to the FlashCopy feature.

For a more in-depth information on using the Command Line Interface and the Script Editor, refer to “Command Line Interface and Script Editor” on page 311.

The following is a list of Command Line Interface commands that are pertinent to FlashCopy:

- **`create flashCopyLogicalDrive`**
- **`set LogicalDrive`**
- **`disableFlashCopy LogicalDrive`**
- **`recreateFlashCopy LogicalDrive`**
- **`delete logicalDrive`**

7.3.1 Create FlashCopyLogicalDrive - CLI

You can create a FlashCopy from a standard logical drive using the Command Line Interface in one of two ways.

- The *simple* method: The result is similar as to that described using the GUI “Creating a FlashCopy drive” on page 204. All that is specified is the base logical drive from which a FlashCopy is created. The resulting FlashCopy logical drive and repository logical drive is created in the same array as the base logical drive using the default naming convention. The repository logical drive will have the default size (20 percent of base), the default warning threshold (50 percent full), and default repository full policy (fail the FlashCopy logical drive).

  See “Example 1” on page 235 for a demonstration of this command.

- The *advanced* method: Using this method you can specify:
  - The FlashCopy logical drive name
  - The repository logical drive name
  - Which of the existing arrays you would like to place the repository logical drive on
  - Specify to place the repository logical drive on an unconfigured array, the RAID level of that array, which drives to use for that array
  - The size of the repository logical drive
  - The warning threshold of the repository logical drive
  - The repository full policy

  Any parameter not specified will result in the default value being used. See “Example 2” on page 237 for a demonstration of this method.

Create FlashCopyLogicalDrive examples

Example 1

In this example we create a FlashCopy of logical drive Data 1 in array 2 (See Figure 7-50 on page 236). The simple method is used so that *all* default values are applied.

We use the following command (for more information on using CLI and the correct syntax refer to “Command Line Interface and Script Editor” on page 311):

```
SMcli 10.1.1.1 -c "create FlashCopyLogicalDrive baseLogicalDrive=\"Data 1\";"
```

Where 10.1.1.1 is the IP address of the storage subsystem that is targeted, and “Data 1” is the base logical drive from which a flash copy is to be made from.

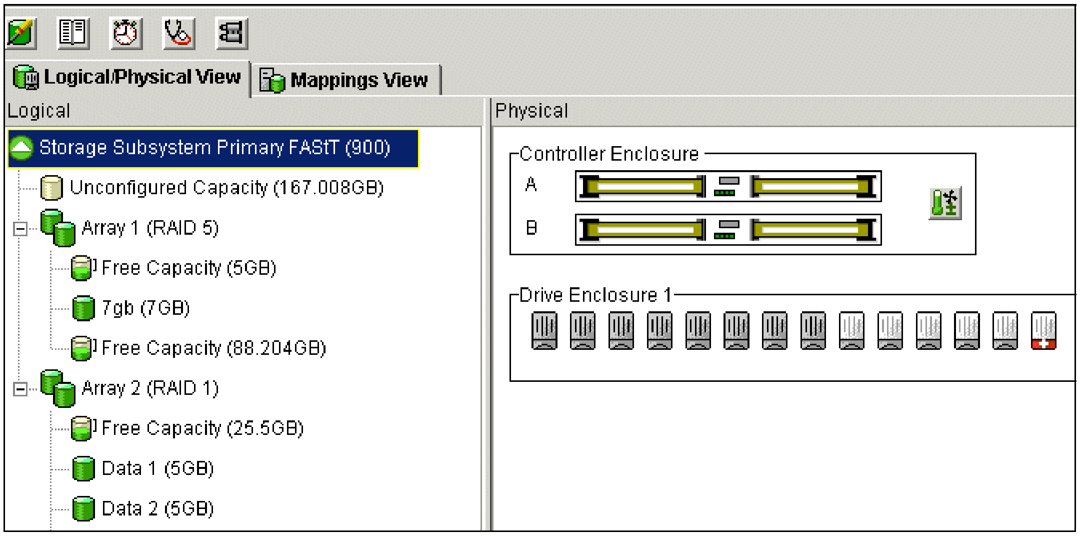

Figure 7-50 Logical drives on array 2

When the command is executed, the syntax is checked and if correct the script is executed. A completion status is given.

A check of the GUI (Figure 7-51) shows that “Data 1-1” FlashCopy logical drive has been created as a child node of “Data 1”. Also the repository logical drive “Data 1-R1” has been created in Array 2, with a default size of 20 percent of the base logical drive.

A check of the repository logical drive properties (Figure 7-52 on page 237) shows that the default warning threshold of 50 percent full and the default repository full policy of fail FlashCopy logical drive have been set.

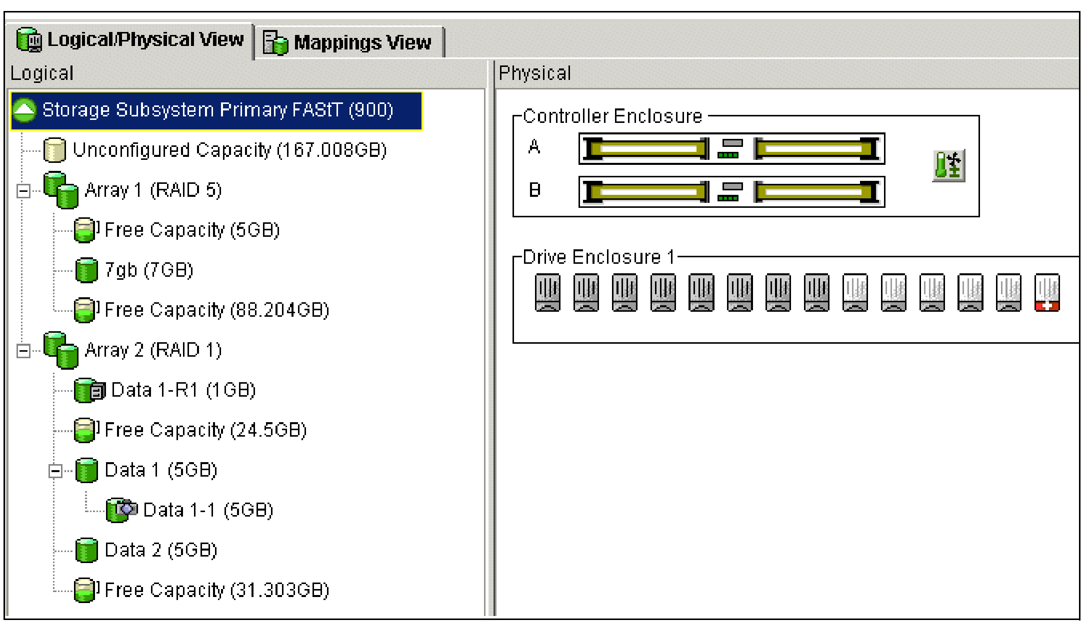

Figure 7-51 Data 1-1 created

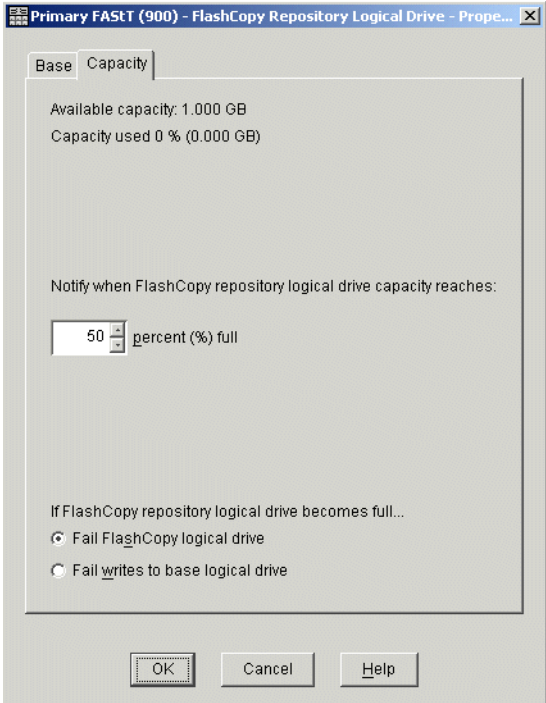

Figure 7-52 Repository logical drive - Properties

Example 2

In this example we create a FlashCopy of logical drive Data 2 in array 2 (See Figure 7-50 on page 236). The advanced method is used to specify the following:

- FlashCopy logical drive name = Flash 2
- Repository logical drive name = Repos 2
- Repository logical drive will be created on:
  - A new array
  - Using three disk drives
  - RAID level 5
  - Be 40 percent of the base size
  - Will warn when it is 75 percent full
  - Will fail writes to the base logical drive

To achieve this we enter the following command:

```
SMcli 10.1.1.1 -c "create FlashCopyLogicalDrive baseLogicalDrive=\"Data 2\"
repositoryDriveCount=3 userLabel=\"Flash 2\" repositoryUserLabel=\"Repos 2\"
repositoryRaidLevel=5 repositoryPercentofBase=40 warningThresholdPercent=75
repositoryFullPolicy=failBaseWrites;"
```

As can be seen in Figure 7-53 on page 238, “Flash 2” FlashCopy logical drive has been created as a child node of “Data 2”. A new RAID 5 array has been created, “Array 3”, using drives 9,10, and 11 and the repository logical drive “Repos 2” has been created, with a size of 40 percent of the base logical drive.

A check of the repository logical drive properties shows that the warning threshold of 75 percent full and the repository full policy of fail writes to base logical drive have been set.

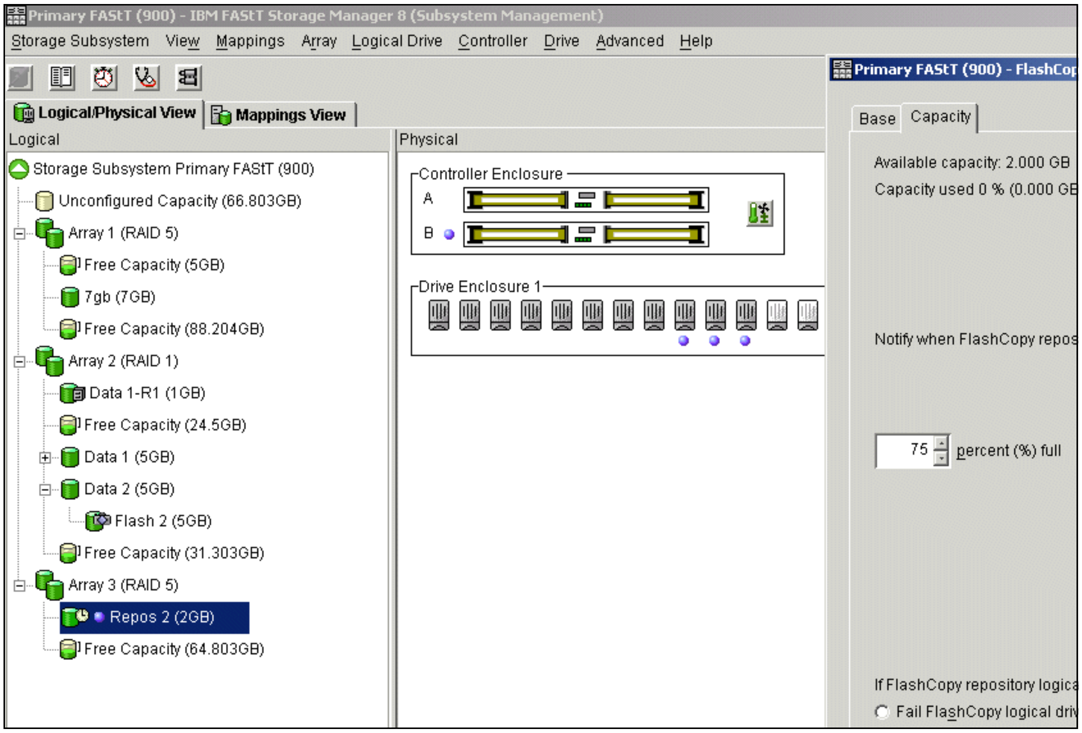

Figure 7-53 Flash 2 and Repos 2 created

7.3.2 Set LogicalDrive - CLI

The **`set logicalDrive`** command can be used to:

- Change a Base, FlashCopy, or repository logical drive name.
- Change a repository logical drive warning threshold.
- Change a repository logical drive full policy.

set logicalDrive examples

Example 1

In Example 1 the name of a FlashCopy logical drive is changed from “Flash 2” to “newname”. This achieved through the following command:

```
SMcli 10.1.1.2 -c "set logicalDrive [\"Flash 2\"] userLabel=\"newname\";"
```

The results are displayed in Figure 7-54 on page 239.

Example 2

In Example 2 the repository logical drive warning percentage has been changed to 25 percent and the repository logical drive full policy has been changed to fail the FlashCopy. This is achieved by using the following command.

```
SMcli 10.1.1.2 -c "set logicalDrive [\"Repos 2\"] warningthresholdPercent=25
repositoryFullpolicy=failFlashCopy;"
```

The results are displayed in Figure 7-54 on page 239.

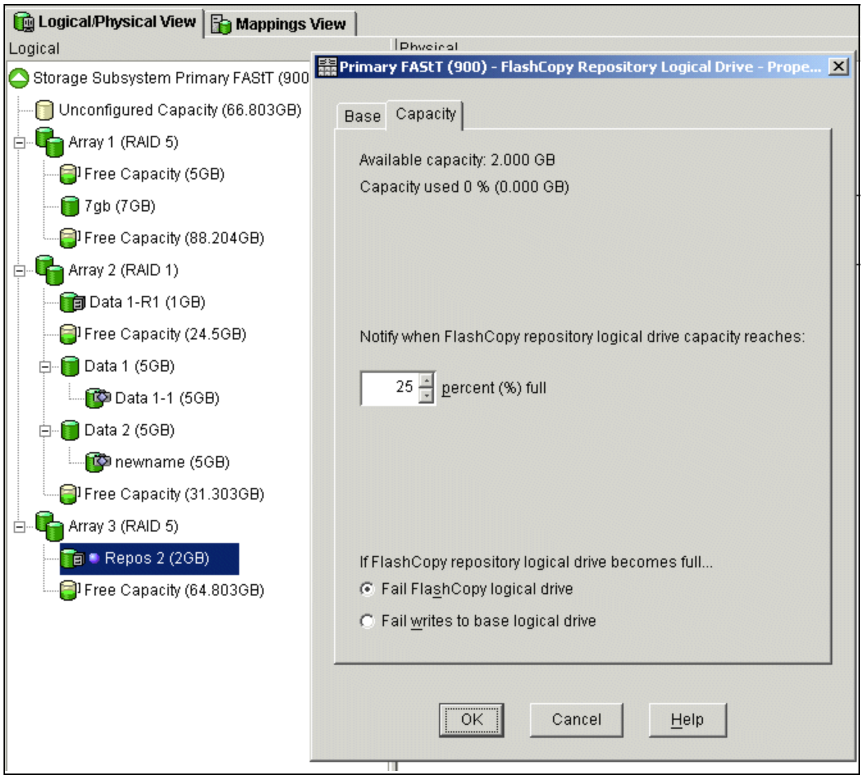

Figure 7-54 set logicalDrive examples

7.3.3 disableFlashCopy LogicalDrive - CLI

The **`disableFlashCopy LogicalDrive`** command is used to disable a FlashCopy as discussed in “Disabling and recreating a FlashCopy logical drive” on page 195.

Example

In this example, FlashCopy logical drive “newname” is disabled using the following command:

```
SMcli 10.1.1.2 -c "disableFlashCopy logicalDrive [\"newname\"];"
```

The result can be seen in Figure 7-55. The FlashCopy logical drive “newname” icon has become a light shade representing disabled.

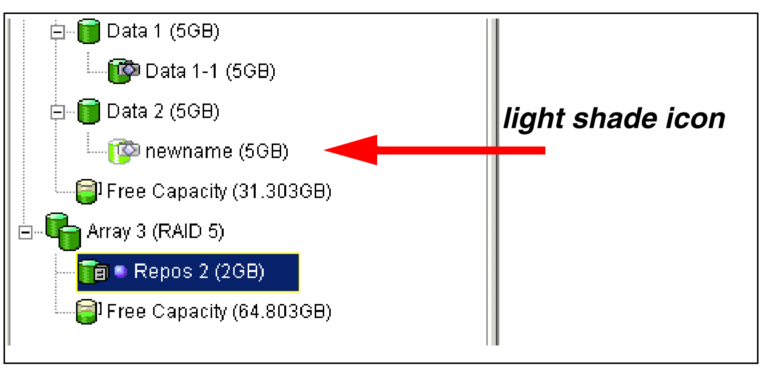

Figure 7-55 FlashCopy logical drive “newname” disabled

7.3.4 recreateFlashCopy LogicalDrive - CLI

The **disableFlashCopy LogicalDrive** command is used to recreate a FlashCopy as discussed in “Disabling and recreating a FlashCopy logical drive” on page 195.

Example

In this example we could just recreate the FlashCopy by specifying only the FlashCopy logical drive name. In this case, the new FlashCopy would be created using the original FlashCopy logical drive name and the repository logical drive parameters. However, using the CLI, we have the ability to recreate the FlashCopy logical drive, renaming it to “Flash2”, with repository warning threshold of 50 percent and a repository full policy of fail writes to base logical drive, with the following command:

```
SMcli 10.1.1.2 -c "recreateFlashCopy logicalDrive[\"newname\"] userLabel=\"Flash 2\"
warningThresholdPercent=50 repositoryFullPolicy=failBaseWrites;"
```

The results are shown in Figure 7-56.

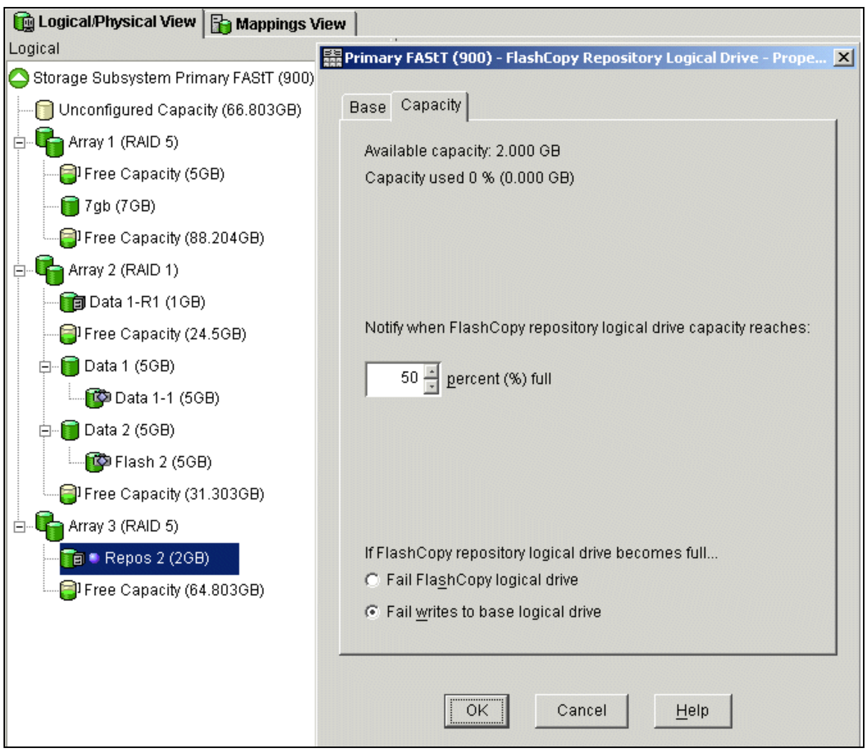

Figure 7-56 recreateFlashCopy LogicalDrive

7.3.5 delete logicalDrive

A FlashCopy logical drive can be deleted using this command. When deleting the FlashCopy logical drive the associated Repository logical drive is also deleted. Below is an example of this command.

Example

Here we provide an example.

```
SMcli 10.1.1.2 -c "delete logicalDrive [\"Flash 2\"];"
```

8

Remote Volume Mirroring

We have introduced Remote Volume Mirroring in “Remote Volume Mirroring (RVM) overview” on page 105. This chapter describes in detail the concepts of RVM, 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 can be used as a disaster recovery option, providing protection for your data 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 different locations.

8.1 Introduction to the Remote Volume Mirror option

The Remote Volume Mirror option is a premium feature that comes with the IBM FAStT Storage Manager 8.4 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 (Figure 8-1).

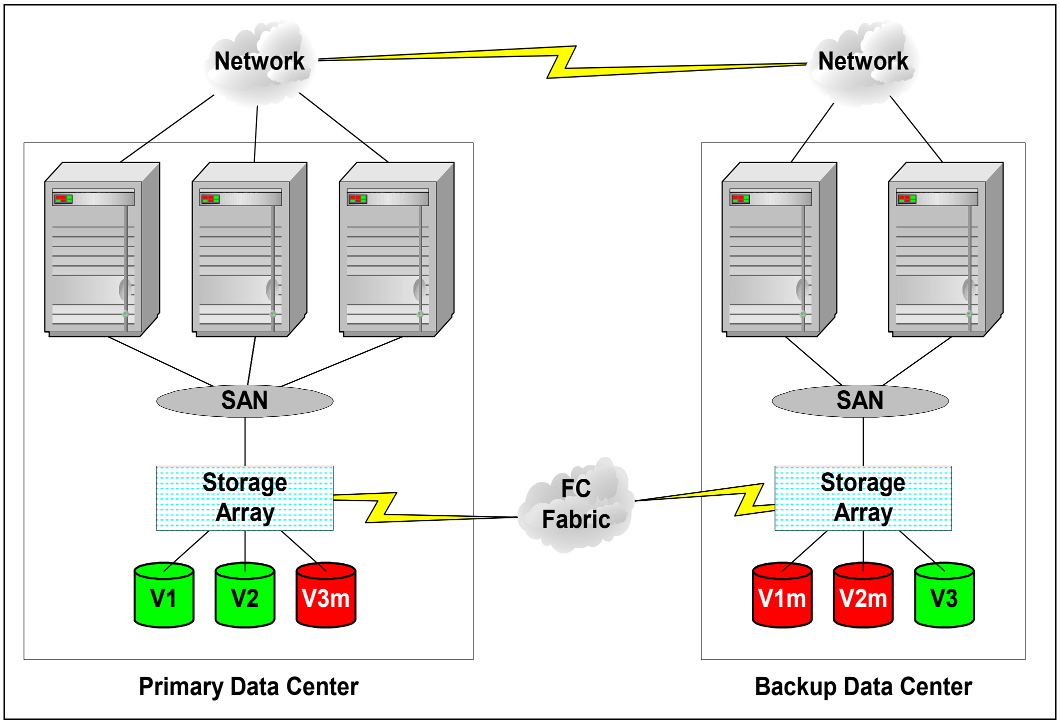

Figure 8-1 Remote Volume Mirroring

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 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.

In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Volume Mirror option enables you to promote a secondary storage subsystem to take over responsibility for normal input/output (I/O) operations.

A mirroring relationship is on a logical drive basis:

- Associates two logical drives (primary and secondary) using Storage Manager software.
- Data is copied to a secondary logical drive in the background.

The mirroring is synchronous. The write must be completed to both volumes before the host receives an I/O complete.

A minimum of two storage subsystems is required. One storage subsystem can have primary volumes being mirrored to arrays on other storage subsystems and hold secondary volumes from other storage subsystems. Also 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.

8.1.1 Primary and secondary logical drives

When you create a Remote Volume Mirror relationship, 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 remote 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 32 with Storage Manager 8.4 on the IBM FAStT900 and 700 subsystems.

8.1.2 Mirror repository logical drives

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.

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.

Notes: See the following notes:

- Two mirror repository logical drives are created for each controller in every storage subsystem that has Remote Volume Mirroring activated. That is one for Controller A and one for Controller B.
- 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.

The capacity is set at 4 MB for each logical drive. The segment size is set at 32KB (or 64 blocks). The controller determines the modification priority. The segment size and modification priority for a logical drive cannot be changed.

8.1.3 Remote Volume Mirror relationships

The primary and secondary roles in a Remote Volume Mirror relationship are 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 (remote)*. While at the same time, other logical drives in *System B (primary)* can be mirrored to logical drives in *System A (remote)*. 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 logical 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 32 using Storage Manager 8.4.

The primary drive accepts the host 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 read or write requests.

The secondary logical drive is unavailable to host applications while in a mirror relationship. In the event of a disaster or unrecoverable error of the primary storage subsystem, a role reversal can be performed to promote the secondary logical drive to the primary logical drive. Hosts mapped to this logical drive can then access the newly promoted logical drive and normal operations can continue.

8.1.4 Logical drive limits

When activating Remote Volume Mirroring, two mirror repository logical drives are created for the subsystem as discussed in the previous section.

Before activating RVM, ensure that creating the mirror repositories will not exceed the logical drive limits of the storage subsystem. Table 8-1 summarizes the logical drive limits for Storage Manager 8.4.

Table 8-1 Storage Manager 8.4 logical drive limits

Model	Maximum number of logical drives supported using SM 8.4	Comment
FAStT 600 Turbo	1024	RVM not supported on FAStT 600 Turbo
FAStT 700	2048	Logical Drives include Standard, FlashCopy Base, FlashCopy Logical Drive, FlashCopy Repository, VolumeCopy Source, VolumeCopy Target, Remote Volume Copy Repository
FAStT 900	2048	

8.1.5 Fabric configuration

The Remote Volume Mirror option requires two dedicated controller host port connections for each storage subsystem that will participate in Remote Volume Mirroring.

When the Remote Volume Mirror option is *activated* on a storage subsystem, the A2 and B2 host-side controller ports become dedicated to Remote Mirror operations. (See Figure 8-2 on page 245 for the host-side mini-bus port locations.) After Activation the A2 and B2 host-side controller ports will no longer permit host system I/O requests. The A2 and B2 host-side controller ports will *only* be able to communicate to other storage subsystems that have the Remote Mirror option activated *and* are connected by the same fabric.

A2 and B2 host-side controllers dedicated Remote 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 Mirroring ports on each controller or zone the remote mirror ports to separate zones in the existing fabric.

The level of redundancy within the fabric depends on the fabric design and Fibre Channel switch configuration. This publication does not specifically address Storage Area Network (SAN) design issues. However, you can find three sample configurations in “Remote Volume Mirroring solution design” on page 276.

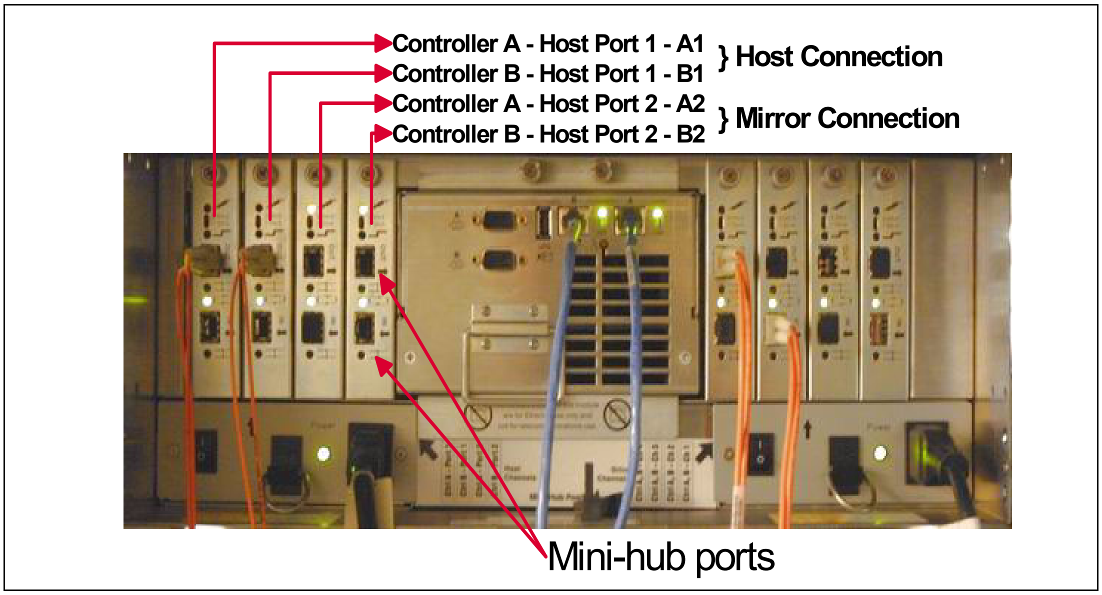

Figure 8-2 Host-side mini-bus - port locations

Distance limits

The distance between primary and remote storage subsystems is limited by the distance limits of Fibre Channel inter-switch links (ISL). See Table 8-2.

Table 8-2 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 distances greater than 10 km can be achieved by using extended fabric features and Fibre Channel extension via protocol conversion such as DWDM and FC-ATM-FC. These options have not been tested or certified with the FAStT Storage Server at the time of writing.

Fabric zoning

It is a mandatory requirement that you keep Remote Volume Mirroring links in a separate zone. 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 8-3 shows one example of zoning a fabric for FAStT and Remote Volume Mirroring.

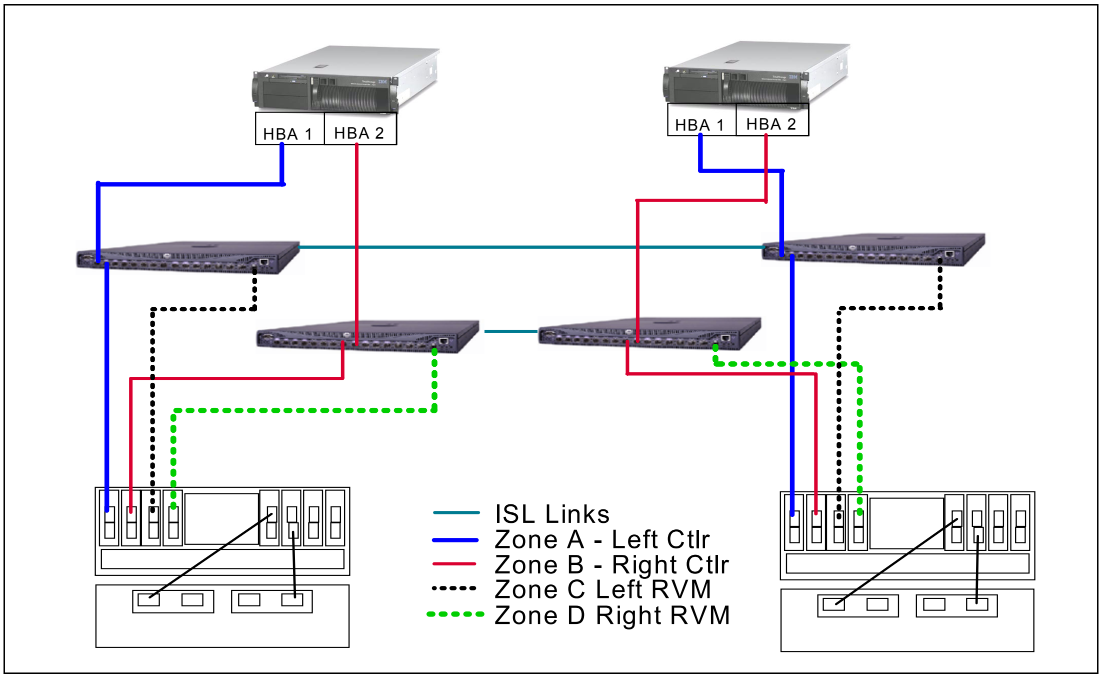

Figure 8-3 Remote Mirroring zoning example

8.1.6 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 systems and applications.

When a primary controller (the controller owner of the primary logical drive) receives a write request from a host, the controller first logs information about the write to a special logical drive called a *mirror repository logical drive*. It then writes the data to the primary logical drive. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the remote site.

Note: The owning primary controller only writes status and control information to the repository logical drive. It is not used to store host data.

After the host write to the primary logical drive is complete and the data has been copied to the secondary logical drive at the remote site, then the controller removes the log record on the mirror repository logical drive. Finally, 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 has been copied to both the primary and secondary logical drives, this mirroring operation is called *synchronous* (refer to Figure 8-4 on page 247 for a logical view of the data flow).

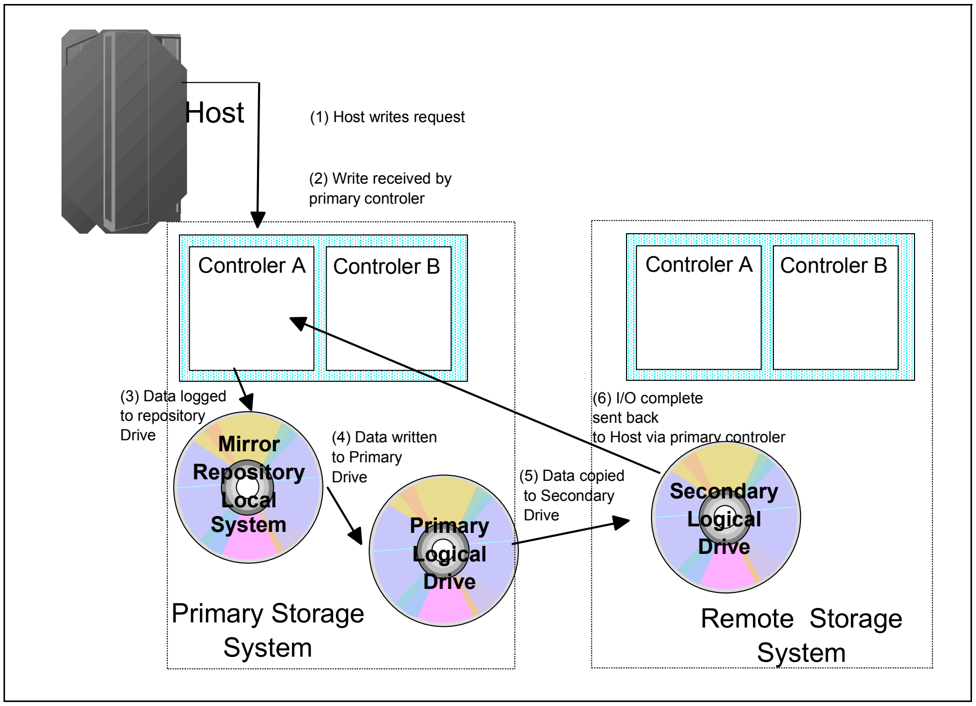

Figure 8-4 Remote Volume Mirror logical data flow

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 site (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 has been stored to physical media on that site.

When a controller receives a read request from a host system, the read request is handled normally and no communication takes place between the primary and secondary storage subsystems.

8.1.7 Link interruptions or secondary Logical Drive errors

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 *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 will 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 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.

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.

8.1.8 Remote Volume Mirroring Status icons

The status of a Remote Mirror indicates whether or not the data on the primary logical drive is replicated identically (fully synchronized) on the secondary logical drive. A mirror status is independent of the component status of the underlying logical drives in the mirrored pair. The primary and secondary logical drive icons in the Logical View indicate both the state of the logical drive as well as the state of the Remote Mirror.

There are three states of a Remote Mirror: Synchronized, Synchronization in Progress, and Unsynchronized.

The mirror status is represented with different icons in the Logical View of the Subsystem Management Console, depending on which storage subsystem you are monitoring and whether the storage subsystem contains the primary logical drive or the secondary logical drive. The table in Figure 8-5 shows a summary of the Remote Mirror status icons.

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 8-5 Remote Mirror status icons

8.1.9 Connectivity and logical drive ownership

A primary controller will only attempt 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 site, 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, or if the storage administrator changes the controller owner of the primary logical drive, the next remote write processed will automatically trigger a matching ownership change on the secondary. 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 will change the controller owner of the secondary logical drive from Controller A to Controller B. Because controller ownership changes on the secondary subsystem are controlled by the primary side, they do not require any special intervention by the storage administrator.

8.1.10 Controller resets and Storage Subsystem power cycles

If 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.

8.1.11 Performance considerations for 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 have a performance impact to the host application. The synchronization priority defines how much processing time is allocated for synchronization activities relative to system performance. The following priority rates are available: Lowest, low, medium, high, and highest.

To learn more about Remote Mirror synchronization rates, refer to “Changing mirror synchronization priorities” on page 266.

> **Note:** The lowest priority rate favors system performance, but the full synchronization will take longer. The highest priority rate favors the full synchronization, but system performance may be compromised. Logical drive size and host I/O rate loads affect the synchronization time comparisons. A full synchronization at the lowest synchronization priority rate will take approximately eight times 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. For more information on

Logical Drive properties refer to “Viewing Remote Volume Mirror properties and status” on page 262.

8.1.12 Restrictions

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 will automatically fail all associated FlashCopies.
- A primary logical drive can be a source or target logical drive in a VolumeCopy. A secondary logical drive cannot be a source or target logical drive unless a role reversal was initiated after the copy had completed. If a role reversal is initiated during a copy in progress, the copy will fail and cannot be restarted.
- A given logical drive may participate in only one mirror relationship.
- On a given FAStT900 or FAStT700 storage subsystem, up to 32 logical drives may participate in mirror relationships. Up to 2048 logical drives on a FAStT900 or FAStT700 storage subsystem are supported when the Remote Mirroring premium feature is activated.

8.2 Remote Volume Mirroring: Step-by-step

This section presents an easy-to-follow, step-by-step procedural guide to common administration tasks for a FAStT Remote Volume Mirroring. The tasks that are covered include:

- Enabling and activating the Remote Volume Mirroring premium feature
- Creating Remote Volume Mirror relationships
- Viewing Remote Volume Mirror properties
- Changing Remote Volume Mirror synchronization priorities
- Mapping a secondary logical drive
- Removing Remote Volume Mirror relationships

The configuration shown in Figure 8-6 on page 251 was used during the writing of this section to demonstrate the Remote Volume Mirroring.

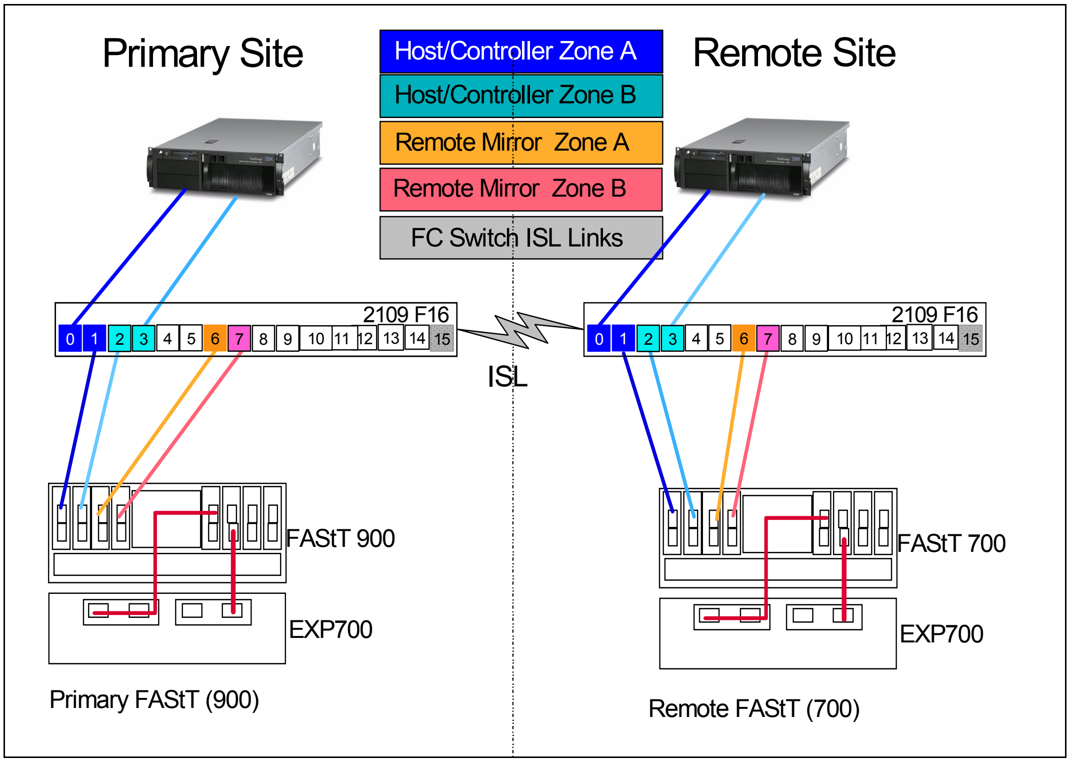

Figure 8-6 Remote Volume Mirroring - demonstration solution

8.2.1 Enabling and activating Remote Volume Mirroring

Before you can use Remote Volume Mirroring, you must first purchase the option (feature key) from IBM for *all* Storage Subsystems participating in Remote Volume Mirroring. This section discusses the process for enabling, activating, and deactivating the Remote Volume Mirror premium feature.

There are four possible states for the Remote Volume Mirror option:

- Disabled/Deactivated

  This is the default state if the feature key has not been installed.

- Enabled/Deactivated

  This is the state when the feature key has been enabled, but Remote Volume Mirroring has not been activated. There are no changes made to the Storage Subsystem configuration and all host connectivity is as per normal.

- Enabled/Activated

  This state makes controller ports A2 and B2 available for mirroring only. No host system I/Os are allowed. Up to 32 mirror pairs can now be created. This is the normal state for Remote Volume Mirroring operations.

- Disabled/Activated

  This state is possible when the premium feature has been disabled before deactivation.

  If mirrored pairs are present when this happens, an `Out of Compliance` error message will occur. Removing the mirrored pairs will remove this error. If the feature is Disabled/Activated, controller ports A2 & B2 will still be unavailable for normal host I/O activity. This is an undesirable state. The status icon shown in Figure 8-7 illustrates this state.

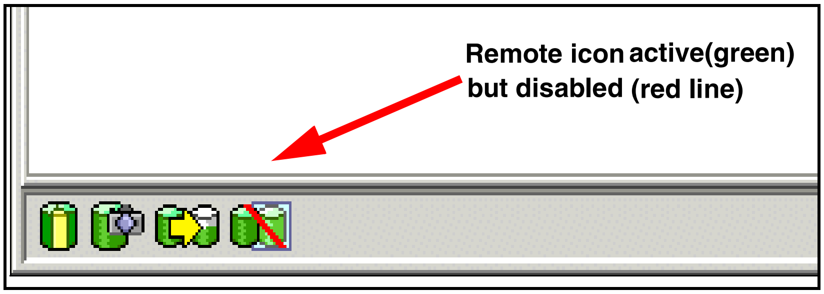

Figure 8-7 Remote Volume Mirror option Disabled/Activated

Enabling the Remote Volume Mirror premium feature

Use this section to:

- Check the premium features list.
- Enable Remote Volume Mirroring.

Remember to enable the premium feature on all participating storage subsystems.

Checking the premium features list

To check the premium option, follow these steps: From the Subsystem Management window, click **Storage Subsystem -> Premium Features -> List**.

The List Premium Features dialog opens (Figure 8-8). It lists the following items:

- Premium features enabled/disabled on the storage subsystem
- Feature enable identifier

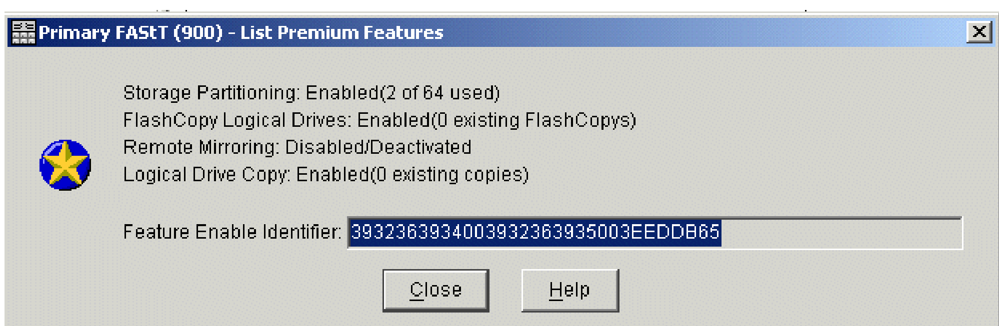

Figure 8-8 List premium features

You can also view the premium feature icon in the lower left of the Subsystem Management window (Figure 8-9 on page 253).

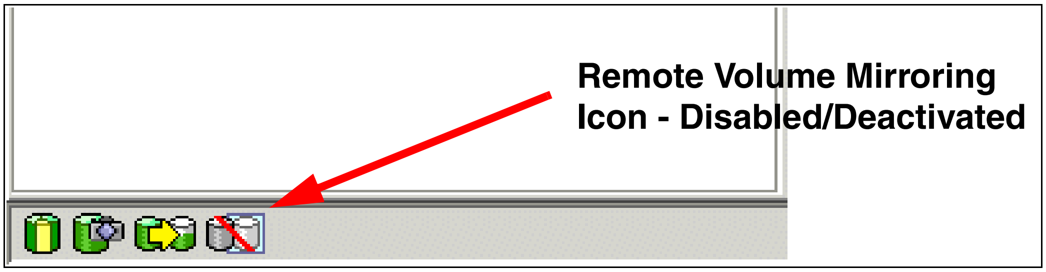

Figure 8-9 Remote Volume Mirroring Feature Icon - Disabled/Deactivated

Enabling Remote Volume Mirroring feature

From the Subsystem Management window click **Storage Subsystem -> Premium Features -> Enable**. The Select Feature Key File window opens. Chose the directory where you stored the **.key* file, then proceed. The Enable Premium Feature confirmation dialog appears (Figure 8-10). Click **Yes** to confirm.

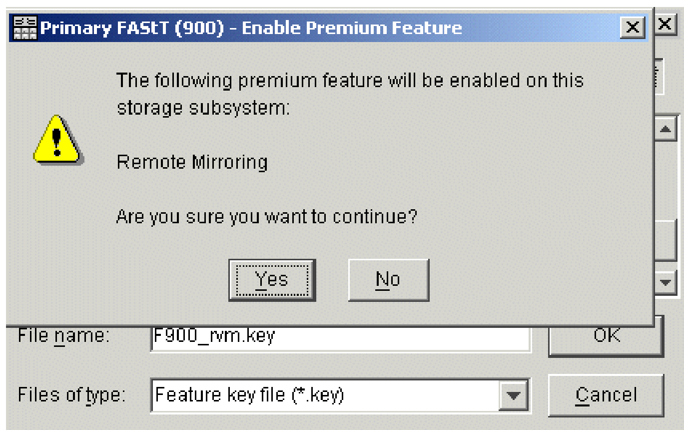

Figure 8-10 Enable Remote Volume Mirroring premium feature - confirmation

To check that the Remote Volume Mirroring premium feature is enabled, you can view the icon in the lower left corner of the Subsystem Management window (Figure 8-11). The red slash is removed.

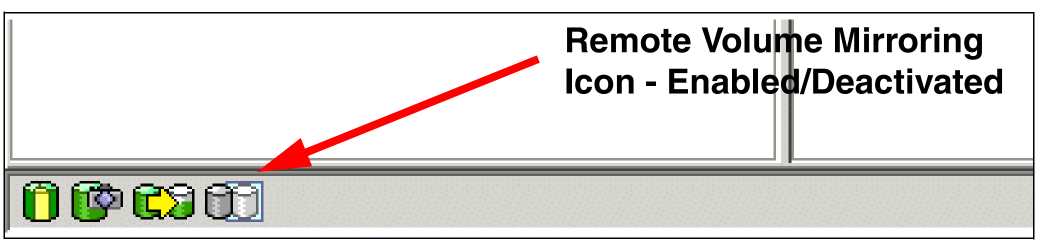

Figure 8-11 Remote Volume Mirroring icon - Enabled/Deactivated

Activating the Remote Volume Mirror option

This is achieved by using the Activate Remote Mirroring Wizard. Remember to activate the premium feature on all participating storage subsystems.

Important: The following information is important:

- When the Remote Mirroring premium feature is activated, it creates two Mirror Repository logical drives per subsystem. The FAStT900 and FAStT700 Storage Subsystems support up to 2048 logical drives. Before activating the Remote Mirroring feature, verify that the number of configured logical drives on your storage subsystem is under the supported limit.
- Ensure that the switch fabric is appropriately configured before beginning this procedure. Refer to “Fabric configuration” on page 244.
- Reservations on secondary logical drives are blocked. However, reservations on primary logical drives are allowed.

Activate the Remote Volume Mirroring feature as follows:

1. Select the **Storage Subsystem -> Remote Mirroring -> Activate** pull-down menu option, or select **Remote Mirroring -> Activate** from the right-mouse pop-up menu. Refer to Figure 8-12.

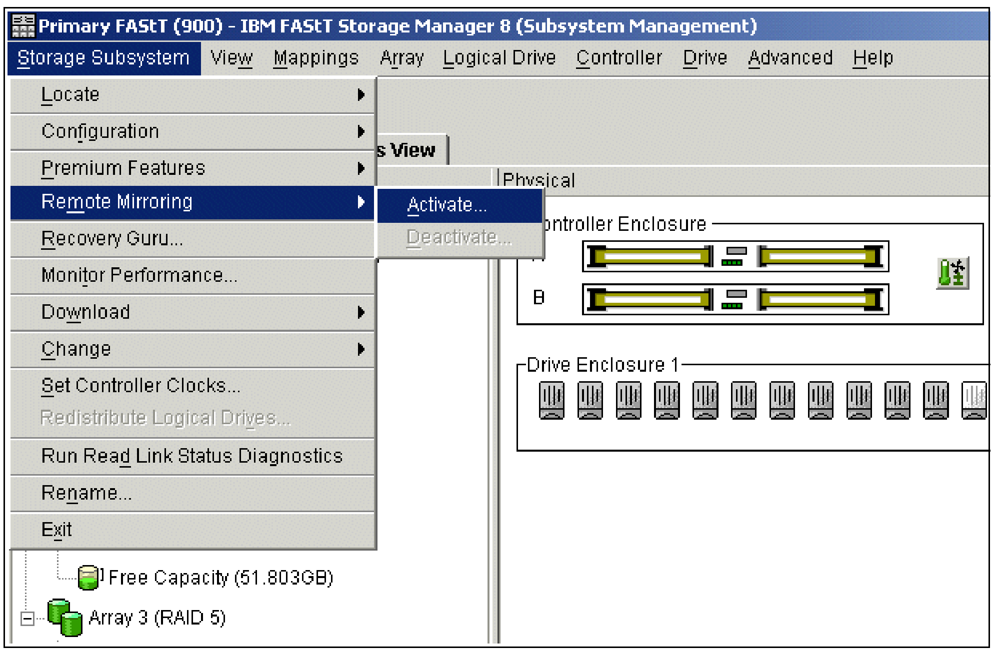

Figure 8-12 Activating Remote Mirroring

2. The Activate Remote Mirroring Wizard - Introduction dialog is opened (Figure 8-13 on page 255).

   In this dialog, you choose where to allocate capacity for the mirror repository logical drives. Read the information on the screen. You can choose to create these in either of the following:

   – An existing array with sufficient free space. If you choose this option, highlight the desired array and click **Next** to proceed to the next screen.

   – An unconfigured array. Click **Next** to proceed to the create array screen.

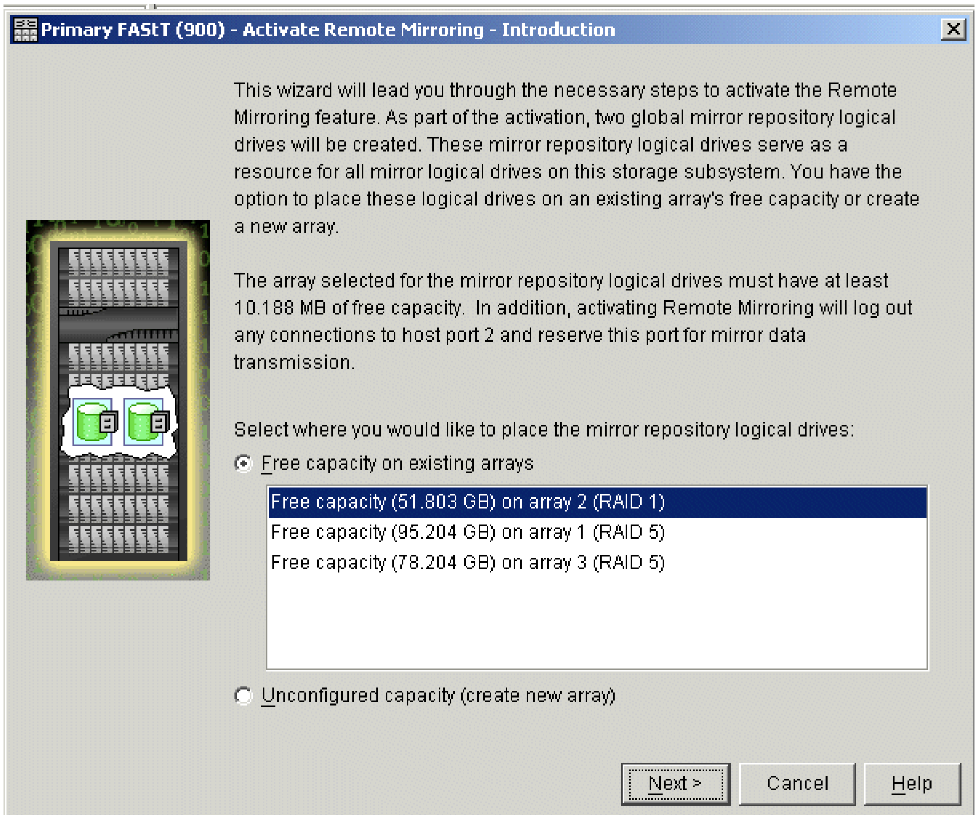

Figure 8-13 Activate Remote Mirroring Wizard - Introduction

After choosing where to create the mirror repository logical drives, the Activate Remote Mirroring Wizard Preview window is presented.

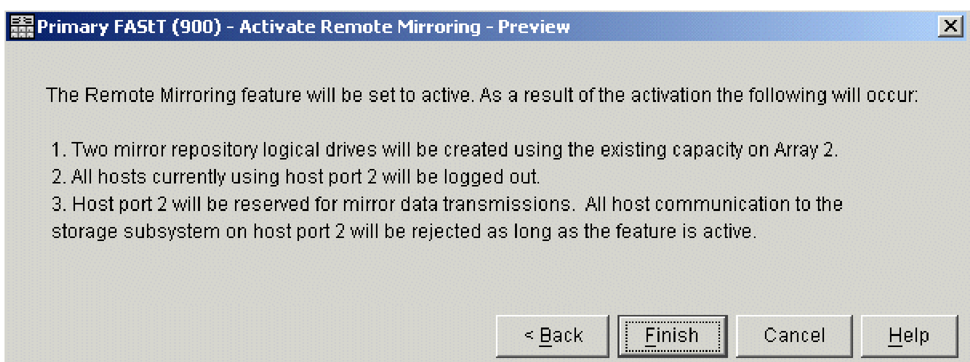

Figure 8-14 Activate Remote Mirroring Wizard - Preview

Read the information and click **Next** to continue.

The Activate Remote Mirroring Wizard Completed message is displayed (Figure 8-15 on page 256). Click **OK** to close the dialog.

Figure 8-15 Activate Remote Mirroring Wizard - Completed

Two mirror repository logical drives, one for each controller, are created as a resource for the controller managing the Remote Mirror. Host port 2 on each controller is now dedicated to Remote Mirroring communication. The dedicated host ports will not accept read/write requests from a host application.

The remote volume mirroring status icon at the lower left of the Subsystem Management window shows Enabled/Activated (Figure 8-16).

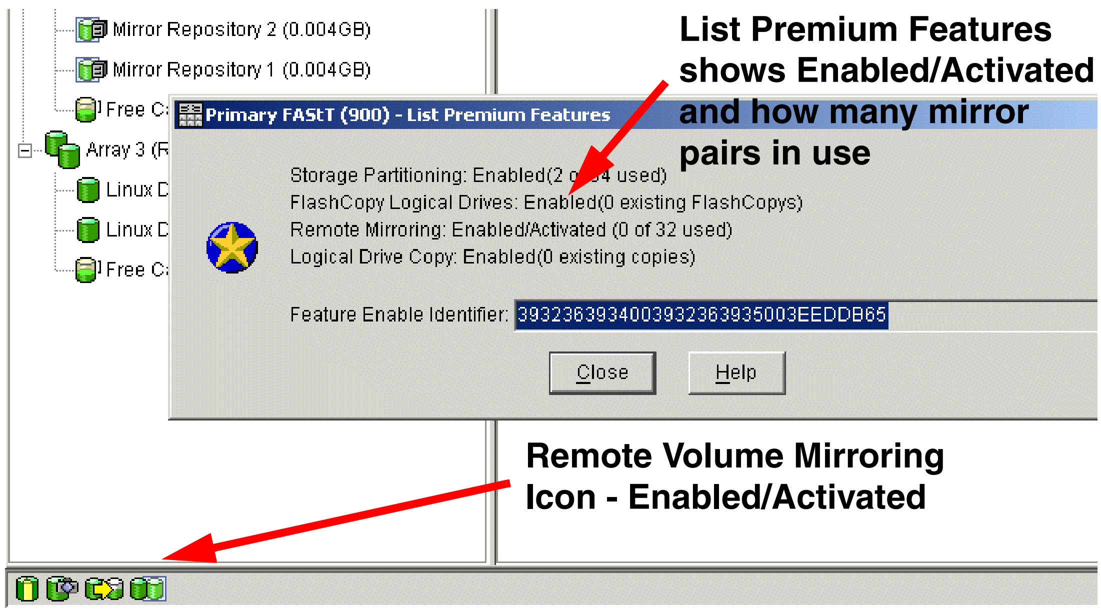

Figure 8-16 Remote Volume Mirroring - Enabled/Activated

8.2.2 Creating Remote Volume Mirror relationships

Once you have enabled and activated the Remote Volume Mirror option, you are ready to create Remote 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.

Note: The secondary logical drive must be equal to, or greater than, the size of the primary logical drive.

The primary and secondary (remote) storage subsystems must be connected as described in “Fabric configuration” on page 244.

The Create Remote Mirror wizard helps you quickly define a primary logical drive and a secondary logical drive and synchronize the data between the two logical drives.

To launch the Create Remote Mirror Wizard, select a logical drive in the Logical View of the designated primary storage subsystem and then select the **Logical Drive -> Remote Mirroring -> Create** pull-down menu option or select **Create Remote Mirror** from the right-mouse pop-up menu (Figure 8-17).

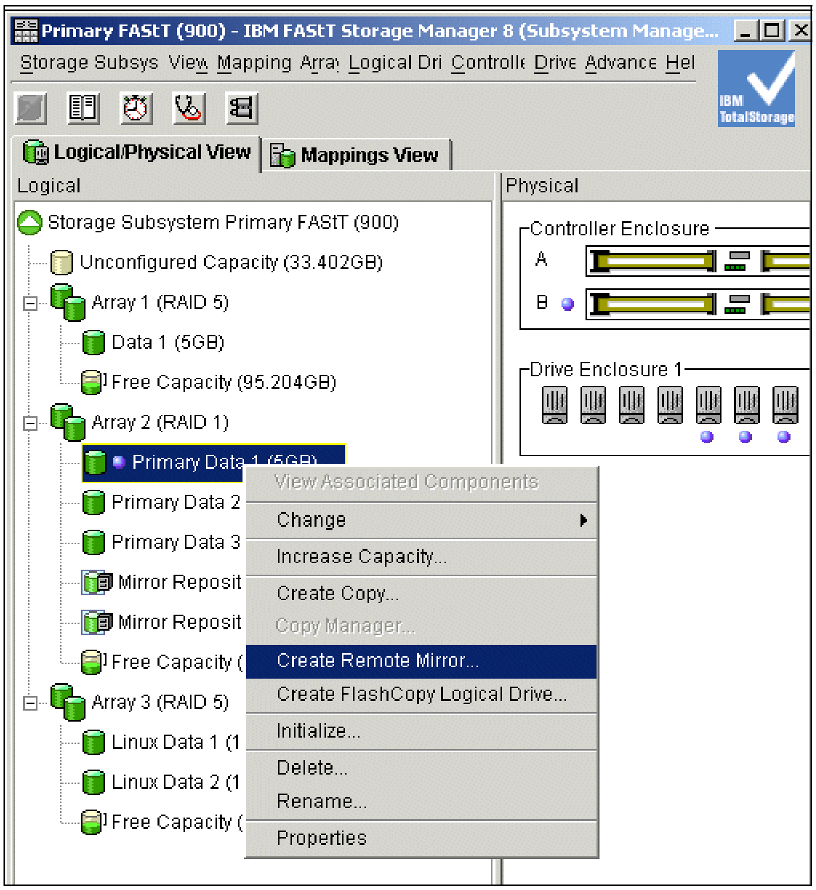

Figure 8-17 Select primary logical drive from Logical View

This launches the Create Remote Mirror - Introduction window (Figure 8-18 on page 258).

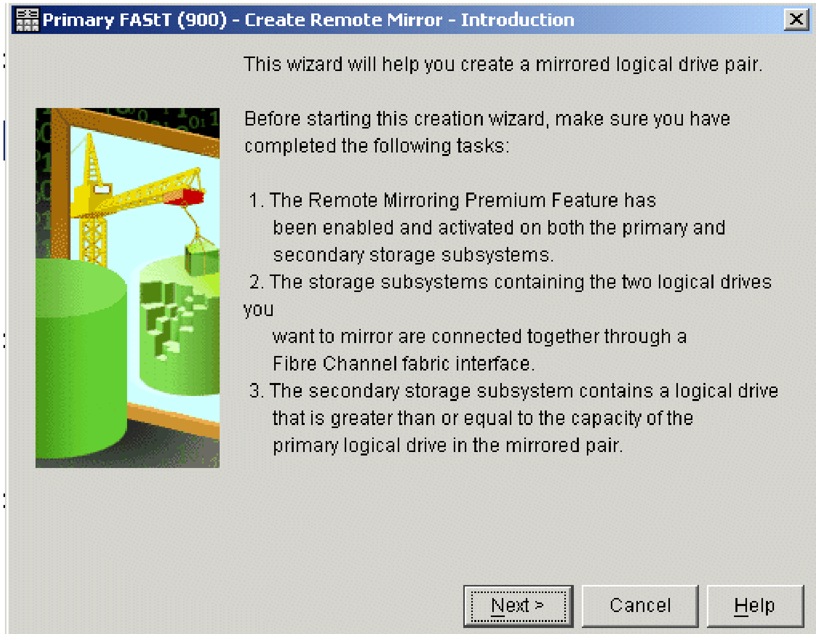

Figure 8-18 Create Remote Mirror - Introduction

Read the information displayed on the screen and click **Next**.

The Create Remote Mirror - Select Primary Logical Drive window is displayed. Use this screen to select the logical drive on the local storage subsystem to be used as the primary logical drive in the mirrored pair.

This screen is displayed when the following conditions exist:

- There are less than 32 Remote Mirrors on the local FAStT900 or FAStT700.
- There are primary logical drive candidates on the local storage subsystem.

Only valid primary logical drive candidates are displayed in the table. The valid primary logical drive candidates shown in the table are sorted alphabetically (Figure 8-19).

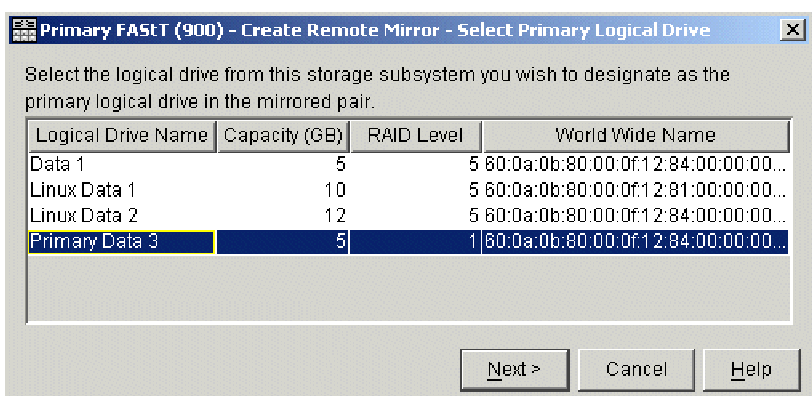

Figure 8-19 Create Remote Mirror - Select Primary Logical Drive

Important: Logical drives that are not valid primary logical drive candidates will not appear in the list (see “Restrictions” on page 250).

Select the primary logical drive to be included in the mirrored pair. You may select only one. Click **Next** to continue the mirror creation process.

The next screen displayed is the Create Remote Mirror Wizard - Selecting a Storage Subsystem for the Secondary Logical Drive (Figure 8-20).

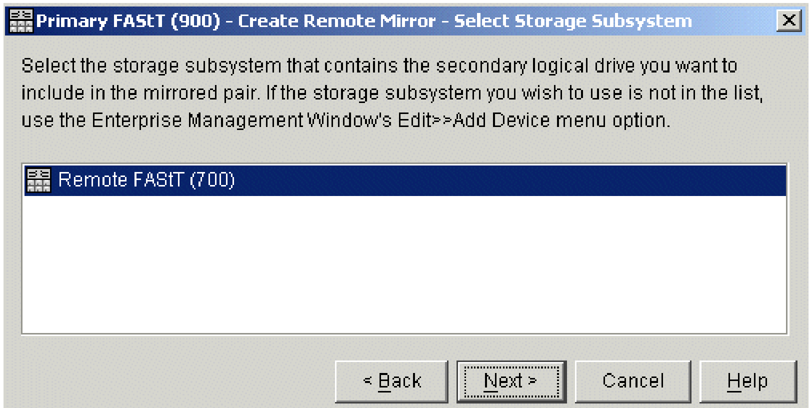

Figure 8-20 Create Remote Mirror - Select Storage Subsystem

Select the storage subsystem that contains the secondary logical drive you want to include in the mirrored pair.

Click **Next** to continue.

The next screen displayed lets you select a secondary logical drive for the Remote Mirror pair (Figure 8-21).

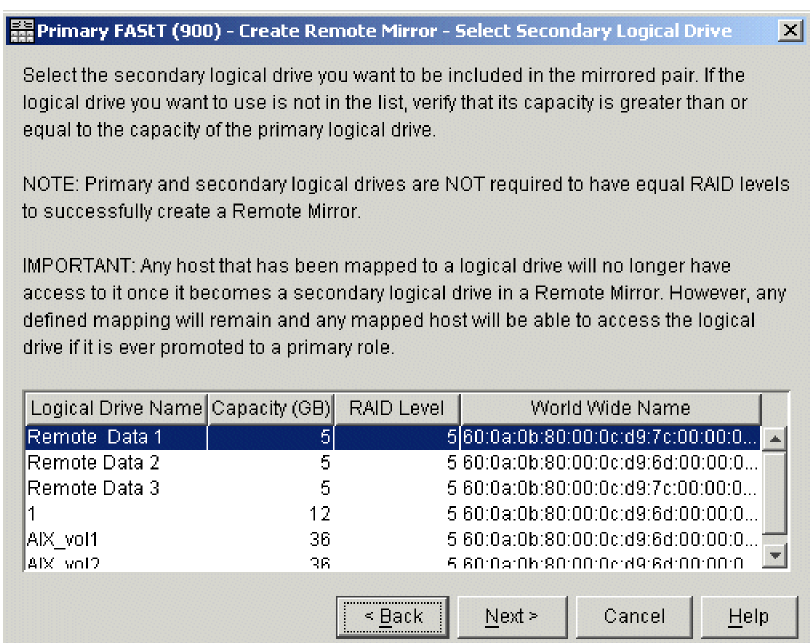

Figure 8-21 Create Remote Mirror - Select Secondary Logical Drive

The logical drives shown in the table are sorted by capacity, starting with the logical drive nearest to the primary logical drive's capacity. Logical drives with identical capacity will be sorted alphabetically. Selected primary and secondary logical drives can have different RAID levels, segment sizes, and caching parameters.

Important: Logical drives that are not valid secondary logical drive candidates will not appear in the list (see “Restrictions” on page 250).

Select the secondary logical drive to be included in the mirrored pair. You may select only one. Click **Next** to proceed. The next screen prompts you to set the synchronization priority (Figure 8-22). Use this screen to set the rate at which the controller owner of the primary logical drive will synchronize data with the secondary logical drive. The controller owner of the primary logical drive uses the synchronization priority, but a synchronization priority is also set on the secondary logical drive at this time in case it is promoted to a primary logical drive.

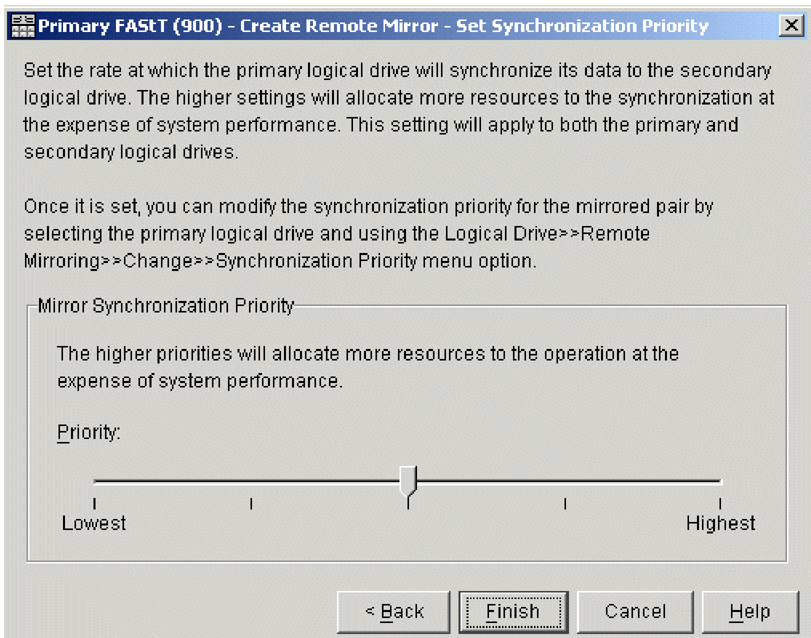

Figure 8-22 Create Remote Mirror - Set Synchronization Priority

This setting applies to both primary and secondary logical drives.

Important: If you decide later that the synchronization priority is too low or too high, you can change it. For more information see “Changing mirror synchronization priorities” on page 266.

Select **Finish** to set the synchronization priority and finish the creation of the Remote Mirror. The Creation Successful message dialog is displayed (Figure 8-23).

Figure 8-23 Create Remote Mirror - Creation Successful

If you choose Yes, then you will be taken back to the Create Remote Mirror - Select Primary Logical Drive window. You can create more mirror pairs if required. If you have no more mirror pairs to create, then select **No** to finish.

The Create Remote Mirror - Completed message is displayed (Figure 8-24).

Read the information and click **OK** to close the dialog.

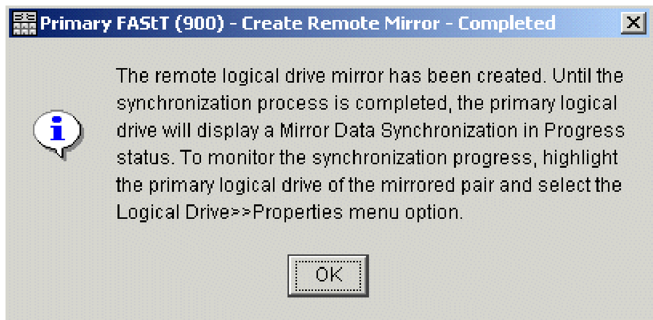

Figure 8-24   Create Remote Mirror - Completed

Looking at the logical view of the primary and secondary storage subsystems shows the synchronization in progress; icon on the primary, as well as the secondary logical drive appears as a child of the primary logical drive. See Primary and secondary logical views.

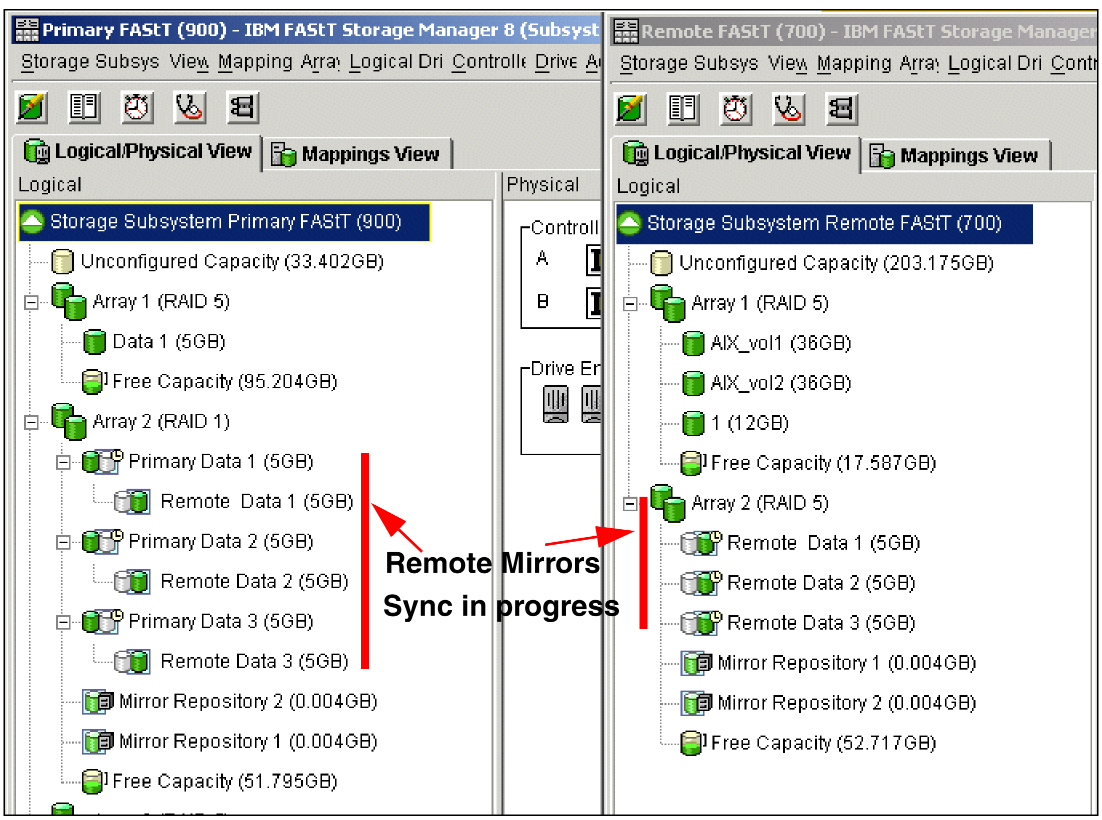

Figure 8-25   Primary and secondary logical views

8.2.3 Viewing Remote Volume Mirror properties and status

Select either the **Logical Drive -> Properties** pull-down menu option, or **Properties** from the right-mouse pop-up menu to view the Logical Drive Properties dialog. If the Remote Mirroring premium feature is enabled and you have selected a primary logical drive or a secondary logical drive, then a Mirroring tab is available on this dialog that shows mirroring-related information (Figure 8-26).

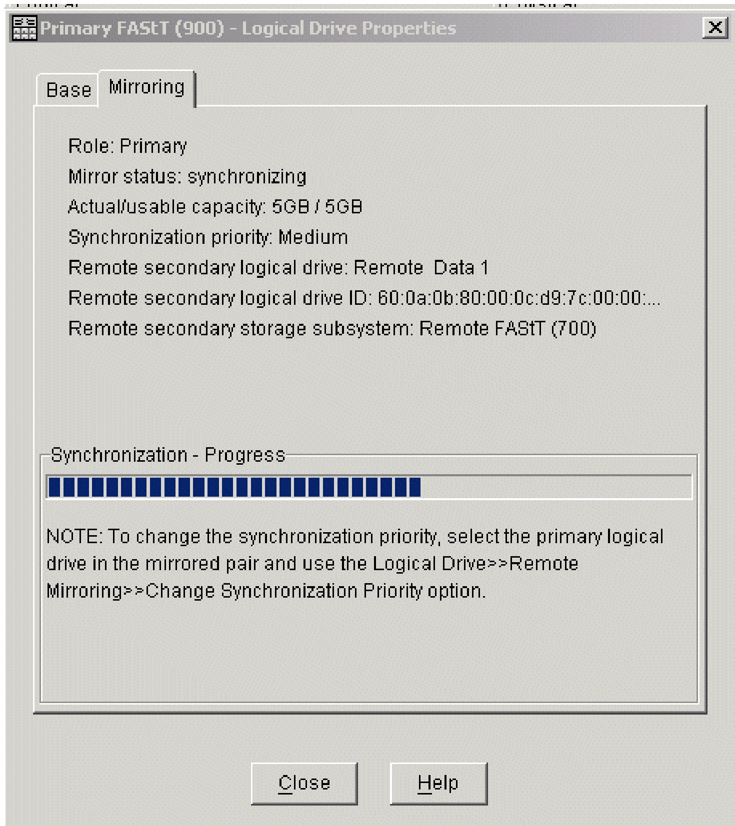

Figure 8-26 Viewing Mirroring properties

The properties displayed include:

- Role (primary or secondary)
- Mirror status (synchronized, synchronization in progress, or unsynchronized) (not displayed for remote secondary logical drives)
- Actual/usable capacity
- Synchronization priority: Lowest, low, medium, high, highest (not displayed for local or remote secondary logical drives)

If you have selected a secondary logical drive, the following properties will display:

- (Primary) logical drive: User label of the remote primary logical drive
- (Primary) World-wide Name (WWN): World Wide Name (WWN) of the remote primary logical drive
- (Primary) storage subsystem: WWN of the remote storage subsystem containing the primary logical drive

If you have selected a primary logical drive, the following properties will display:

- (Secondary) logical drive: User label of the remote secondary logical drive
- (Secondary) World-wide Name (WWN): WWN of the remote secondary logical drive
- (Secondary) storage subsystem: WWN of the remote storage subsystem containing the secondary logical drive

Note: If the remote logical drive in the mirrored pair is in an Unresponsive state, the following fields will display in place of the fields shown above:

- Unresponsive remote logical drive: User label of the remote logical drive, from last known data in the cache
- Unresponsive remote logical drive world-wide name (WWN): WWN of the remote logical drive, from last known data in the cache
- Unresponsive remote storage subsystem: WWN of the remote storage subsystem, from last known data in the cache

The progress bar at the bottom of the Mirroring tab of the Logical Drive Properties dialog displays the progress of a full synchronization.

Select **Close** to exit the dialog.

Remote Volume Mirroring can also be monitored by viewing the Storage Subsystem Profile.

To view the Storage Subsystem Profile, select either the **View -> Storage Subsystem Profile** pull-down menu option, or **View Profile** from the right-mouse pop-up menu. The Storage Subsystem Profile dialog is displayed (Figure 8-27).

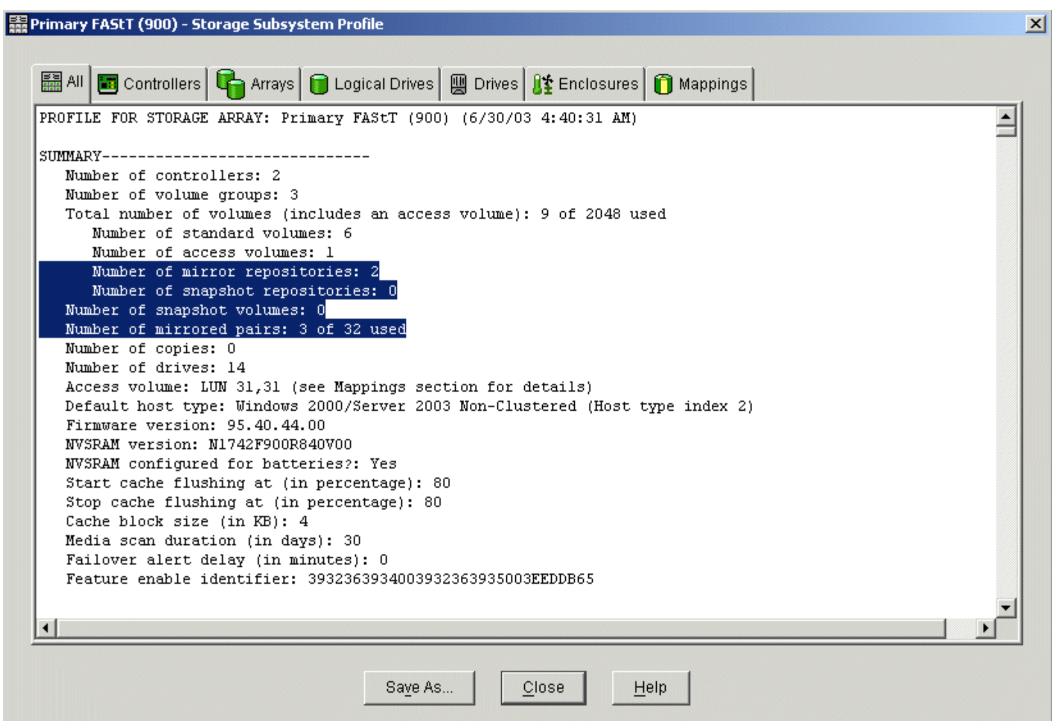

Figure 8-27 Storage Subsystem Profile - All tab

The All tab displays a summary of available Storage Subsystem Profile data, and includes a summary of Mirror Repositories and number of mirror pairs in use.

The Logical Drives tab is divided into five sections, with separate tabs providing access to information for each type of logical drive.

The Repositories tab contains a section of logical drive information that displays FlashCopy and mirror repository logical drive properties. See Figure 8-28.

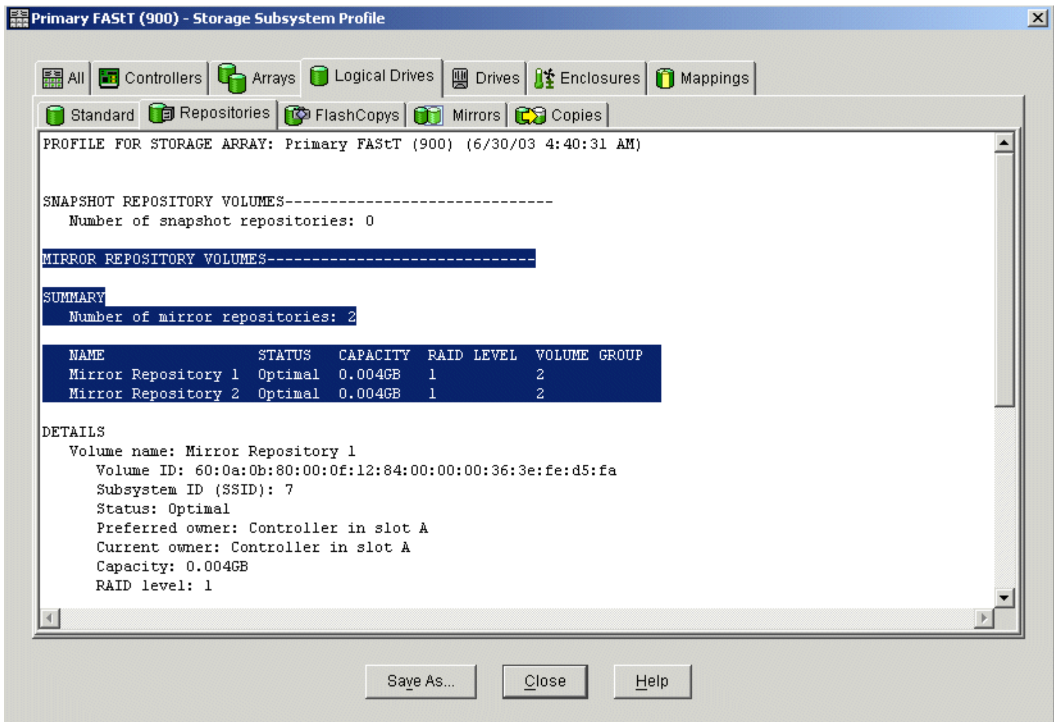

Figure 8-28 Storage Subsystem Profile - Repositories tab

The Mirrors tab is a section of logical drive information that displays Remote Mirror properties. The Mirrors tab is only available when the Remote Mirroring premium feature is enabled and activated. See Figure 8-29 on page 265.

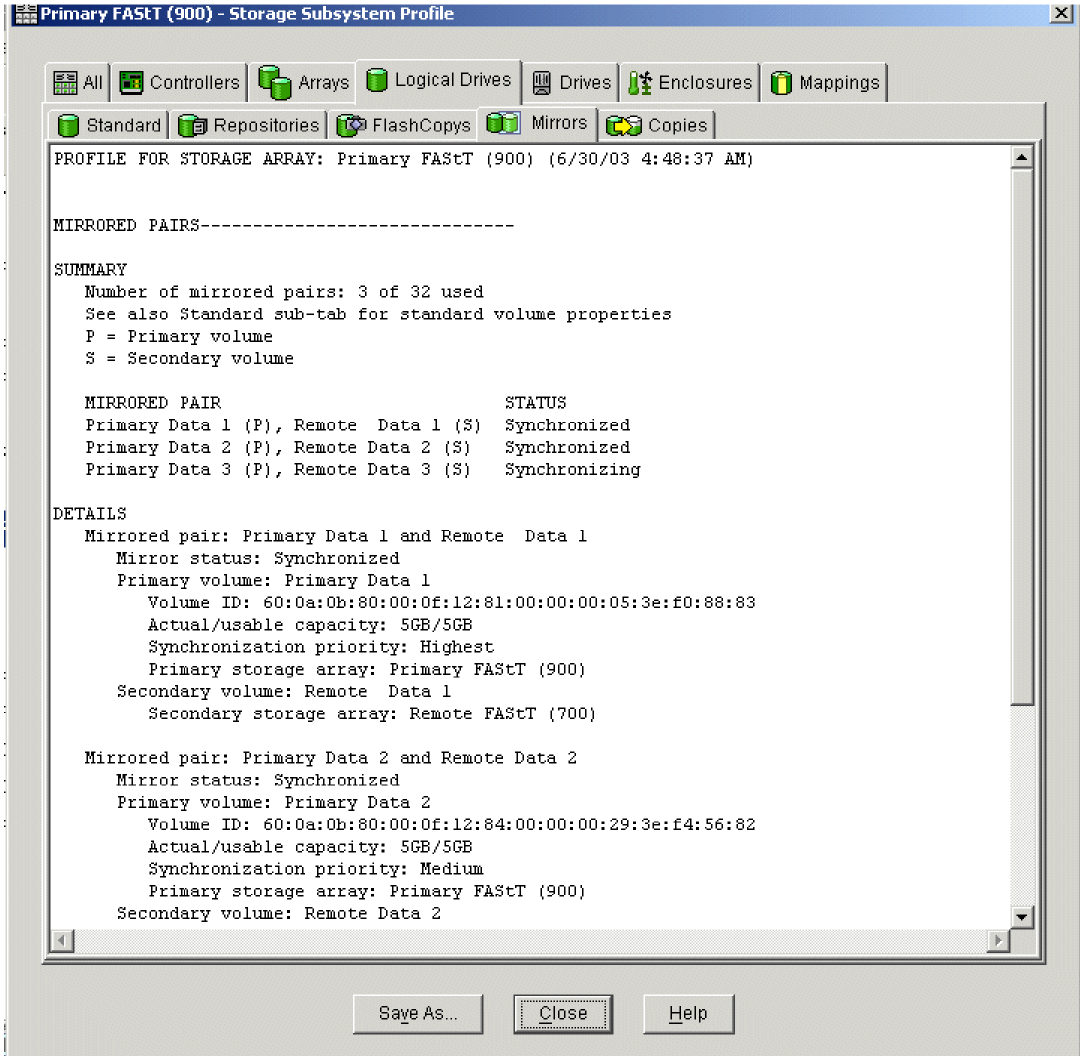

Figure 8-29 Storage Subsystem Profile - Mirrors tab

The View Associated Components window (Figure 8-30 on page 266) 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**.

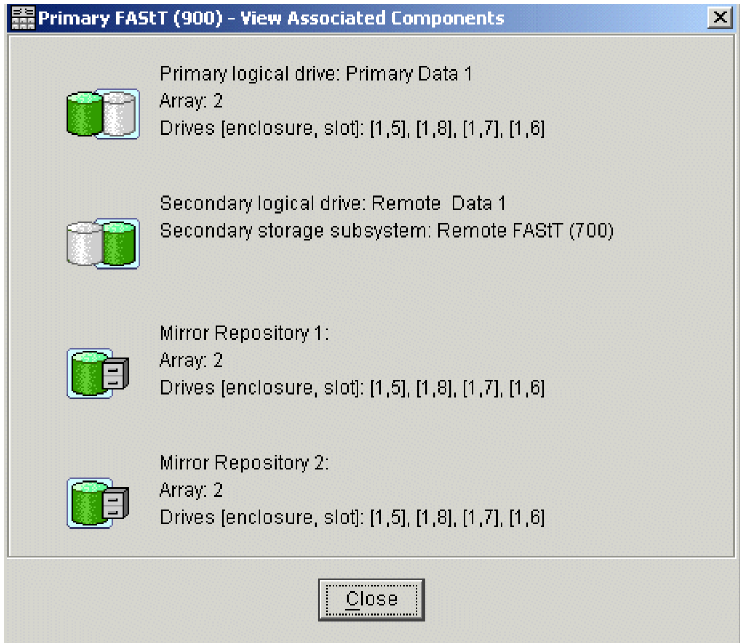

Figure 8-30 View Associated Components

8.2.4 Changing mirror synchronization priorities

Use this option to change the synchronization priority the controller owner of the primary logical drive uses when performing a full synchronization. The controller owner of the primary logical drive uses the synchronization priority, so this option is not available for secondary logical drives.

Select a primary logical drive in the Subsystem Management Window Logical View, then select either the **Logical Drive -> Remote Mirroring -> Change -> Synchronization Priority** pull-down menu option, or **Change -> Synchronization** priority from the right-mouse pop-up menu.

The Change Synchronization Priority dialog is displayed. The Select logical drives list box lists all primary logical drives in the storage subsystem. The primary logical drive that you selected to invoke the option is highlighted (see Figure 8-31 on page 267).

Select the primary logical drives for which you want to change the synchronization priority in the Select logical drives list box. To select multiple logical drives, use the Ctrl + click or Shift + click keyboard options. To select all of the primary logical drives, select the **Select All** button.

In the Select synchronization priority area, use the slider bar to select the appropriate synchronization priority.

There are five choices: Lowest, low, medium, high, and highest. These settings are shown with five tick marks on the slider bar. The lowest and highest settings are labeled. The higher the setting, the more the primary controller will prioritize a full synchronization over host I/O, and the larger the impact to system performance. Select **OK** to proceed.

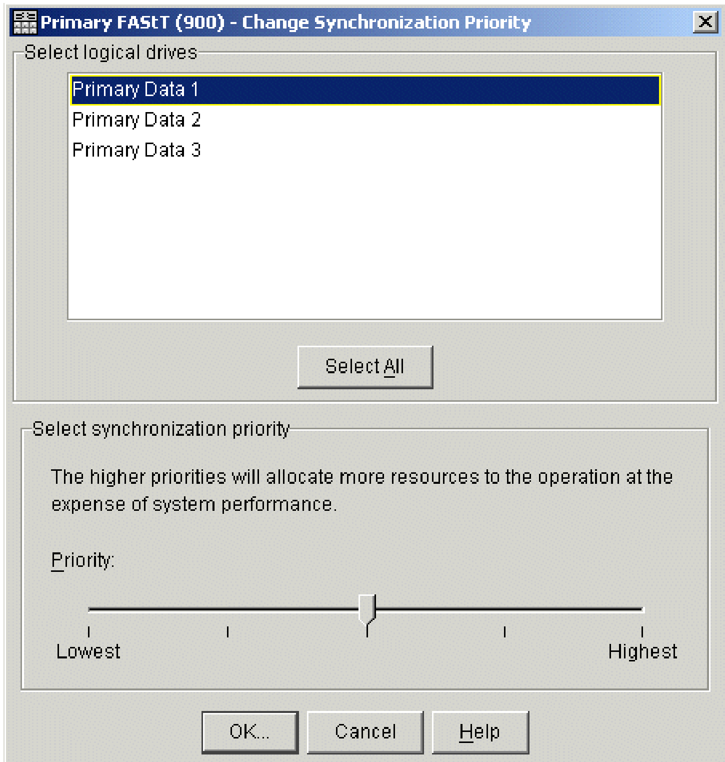

Figure 8-31 Change Synchronization Priority

The Change Synchronization Priority - Confirmation dialog is displayed (Figure 8-32).

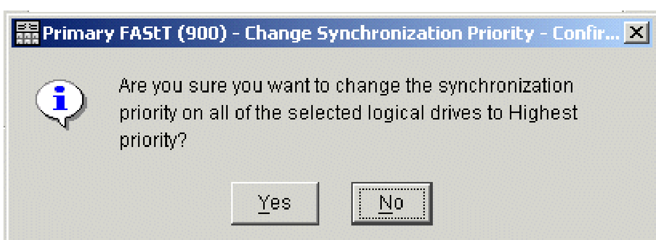

Figure 8-32 Change Synchronization Priority - Confirmation

If the parameters you have chosen are correct, then select **Yes** to continue. The `Completed` message appears. Click **OK** to return to the Logical View.

8.2.5 Mapping a secondary drive

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.

8.2.6 Removing mirror relationships

Use the Remove Mirror Relationship option to remove the mirror relationship between the two logical drives in a Remote Mirror.

This option does not delete the primary logical drive, secondary logical drive, or the mirror repository logical drives that support mirroring for the storage subsystems. Data on the logical drives is not affected. As a result of this operation the primary logical drive and secondary logical drive become standard, host-accessible, non-mirrored logical drives.

This option is only available for the local logical drive (primary or secondary) that is present in the storage subsystem you are currently managing. This option is not available if you select a remote secondary logical drive in the logical view.

The mirror relationship is first removed on the local storage subsystem, then on the remote storage subsystem. If the mirror relationship is successfully removed on the local storage subsystem but cannot be removed on the remote storage subsystem because of a communication problem, then an error message is displayed. The error message shows the name of the remote storage subsystem with the orphaned mirrored logical drive, and the name of the logical drive. Open the Subsystem Management window for the remote storage subsystem, select the specified logical drive, and remove the mirror relationship.

If the mirror relationship was successfully removed on the secondary side, but not the primary side, then the first I/O write to the primary logical drive will cause the mirror state to go unsynchronized. The primary Subsystem Management Window Logical View will also then show an unresponsive remote secondary logical drive. Remove the mirror relationship on the primary storage subsystem to correct the problem.

If the mirror relationship was successfully removed on the primary side, but not on the secondary side, then there will not typically be a change to the representation of the secondary logical drive in the secondary Subsystem Management Window Logical View. Remove the mirror relationship on the secondary storage subsystem to correct the problem and make the logical drive accessible by hosts.

To remove a remote mirror relationship, select a primary or secondary Logical Drive in the Logical View (whichever is local), then select either the **Logical Drive -> Remote Mirroring >> Remove Mirror Relationship** pull-down menu option, or **Remove Mirror Relationship** from the right-mouse pop-up menu.

The Remove Mirror Relationship dialog is displayed (Figure 8-33).

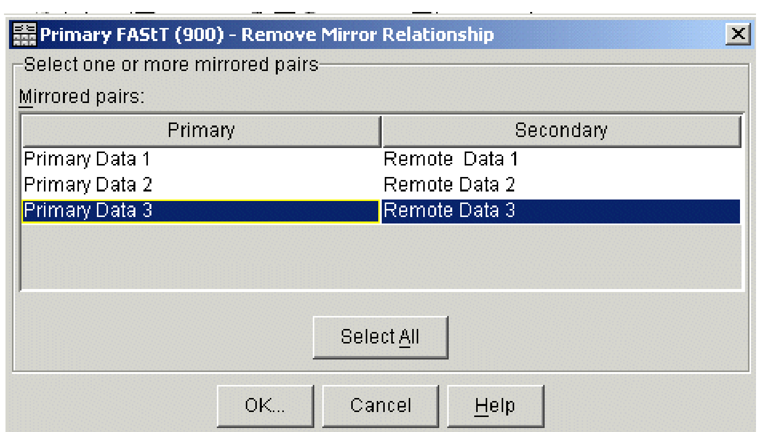

Figure 8-33 Remove Mirror Relationship

Select the mirrored logical drive pairs for which you wish to remove the mirror relationship. To select multiple mirrored pairs, use the Ctrl + click or Shift + click keyboard options. To select all of the mirrored pairs, select the **Select All** button. Select **OK**. The Remove Mirror Relationship - Confirmation dialog is displayed (see Figure 8-34).

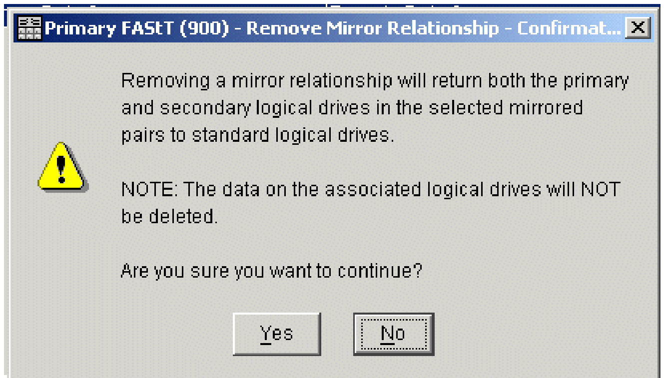

Figure 8-34 Remove Mirror Relationship - Confirmation

Select **Yes**. The Remove Mirror Relationship - Complete dialog is displayed. See Figure 8-35. When the mirror relationship is removed, the logical drives in the mirrored pairs that you selected are no longer part of a Remote Mirror. All data on the logical drives remains intact and available. Logical drives that were the secondary logical drive in the mirrored pair are now accessible by hosts that are mapped to them.

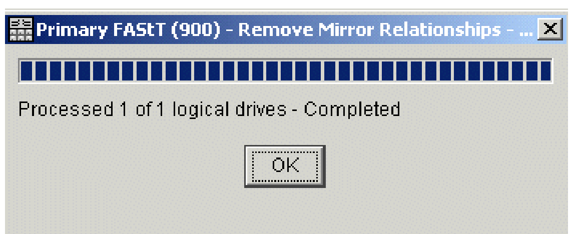

Figure 8-35 Remove Mirror Relationships - Complete

8.3 Using Remote Volume Mirroring for disaster recovery

As modern business pressures increasingly require 24-hour data availability, system administrators are required to ensure that critical data is safeguarded against potential disasters. Additionally, storage administrators are searching for methods to migrate from one host to another or from one storage subsystem to another, with as little disruption as possible.

Remote Volume Mirroring of data is one methodology that can be implemented to assist in disaster recovery and business continuity. Once critical data has been mirrored from a primary storage subsystem to a remote storage subsystem, primary and secondary logical drives can have their roles reversed so that the copied data can be accessed from the remote storage subsystem.

This next section discusses how to reverse the roles of the primary and secondary logical drives for disaster recovery. This section does *not* cover building and implementing a complete disaster recovery plan.

8.3.1 Reversing the roles of primary and secondary logical 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 logical drive to a primary logical drive

Use this option to perform a role reversal between the two paired logical drives in a Remote Mirror. This option promotes a selected secondary logical drive to become the primary logical drive of the mirrored pair. If the associated primary logical drive can be contacted, the primary logical drive is automatically demoted to be the secondary logical drive. Use this option when a catastrophic failure has occurred to the storage subsystem that contains the primary logical drive and you want to promote the secondary logical drive so that hosts can access data and business operations can continue.

Important: The following is important.

- When the secondary logical drive becomes a primary logical drive, any hosts that are mapped to the logical drive through a logical drive-to-LUN mapping will now be able to read or write to the logical drive.
- If a communication problem between the secondary and primary sites prevents the demotion of the primary logical drive, an error message is displayed. However, you are given the opportunity to proceed with the promotion of the secondary logical drive, even though this will lead to a Dual Primary Remote Mirror status condition. For more information see “Forcing a secondary to change to a primary logical drive” on page 271.

Select the secondary logical drive in the Logical View, then select either the **Logical Drive -> Remote Mirroring -> Change -> Role to Primary** pull-down menu option, or **Change -> Role to Primary** from the right-mouse pop-up menu (Figure 8-36).

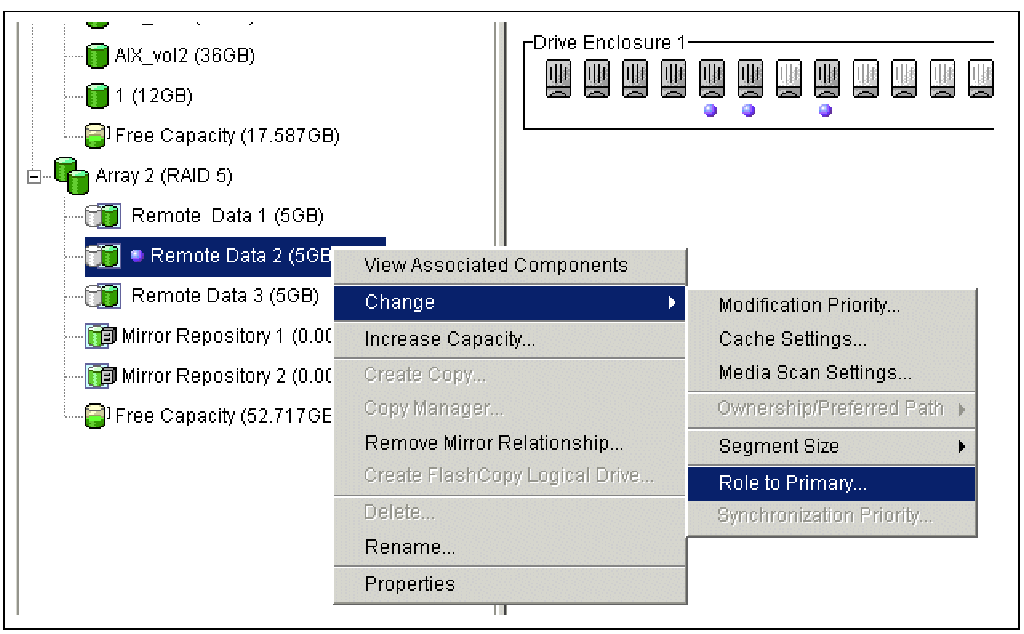

Figure 8-36 Change role - Secondary to primary

The Change to Primary dialog is displayed (Figure 8-37 on page 271).

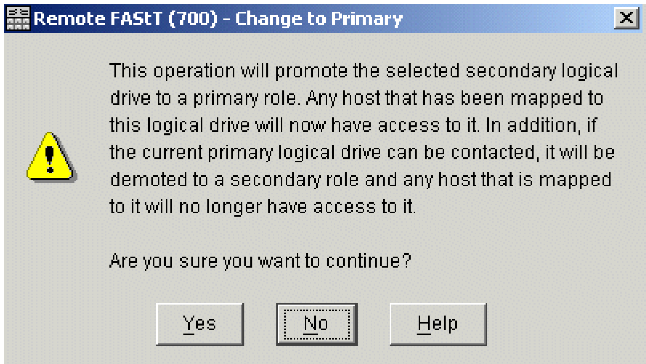

Figure 8-37 Change to Primary - Confirmation

Select **Yes** to proceed. The secondary logical drive is promoted to be the primary logical drive in the Remote Mirror. If the controller owner of the primary logical drive can be contacted, the primary logical drive is automatically demoted to be the secondary logical drive in the Remote Mirror.

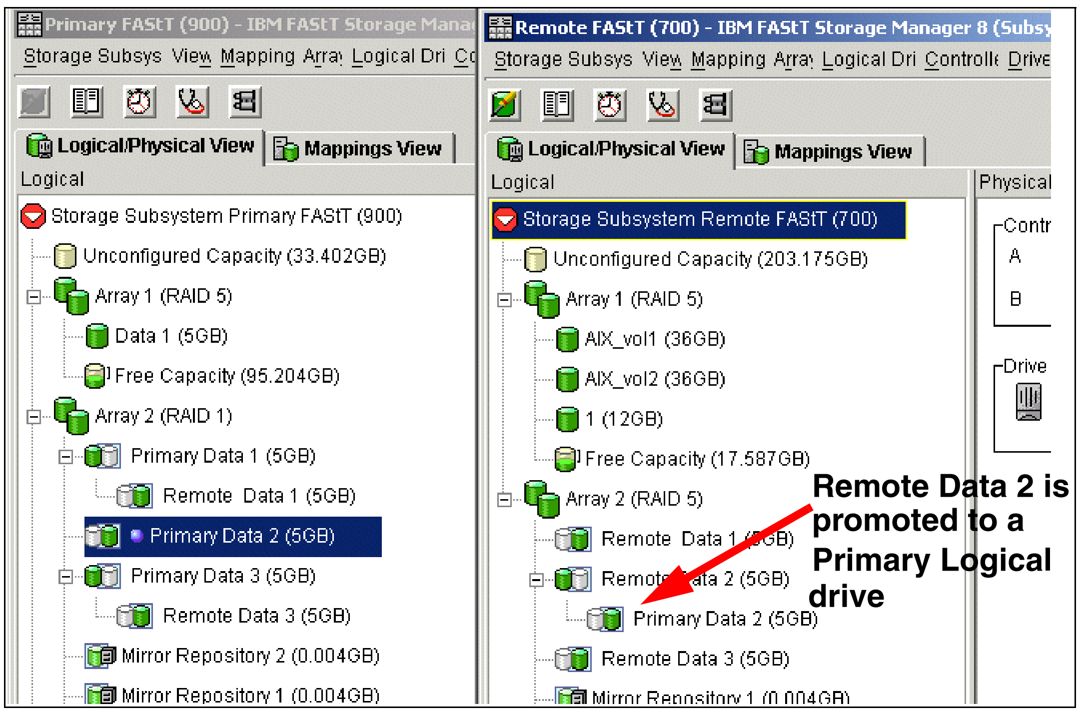

Figure 8-38 Change secondary to primary completed

In the example illustrated in Figure 8-38, the Logical View shows the secondary logical drive *Remote Data 2* in the remote FAStT (700) storage subsystem has been promoted to primary. The secondary logical drive *Primary Data 2*, from the Primary FAStT (900) storage subsystem is displayed as a child node of the Remote Data 2 in the Subsystem Management window.

Forcing a secondary to change to a primary logical drive

If, when attempting to promote a secondary logical drive to a primary logical drive, a communication failure between the primary and remote storage subsystems exists, then an error dialog will appear (Figure 8-39 on page 272).

Select **Yes** on the Change to Primary - Error dialog to force the software to promote the selected secondary logical drive to a primary role.

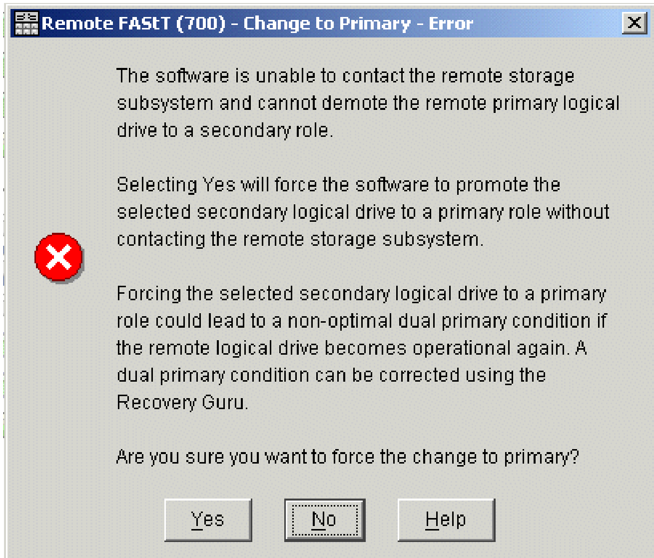

Figure 8-39 Force secondary to primary - confirmation

Use this option when a catastrophic failure has occurred to the storage subsystem that contains the primary logical drive, the primary storage subsystem cannot be contacted, and you want to promote the secondary logical drive so that hosts can access data and business operations can continue. Or consider using this option when no catastrophic failure has occurred to the storage subsystem that contains the primary logical drive, but you want to perform a role reversal for some other reason and a communication problem between the local and remote sites is preventing the demotion of the primary logical drive at the remote site.

Figure 8-40 shows a logical volume *Remote Data 2* on a remote subsystem that was forced to primary. Because the connection with the primary is lost, it is shown as unsynchronized.

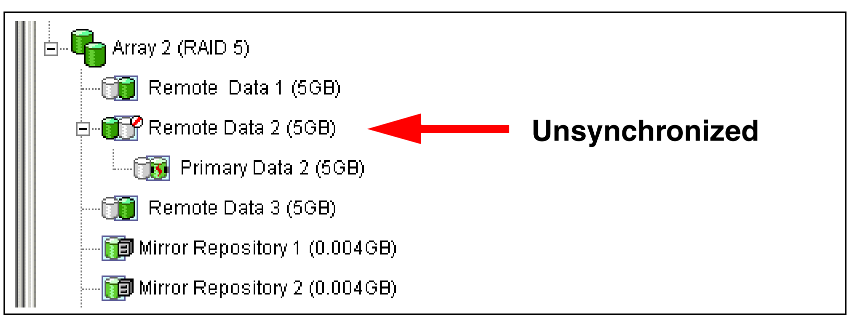

Figure 8-40 Unsynchronized secondary

When the remote storage subsystem has recovered and any communication problems have been resolved, a Dual Primary error condition will be raised. Click the **Recovery Guru** to resolve the condition (see Figure 8-41 on page 273). Follow the procedure to remove and re-create the mirror relationship.

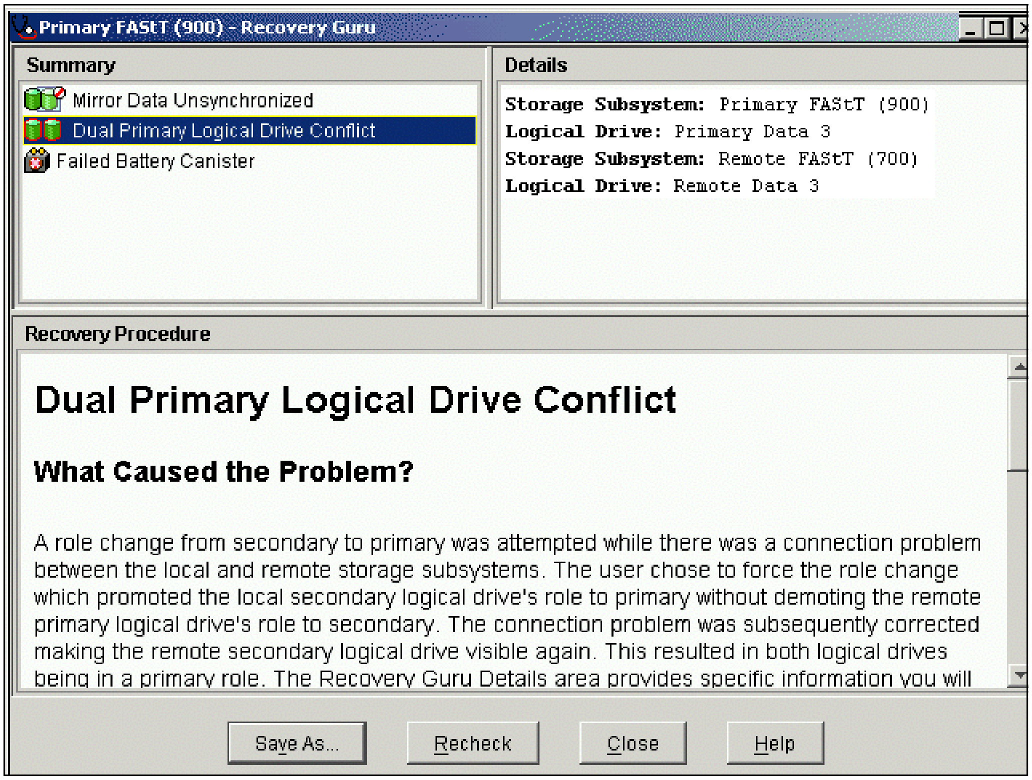

Figure 8-41 Recovery Guru - Dual primary logical drive conflict

To avoid the Dual Primary error condition and subsequent recovery steps when no catastrophic failure has occurred to the storage subsystem containing the primary logical drive, wait until the connection between the storage subsystems is operational to perform the role reversal.

Changing a primary to a secondary logical drive

Use this option to perform a role reversal between the two paired logical drives in a Remote Mirror. This option demotes a selected primary logical drive to become the secondary logical drive of the mirrored pair. If the associated secondary logical drive can be contacted, the secondary logical drive is automatically promoted to be the primary logical drive. Use this option for role reversals during normal operating conditions. You can also use this option during a Recovery Guru procedure for a Dual Primary Remote Mirror status condition.

> **Important:** Any hosts that are mapped to the primary logical drive through a logical drive-to-LUN mapping will no longer be able to 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 will be written to the logical drive.

If a communication problem between the primary and secondary sites prevents the promotion of the secondary logical drive, an error message is displayed. However, you are given the opportunity to proceed with the demotion of the primary logical drive, even though this will lead to a Dual Secondary Remote Mirror status condition.

For more information see “Forcing a primary to become a secondary logical drive” on page 274.

If the selected primary logical drive has associated FlashCopy logical drives, demoting this logical drive to a secondary role will cause the associated FlashCopy logical drives to fail.

To demote a primary logical drive to the secondary logical drive role, select the primary logical drive in the Logical View, then select either the **Logical Drive ->Remote Mirroring -> Change -> Role to Secondary** pull-down menu option, or **Change -> Role to Secondary** from the right-mouse pop-up menu.

The Change to Secondary Confirmation dialog is displayed. Read the information and select **Yes** to proceed.

The primary logical drive is demoted to be the secondary logical drive in the Remote Mirror. If the controller owner of the secondary logical drive can be contacted, the secondary logical drive is automatically promoted to be the primary logical drive in the Remote Mirror.

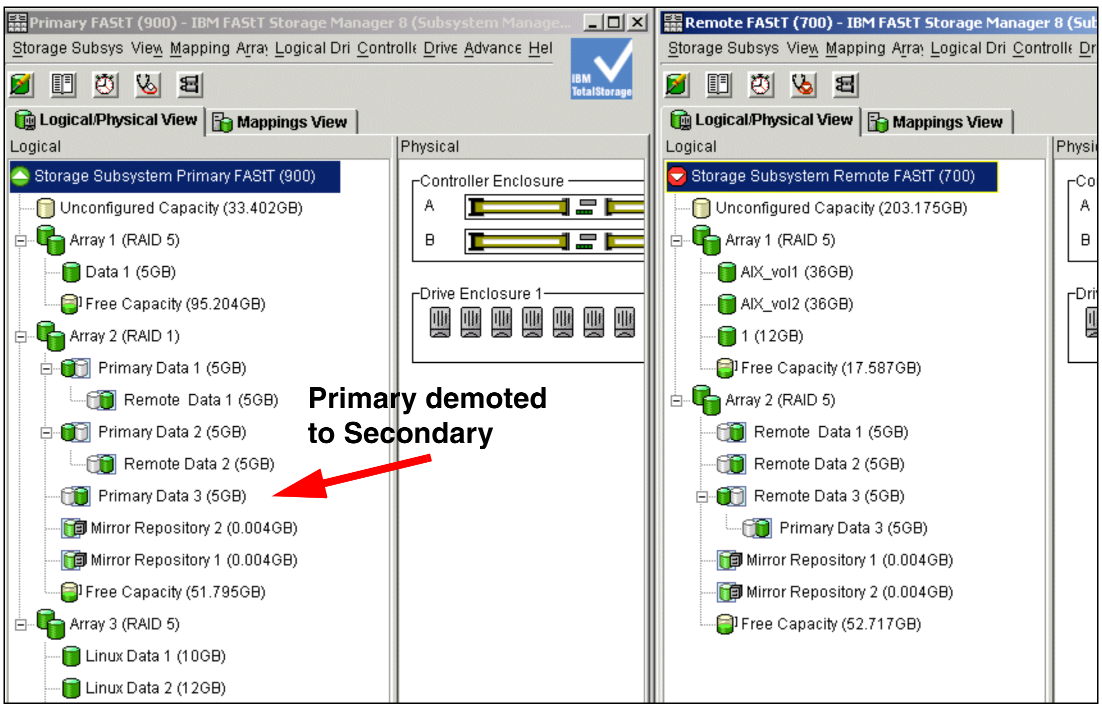

Figure 8-42 Change primary to secondary completed

Figure 8-42 illustrates that the Primary Logical Drive *Primary Data 3* in the Primary FAStT storage subsystem has been demoted to secondary. The (former) Secondary Logical Drive *Remote Data 3*, from the Remote FAStT storage subsystem, is now the primary logical drive of the Remote Mirror and is no longer displayed as a child node in the Logical View for the Primary FAStT storage subsystem.

Forcing a primary to become a secondary logical drive

If, when attempting to demote a primary logical drive to a secondary logical drive, a communication failure between the primary and remote storage subsystems exists, then an error dialog will appear (see Figure 8-43 on page 275).

Select **Yes** on the Change to Secondary - Error dialog to force the software to demote the selected primary logical drive to a secondary role.

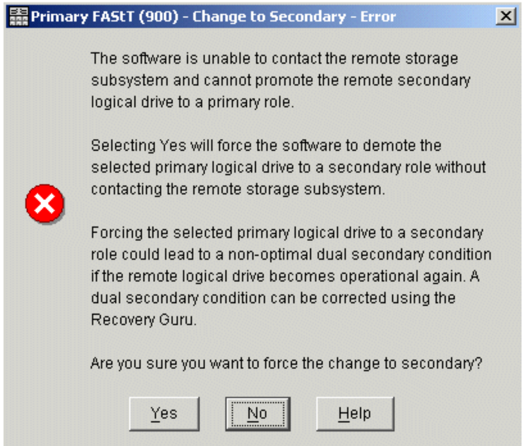

Figure 8-43 Force primary to secondary - Confirmation

Consider using this option when you want to perform a role reversal and a communication problem between the local and remote sites is preventing the promotion of the secondary logical drive at the remote site.

When the communication problem has been resolved, there will be a Dual Secondary error condition. Click the **Recovery Guru** to resolve this condition (Figure 8-44).

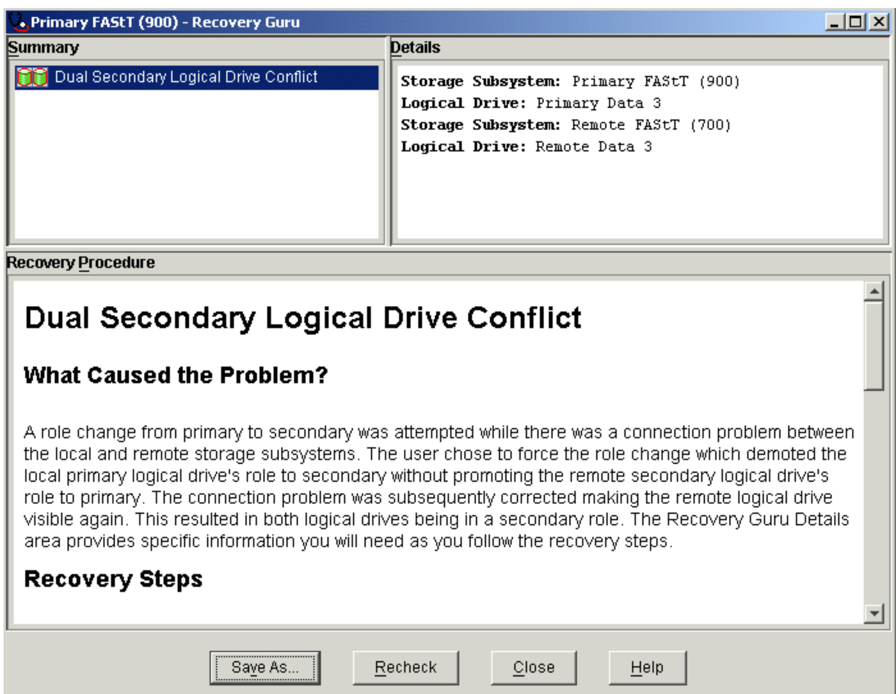

Figure 8-44 Recovery Guru - Dual Secondary Drive Conflict

Once the communication problem between the storage subsystems has been resolved, the Recovery Guru will instruct you to either promote the preferred secondary to primary if the status is synchronized, or remove and re-create the mirror relationship if the status is unsynchronized.

To avoid the Dual Secondary error condition and subsequent recovery steps, wait until the connection between the storage subsystems is operational to perform the role reversal.

8.3.2 Re-establishing Remote Volume Mirroring after disaster recovery

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 "Remote Volume Mirroring: Step-by-step" on page 250.
2. Ensure that storage partitioning is properly defined on the recovered primary site so that it can take over normal operation from the secondary site.
3. Ensure that the host software is properly configured so that the host systems at the recovered primary site can take over I/O from the secondary site host systems.
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 "Changing a primary to a secondary logical drive" on page 273.

The Remote Volume Mirror configuration is now optimal.

8.4 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 solution, 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 mirroring 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 volume 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 all participating storage subsystem.
   – Ensure that the required primary and secondary logical drives are created on the primary and remote storage subsystems.

8.4.1 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 8-45 on page 278. The primary storage subsystem and remote 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.

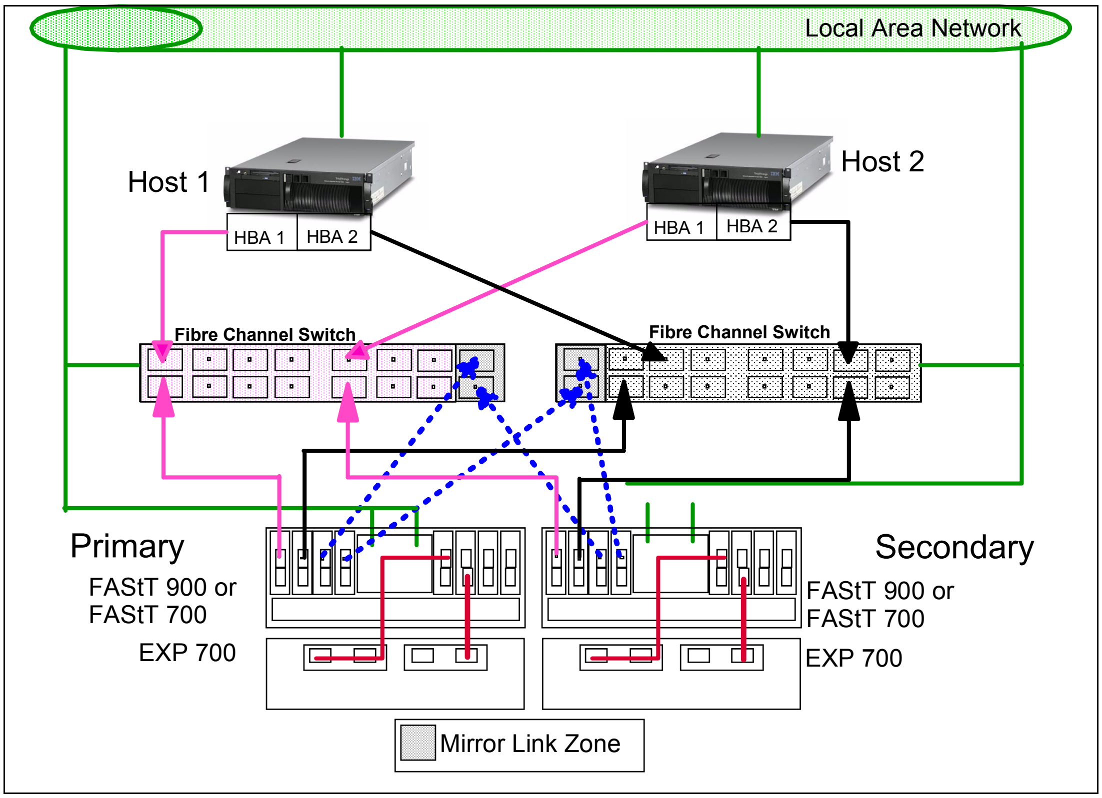

Figure 8-45 Simple departmental solution

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.

8.4.2 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 an Inter Switch Link (ISL) forming a Fibre Channel fabric as shown in Figure 8-46 on page 279. Using standard Fibre Channel long wave Gigabit Interface Converters (GBICs) and single mode cable, the primary storage subsystem and remote storage subsystem have a maximum connection distance of up to 10 km (6.25 mi).

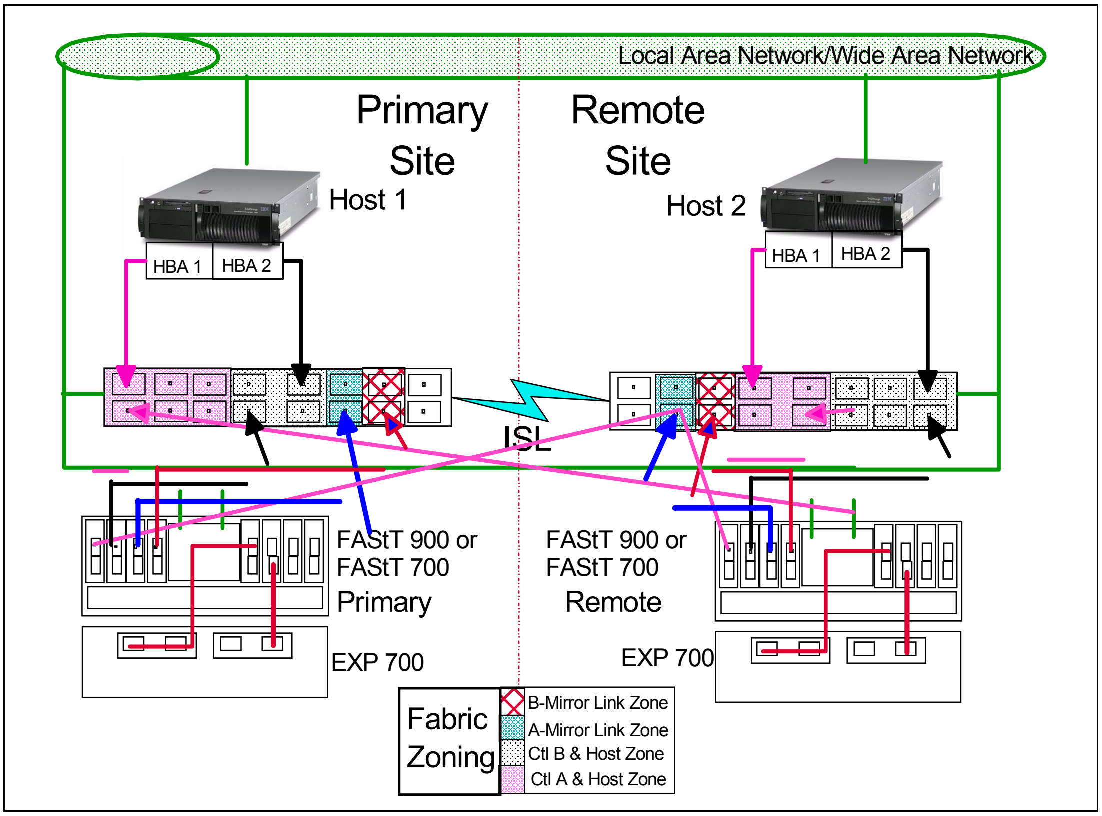

Figure 8-46 Intersite redundant fabric

The configuration provides for redundancy based on a site failover. The fabric has been zoned into four separate zones and each switch shares in all four zones.

We recommend the zone in this solution as a 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 publication, 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.

8.4.3 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 Inter Switch Links (ISLs) forming Fibre Channel fabrics, as shown in Figure 8-47 on page 280. The primary storage subsystem and remote storage subsystem have a maximum connection distance of up to 10 km. (6.25 mi.).

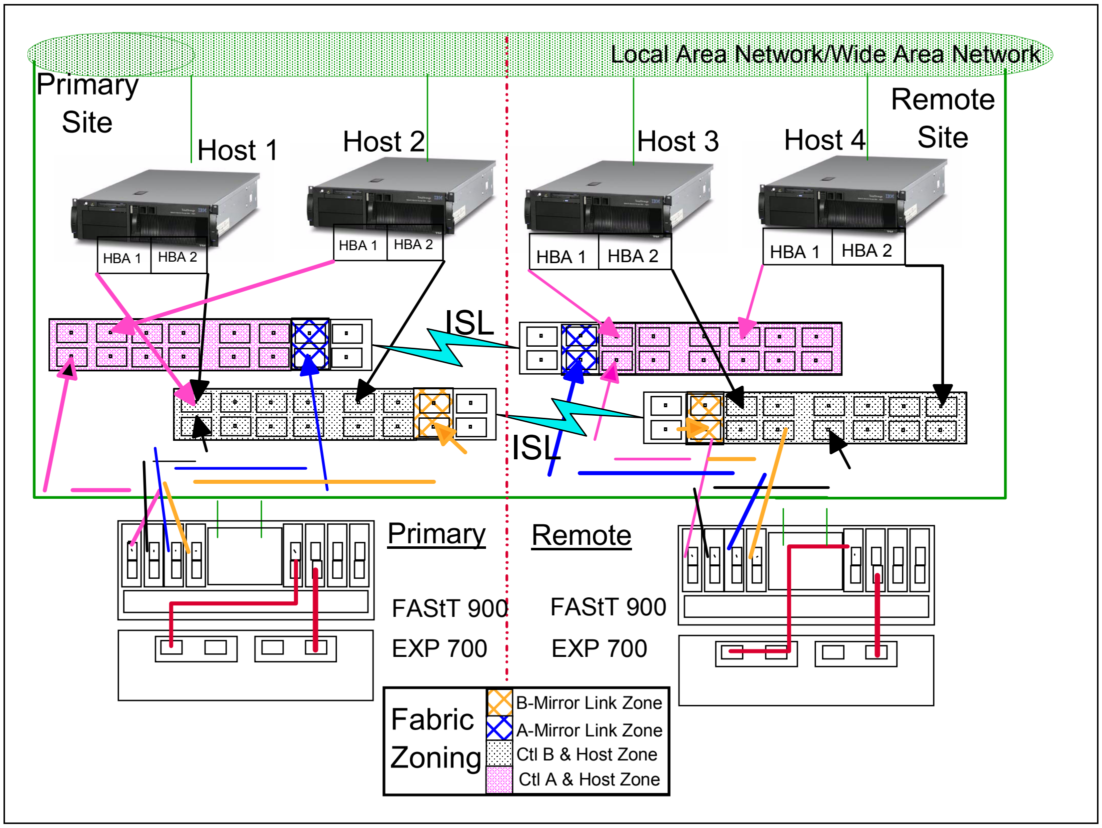

Figure 8-47 Intersite high availability solution

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 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 giving an up-to-date copy at all times, it is still possible through programming or human error to corrupt data and for 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.

9

VolumeCopy

This chapter covers the usage of the VolumeCopy option. This premium feature, new with the FAStT Storage Manager 8.4, is introduced in “VolumeCopy overview” on page 102.

The first part of this chapter presents the various components of VolumeCopy, what they do, and how to set them up.

This is followed by a step-by-step guide on how to use VolumeCopy, first using the Storage Manager client graphical interface, then through the CLI.

9.1 Introduction to VolumeCopy

The VolumeCopy premium feature is used to copy data from one logical drive (the source logical drive) to another logical drive (the target logical drive) in a *single* storage subsystem.

This feature can be used to copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives, to back up data, or to restore FlashCopy logical drive data to the base logical drive. This premium feature includes a Create Copy Wizard, to assist in creating a VolumeCopy; and a Copy Manager, to monitor logical drive (volume) copies after they have been created.

The VolumeCopy premium feature must be enabled by purchasing a feature key file (refer to “Handling premium features” on page 347).

> **Note:** VolumeCopy is the name of the feature; within the SM 8.4 code it is referred to as logical drive copy.

The following sections briefly introduce applications and functions of the VolumeCopy premium feature.

Refer to “VolumeCopy: Step-by-step” on page 287 for detailed “how to” procedures using VolumeCopy.

9.1.1 Copying data for greater access

Some of the advantages and uses of VolumeCopy premium feature include the following.

Migration

As your storage requirements for a logical drive change, the VolumeCopy premium feature can be used to copy data to a logical drive in an array that utilizes larger capacity disk drives within the same storage subsystem. This provides an opportunity to move data to larger drives (for example, 73 GB to 146 GB), change to drives with a higher data transfer rate (for example, 1 Gb/s to 2 Gb/s), or to change to drives using new technologies for higher performance.

Backing up data

The VolumeCopy premium feature allows you to create a backup of a logical drive by copying data from one logical drive to another logical drive in the same storage subsystem. The target logical drive can be used as a backup for the source logical drive, for system testing, or to back up to another device, such as a tape drive.

Restoring FlashCopy logical drive data to the base logical drive

If you need to restore data to the base logical drive from its associated FlashCopy logical drive, the VolumeCopy premium feature can be used to copy the data from the FlashCopy logical drive to the base logical drive. You can create a VolumeCopy of the data on the FlashCopy logical drive, then copy the data to the base logical drive.

> **Note:** If you are using Windows 2000 or Linux, it is recommended that you use the VolumeCopy premium feature with the FlashCopy premium feature to restore FlashCopy logical drive data to the base logical drive. Otherwise, the source logical drive and target logical drive can become inaccessible to the host.

9.1.2 Creating and managing VolumeCopy copies

The VolumeCopy premium feature can be managed using the IBM FAStT Storage Manager 8.4 graphical interface or using the relevant command line interface (CLI) commands and scripts.

Create copy wizard

The IBM FAStT Storage Manager 8.4 introduces the Create Copy Wizard. The Create Copy Wizard guides you through the process of selecting a source logical drive from a list of available logical drives, selecting a target logical drive from a list of available logical drives, and setting the copy priority for the VolumeCopy. When you have completed the Wizard dialogs, the VolumeCopy starts and data is read from the source logical drive and written to the target logical drive.

Operation in Progress icons are displayed on the source logical drive and target logical drive while the VolumeCopy has a status of In Progress or Pending. For a detailed description of how to use the Create Copy Wizard refer to “VolumeCopy: Step-by-step” on page 287.

Refer to “Restrictions” on page 285 for more information about source logical drives and target logical drives.

Copy Manager

After you create a VolumeCopy with the Create Copy Wizard, the VolumeCopy can be monitored and managed through the Copy Manager. From the Copy Manager, a VolumeCopy may be re-copied, stopped, or removed; and its attributes, including the copy priority and the target logical drive Read-Only attribute, can be modified. The status of a VolumeCopy can be viewed in the Copy Manager. Also, if you need to determine what logical drives are involved in a VolumeCopy, use the Copy Manager or the Storage Subsystem Profile. A brief description of the functions of the Copy Manager are as follows:

- Set the target logical drive read-only attribute

  This option determines how read and write requests to the target logical drive are handled after a VolumeCopy is complete or if the VolumeCopy fails prior to completing. After the VolumeCopy is complete, the target logical drive automatically becomes read-only to hosts, and write requests to the target logical drive will not take place. When the read-only attribute for a target logical drive is enabled, a lock icon is displayed in the Target Logical Drive column of the Copy Manager. The Read-Only attribute can be changed in the Copy Manager only after the VolumeCopy is completed.

- Re-copy a logical drive

  This option allows you to create a new VolumeCopy of the data on a selected copy pair. Re-copy can be used to create a new VolumeCopy from the beginning if the original VolumeCopy fails or was stopped. You can also use this option for backup purposes; for instance, if the data on the source logical drive changes, you can use Re-copy to duplicate any new data to the target logical drive.

- Stop a VolumeCopy

  This option can stop a VolumeCopy with a status of Pending, In Progress, or Failed. If you decide not to use a particular logical drive as a source logical drive or target logical drive, you can use this option to stop the VolumeCopy before it completes. The logical drives can then be used in a new VolumeCopy. Using this option on a VolumeCopy with a status of Failed will clear the Needs-Attention condition on the storage subsystem.

- Change copy priority

  This option is used to balance I/O activity with VolumeCopy activity on a storage subsystem. You can set the copy priority to a rate that will have the least impact on I/O activity. There are five copy priority rates available: Lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized and the VolumeCopy will take longer. If the copy priority is set to the highest priority rate, the VolumeCopy is prioritized, but I/O activity for the storage subsystem may be affected.

- Remove Copy Pairs

  This option allows you to remove a VolumeCopy from the Copy Manager. After the VolumeCopy is removed, the source logical drive and target logical drive can be used in a new VolumeCopy. When the VolumeCopy is removed, the Read-Only attribute for the target logical drive is also removed.

For a full description on using the Copy Manager refer to “VolumeCopy: Step-by-step” on page 287.

9.1.3 Understanding VolumeCopy

A VolumeCopy will fail all FlashCopy logical drives associated with the target logical drive, if any exist. If you select a base logical drive of a FlashCopy logical drive, you must disable all FlashCopy logical drives associated with the base logical drive before you can select it as a target logical drive. Otherwise, the base logical drive cannot be used as a target logical drive.

> **Important:** A VolumeCopy overwrites data on the target logical drive and automatically makes the target logical drive *read-only* to hosts.

If there are eight logical drive copies with a status of In Progress, any subsequent VolumeCopy will have a status of Pending until one of the eight logical drive copies completes.

VolumeCopy and modification operations

If a modification operation is running on a source logical drive or target logical drive, and the VolumeCopy has a status of In Progress, Pending, or Failed, then the VolumeCopy will not take place. If a modification operation is running on a source logical drive or target logical drive after a VolumeCopy has been created, then the modification operation must complete before the VolumeCopy can start. If a VolumeCopy has a status of In Progress, then any modification operation will not take place.

Modification operations include:

- Defragmenting an array
- Copyback to a drive that is part of an array
- Initialization of a logical drive
- Dynamic Segment Sizing (DSS) change for a logical drive
- Dynamic Reconstruction Rate (DRR) of a drive that is part of an array
- Dynamic RAID Level Migration (DRM) change for an array
- Dynamic Capacity Expansion (DCE), to increase an array's capacity using unconfigured capacity (in the form of unused drives)
- Dynamic Logical Drive Expansion (DVE), to increase a logical drive's capacity using free capacity available on the array of the standard or FlashCopy repository logical drive

Failed VolumeCopy

A VolumeCopy can fail because of a read error from the source logical drive, a write error to the target logical drive, or because of a failure on the storage subsystem that affects the source logical drive or target logical drive (such as a Remote Mirror role reversal). A critical event is logged in the event log when the VolumeCopy fails, and a Needs-Attention icon is displayed in the Enterprise Management window. While a VolumeCopy has this status, the host has read-only access to the source logical drive, and read and write requests to the target logical drive will not take place until the failure is corrected by using the Recovery Guru.

Preferred controller ownership

During a VolumeCopy, the same controller must own both the source logical drive and target logical drive. If both logical drives do not have the same preferred controller when the VolumeCopy starts, the ownership of the target logical drive is automatically transferred to the preferred controller of the source logical drive. When the VolumeCopy is completed or is stopped, ownership of the target logical drive is restored to its preferred controller. If ownership of the source logical drive is changed during the VolumeCopy, ownership of the target logical drive is also changed.

Failed controller

Controller ownership must be manually changed to the alternate controller to allow the VolumeCopy to complete under the following conditions:

- If a VolumeCopy has a status of In Progress, and
- Its preferred controller fails, and
- The ownership transfer does not occur automatically in the failover.

Restrictions

The following restrictions apply to the source logical drive, target logical drive, and the storage subsystem:

- The source logical drive is available for read I/O activity only while a VolumeCopy has a status of In Progress or Pending. Write requests are allowed after the VolumeCopy is completed.

  > **Tip:** Because all write requests to the source volume are rejected when the VolumeCopy is taking place, it is essential to minimize the time it takes to complete the copy operation. This is why VolumeCopy must always be used in conjunction with FlashCopy; that is make first a FlashCopy of the source logical drive, then perform a VolumeCopy of the FlashCopy (see Figure 4-18 on page 104).

- A logical drive can be used as a target logical drive in only one VolumeCopy at a time.
- The maximum allowable number of logical drive copies per storage subsystem is dependent upon the number of target logical drives available on your storage subsystem.
- A storage subsystem can have up to eight VolumeCopies running at any given time.
- The target logical drive capacity must be equal to or greater than the source logical drive capacity.
- A source logical drive can be a standard logical drive, FlashCopy logical drive, FlashCopy base logical drive, or a Remote Mirror primary logical drive.
- A target logical drive can be a standard logical drive, a base logical drive of a disabled or failed FlashCopy logical drive, or a Remote Mirror primary logical drive.

Important: If you choose a base logical drive of a FlashCopy logical drive as your target logical drive, you must disable all FlashCopy logical drives associated with the base logical drive before you can select it as a target logical drive. Otherwise, the base logical drive cannot be used as a target logical drive.

Logical drives that have the following statuses cannot be used as a source logical drive or target logical drive:

- A logical drive that is reserved by the host cannot be selected as a source or target logical drive.
- A logical drive that is in a modification operation.
- A logical drive that is the source logical drive or target logical drive in another VolumeCopy with a status of Failed, In Progress, or Pending.
- A logical drive with a status of Failed.
- A logical drive with a status of Degraded.

9.1.4 VolumeCopy and performance considerations

During a VolumeCopy, data is read from the source logical drive and written to the target logical drive in the same storage subsystem. Because the VolumeCopy diverts controller processes resources from I/O activity, it can affect I/O activity on the storage subsystem. The copy priority defines how much processing time is allocated for a VolumeCopy versus I/O activity.

VolumeCopy priority rates

Several factors contribute to system performance, including I/O activity, logical drive RAID level, logical drive configuration (number of drives in the array or cache parameters), and logical drive type (FlashCopy logical drives may take more time to copy than standard logical drives).

You can select the copy priority while you are creating a new VolumeCopy, or you can change it later using the Copy Manager. The following copy priority rates are available: Lowest, low, medium, high, highest.

Note: The Lowest priority rate supports I/O activity, but the VolumeCopy will take longer. The Highest priority rate supports the VolumeCopy, but I/O activity may be affected. Remember, however, that VolumeCopy should always be used in conjunction with FlashCopy.

Viewing the progress of a copy operation

The progress of the VolumeCopy is displayed in the status bar at the bottom-right of the Copy Manager. The status bar shows an estimate of the time remaining for the VolumeCopy to complete. You can view the progress of a VolumeCopy in the Copy Manager only when a VolumeCopy has a status of In Progress.

9.2 VolumeCopy: Step-by-step

This section presents an easy-to-follow, step-by-step procedural guide to common administration tasks for a FAStT VolumeCopy. The tasks that are covered include:

- Checking the status of the VolumeCopy premium feature
- Creating a VolumeCopy pair
- Viewing Logical Copy Drive Properties
- Using the Copy Manager
- Changing the Copy priority
- Stopping a Logical Copy
- Re-copying a Logical Copy Drive
- Setting the Target Logical Copy Drive Read-Only Attributes
- Removing Logical Copy pairs
- Viewing the Storage Subsystem Profile

9.2.1 Checking the status of the VolumeCopy premium feature

Use this procedure to view a list of premium features on the storage subsystem and to verify that the VolumeCopy premium 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: From the Subsystem Management window, click **Storage Subsystem -> Premium Features -> List**; the List Premium Features dialog opens as shown in Figure 9-1. It lists the following items:

- Premium features enabled/disabled on the storage subsystem
- Feature Enable Identifier

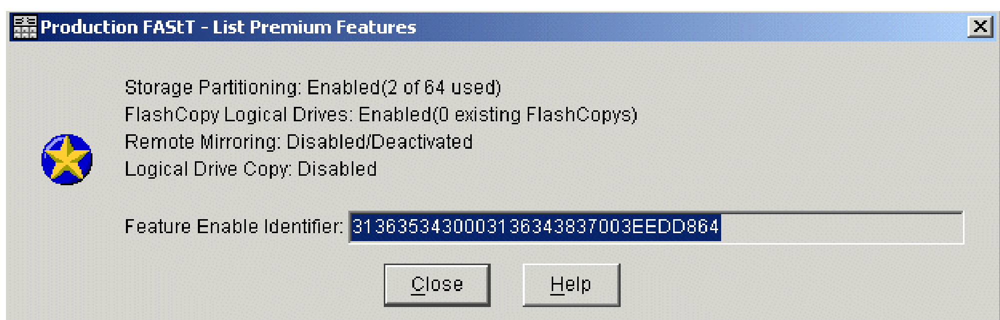

Figure 9-1 List Premium Features

You can also view the Premium Feature icon in the lower left of the Subsystem Management window (Figure 9-2).

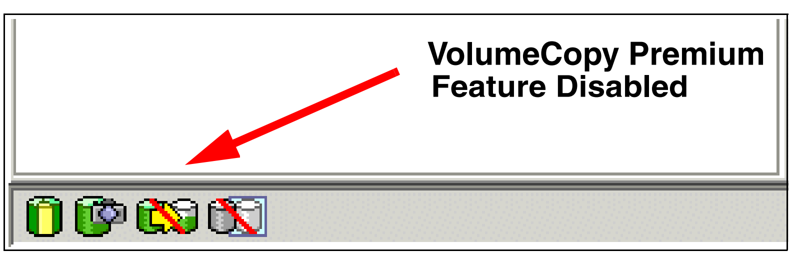

Figure 9-2 Logical Copy Feature icon - Disabled

9.2.2 Enabling the VolumeCopy premium feature

From the Subsystem Management window, click **Storage Subsystem -> Premium Features -> Enable**. The Select Feature Key File window opens. Chose the directory where you stored the *.key file, then proceed (for details on how to obtain the key file, refer to “Handling premium features” on page 347). The Enable Premium Feature confirmation dialog appears (Figure 9-3). Click **Yes** to confirm.

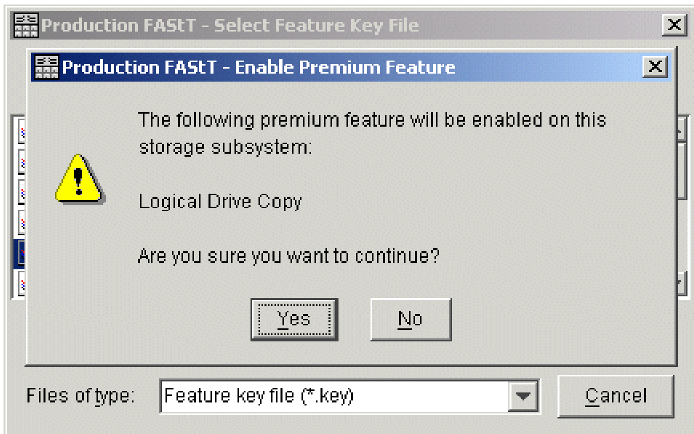

Figure 9-3 Enable VolumeCopy Premium Feature - Confirmation

To check the VolumeCopy premium feature is enabled, look at the icon in the lower left corner of the Subsystem Management window (the slash is removed).

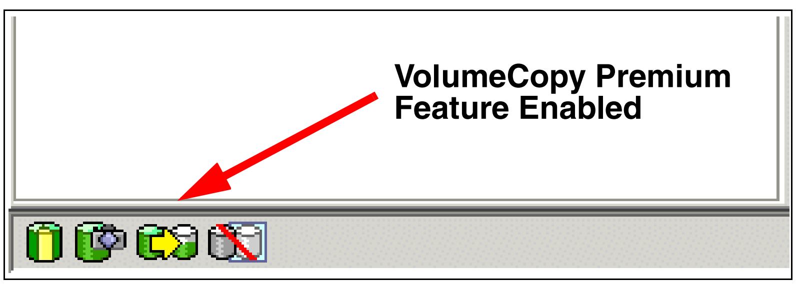

Figure 9-4 Logical Copy Feature icon - Enabled

9.2.3 Creating a VolumeCopy pair

This section describes in detail how to create VolumeCopies using the Create Copy Wizard from the Subsystem Management console. The Wizard will guide you through the VolumeCopy creation process.

> **Important:** Remember that it is recommended to first make a FlashCopy of the logical drive, and then use the FlashCopy as the source of the VolumeCopy.

A VolumeCopy will:

- Overwrite data on the target logical drive.
- Automatically make the target logical drive read-only to hosts. You may want to keep this attribute enabled in order to preserve the data on the target logical drive. For example, if you are using the target logical drive for backup purposes, if you are copying data from one array to a larger array for greater accessibility, or if you are using the data on the target

logical drive to copy back to the base logical drive of a disabled or failed FlashCopy logical drive. If you decide not to preserve the data on the target logical drive after the VolumeCopy is completed, use the Copy Manager to disable the Read-Only attribute for the target logical drive.

Note: If you are using Windows 2000 or Linux, it is recommended that you use the VolumeCopy premium feature with the FlashCopy premium feature to restore FlashCopy logical drive data to the base logical drive. Otherwise, the source logical drive and target logical drive can become inaccessible to the host.

To launch the Create Copy Wizard use the following procedure:

1. Stop all I/O activity to the source logical drive and target logical drive.
2. Unmount any file systems on the source logical drive and target logical drive.
3. Select the source logical drive in the Logical View of the Subsystem Management window.
4. Select either the **Logical Drive -> Copy -> Create** pull-down menu option, or **Create Copy** from the right-mouse pop-up menu.

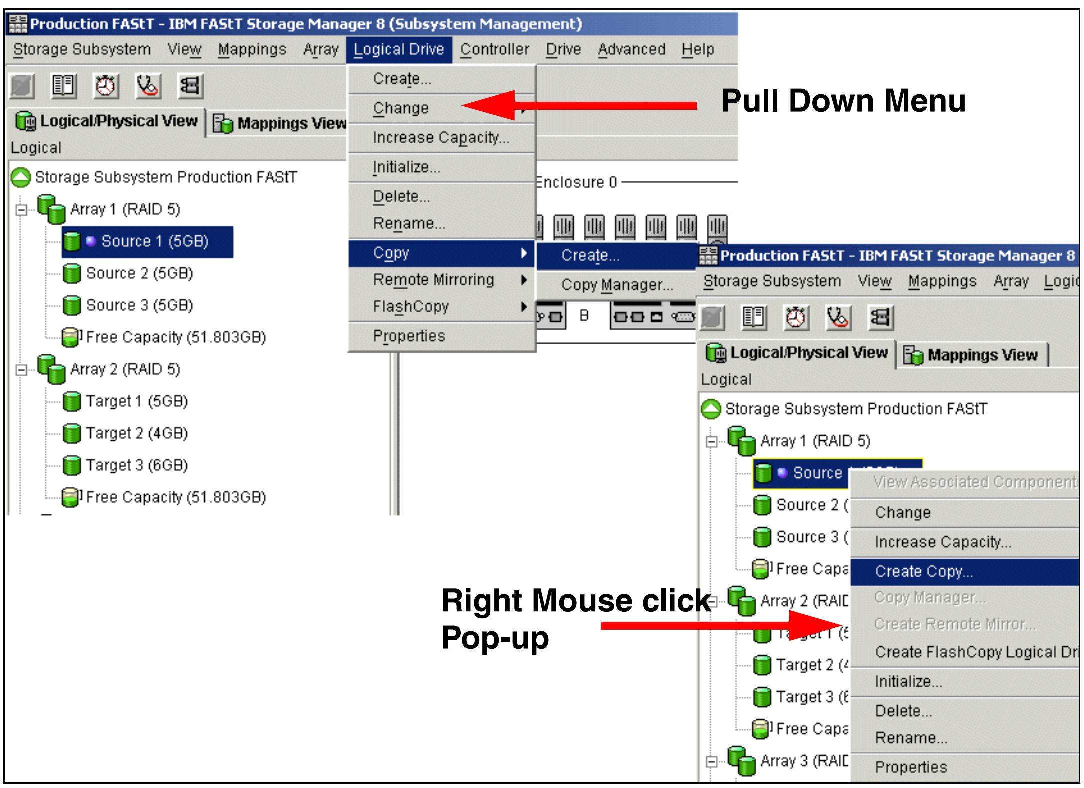

Figure 9-5 Launching Create Copy Wizard from Subsystem Management window

The Create Copy Wizard - Introduction dialog is displayed. The logical drive you selected to start the Wizard is highlighted in the table if it is a valid source logical drive. Follow the directions on each wizard dialog and select **Next** when you are ready to move to the next wizard dialog.

Note: Each wizard dialog has context-sensitive help. Select **Help** to receive help applicable for that particular dialog.

Operation in progress icons are displayed on the source logical drive and target logical drive while the VolumeCopy has a status of In Progress or Pending.

Create Copy wizard - introduction

The first dialog of the Create Copy Wizard defines a VolumeCopy and allows you to choose a source logical drive (Figure 9-6).

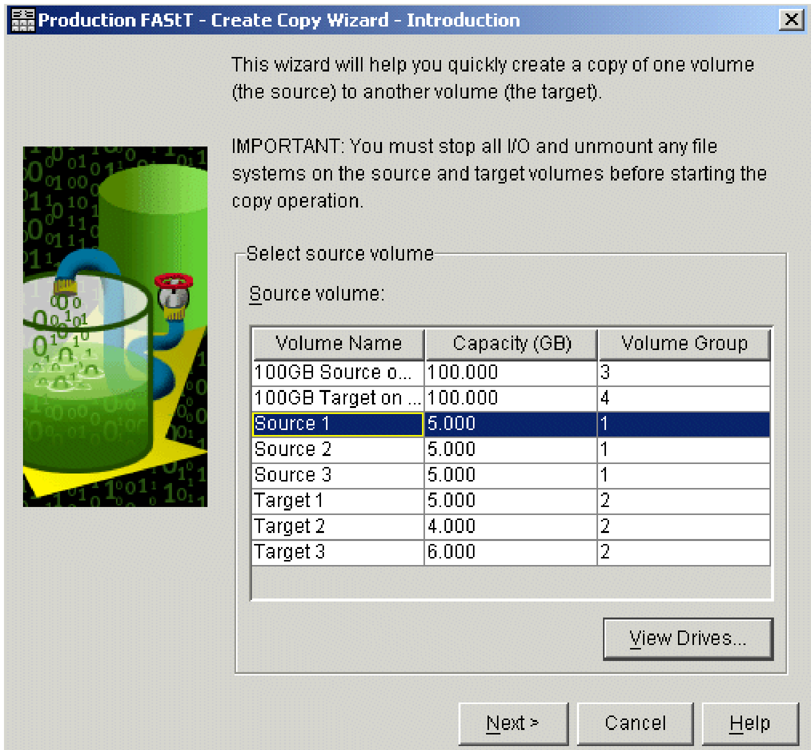

Figure 9-6 Create Copy Wizard - Introduction

The source logical drive is available for read I/O requests only while a VolumeCopy has a status of In Progress, Pending, or Failed. Write requests are allowed after the VolumeCopy is completed. All valid source logical drives on the storage subsystem are displayed in the list in alphanumeric order; the logical drive you selected to start the wizard is highlighted.

You can also create a source logical drive specifically for the VolumeCopy.

Highlight the preferred source logical drive and click **Next**.

Note: The following logical drives are not valid source logical drives and will not appear on the list:

- Remote Mirror secondary logical drive
- Remote Mirror repository logical drive
- FlashCopy repository logical drive
- Failed logical drive
- Missing logical drive
- A logical drive currently in a modification operation
- A logical drive holding a legacy or persistent reservation
- A logical drive that is a source logical drive or target logical drive in another VolumeCopy with a status of In Progress, Failed, or Pending.

Create Copy Wizard - Select target and set copy priority

After successfully selecting the source logical drive, the Select Target and Set Copy Priority window appears (Figure 9-7). This window allows you to select the Target Logical Drive and set the copy priority for the VolumeCopy.

All valid target logical drives on the storage subsystem are displayed in the list.

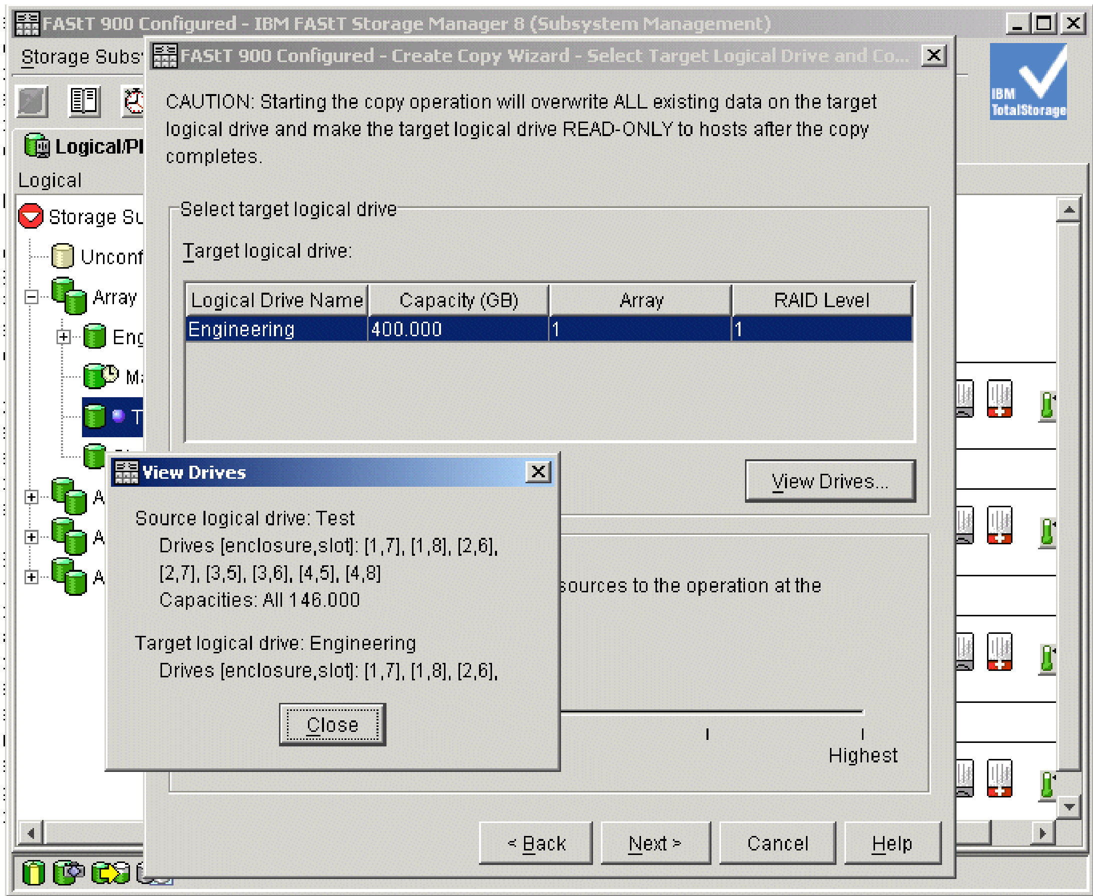

Figure 9-7 Create Copy Wizard - Select target volume and copy priority

Remember that a VolumeCopy will overwrite data on the target logical drive and automatically make the target logical drive read-only to hosts. After the VolumeCopy completes, use the Copy Manager to disable the Read-Only attribute for the target logical drive. If you have used the target logical drive in a VolumeCopy before, ensure that you no longer need that data or have backed it up.

Tip: After you have selected or created a target logical drive, give it a unique name so that it is easily recognizable in the Logical View. For example, if the source logical drive name is Accounting, call the target logical drive Accounting-Copy. That way you can quickly identify the source logical drives and target logical drives available on the storage subsystem.

Note: The following logical drives are not valid target logical drives and will not appear on the list:

- Remote Mirror secondary logical drive
- Remote Mirror repository logical drive
- FlashCopy logical drive
- FlashCopy repository logical drive
- Base logical drive of an optimal FlashCopy logical drive
- Failed logical drive
- Missing logical drive
- A logical drive with a status of Degraded
- A logical drive currently in a modification operation
- A logical drive holding a legacy or persistent reservation
- A logical drive that is a source logical drive or target logical drive in another VolumeCopy that has a status of In Progress, Failed, or Pending

The View Drives dialog can be launched while in the Create Copy Wizard - Introduction or Select Target and Copy Priority windows. See Figure 9-7 on page 291. From the View Drives pop-up you can view drive, slot, and capacity of the source and target logical drives. This information can be used to copy high-access logical drives to different arrays for increased accessibility, or to copy logical drives to an array that uses larger capacity drives.

Once you have selected the target logical drive and set the copy priority, click **Next** to continue.

Create Copy wizard - preview

The next window to appear is the Create Copy Wizard - Preview window (see Figure 9-8). This provides an overview of the selected source logical drive, target logical drive, and copy priority for the VolumeCopy, and allows you to return to previous dialogs to edit the parameters.

Type `Yes` to confirm the operation and click **Finish** to start the copy operation.

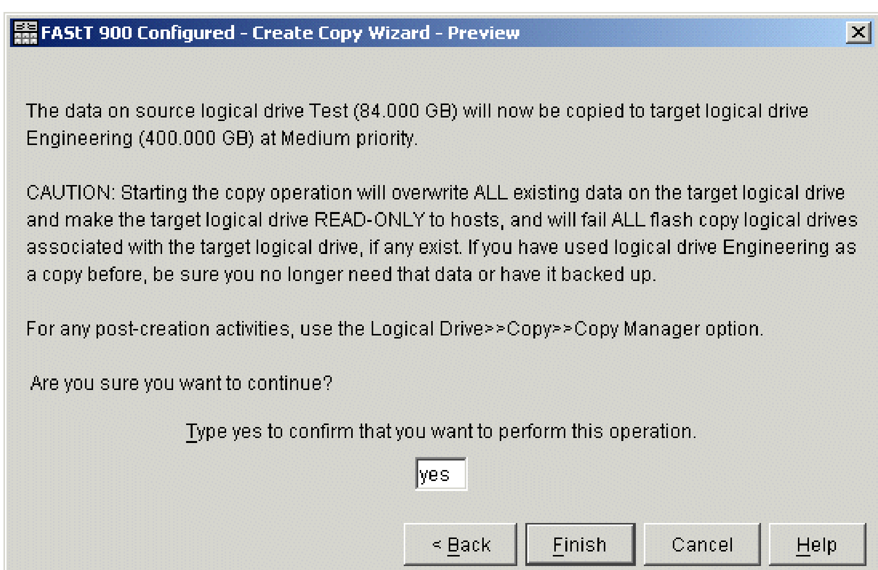

Figure 9-8 Create Copy Wizard - Preview

When you select **Finish**, the Copy Started dialog (Figure 9-9) is displayed to verify that the VolumeCopy has begun.

You have the option to create a new VolumeCopy or exit the wizard.

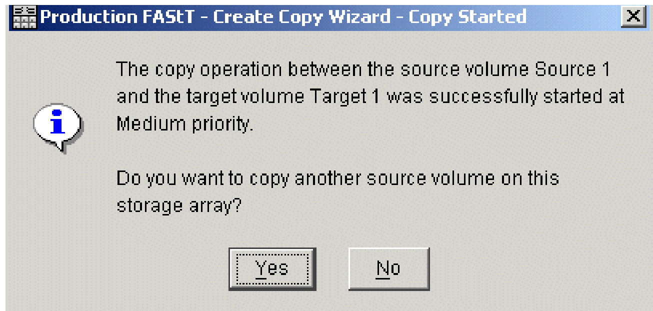

Figure 9-9 Create Copy Wizard - Copy Started message

To close the Copy Started dialog, you must select one of the following:

- Yes - Create a new VolumeCopy.
- No - Exit the wizard.

Operation in progress icons are displayed on the source logical drive and target logical drive while the VolumeCopy has a status of In Progress or Pending.

9.2.4 Viewing VolumeCopy properties

Use the Logical Drive Properties dialog to see information about a selected source logical drive or target logical drive, such as the copy status, copy priority, completion timestamp, and whether the Read-Only attribute is enabled on the target logical drive.

This dialog displays the logical drive's role (source or target), copy status, copy priority, start timestamp, the read-only attribute for the target logical drive, and the world-wide name for the source or target logical drive.

You can only view the progress of a VolumeCopy with a status of In Progress. The procedure to view VolumeCopy properties is as follows:

- Select the source logical drive or target logical drive in the Logical View.
- Select either the **Logical Drive -> Properties** pull-down menu option, or **Properties** from the right-mouse pop-up menu (Figure 9-10 on page 294).

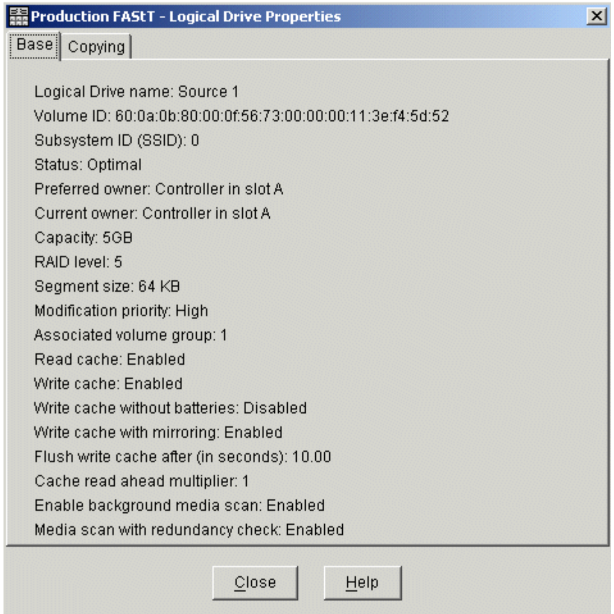

Figure 9-10 Logical Drive Properties - Base tab

- ► If the logical volume you have selected is part of a VolumeCopy relationship, then the Copying tab is present. Select the **Copying** tab (Figure 9-11).

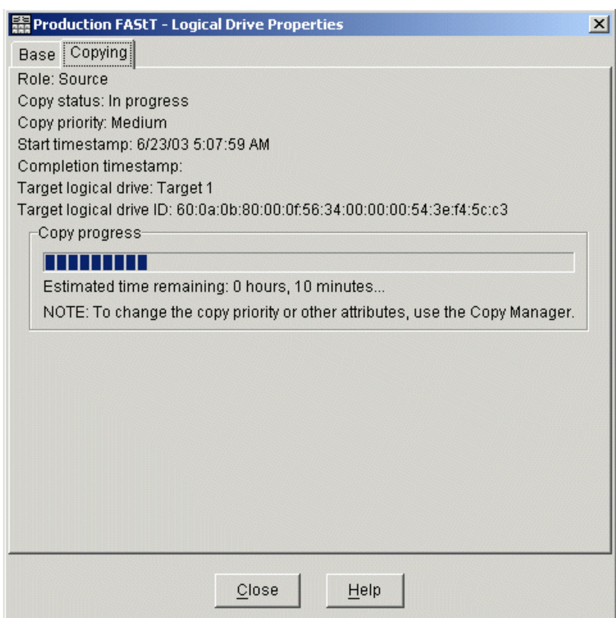

Figure 9-11 Logical Drive Properties - Copying tab

The Copy Properties tab is displayed. The logical drive properties displayed include:

- – Role: Source or Target
- – Copy status: Pending, In Progress, Completed, Stopped, Failed
- – Copy priority: Lowest, Low, Medium, High, Highest
- – Start timestamp: MM/DD/YY HH/MM/SS
- – Completion timestamp: MM/DD/YY HH/MM/SS
- – Read-Only: Enabled/Disabled (only if target logical drive is selected)

– Source/Target Logical Drive: Logical Drive Name
– Source/Target Logical Drive ID: WWN (World Wide Name)

If a source logical drive has multiple target logical drives, then the details will be repeated for each target logical drive.

- Select to exit the dialog.

Tip: If a Logical Drive is part of a VolumeCopy relationship, by right clicking the logical drive, the Go To Source/Target Logical Drive menu item is available. This is a convenient way to find to the copy pair partner. See Figure 9-12 on page 295.

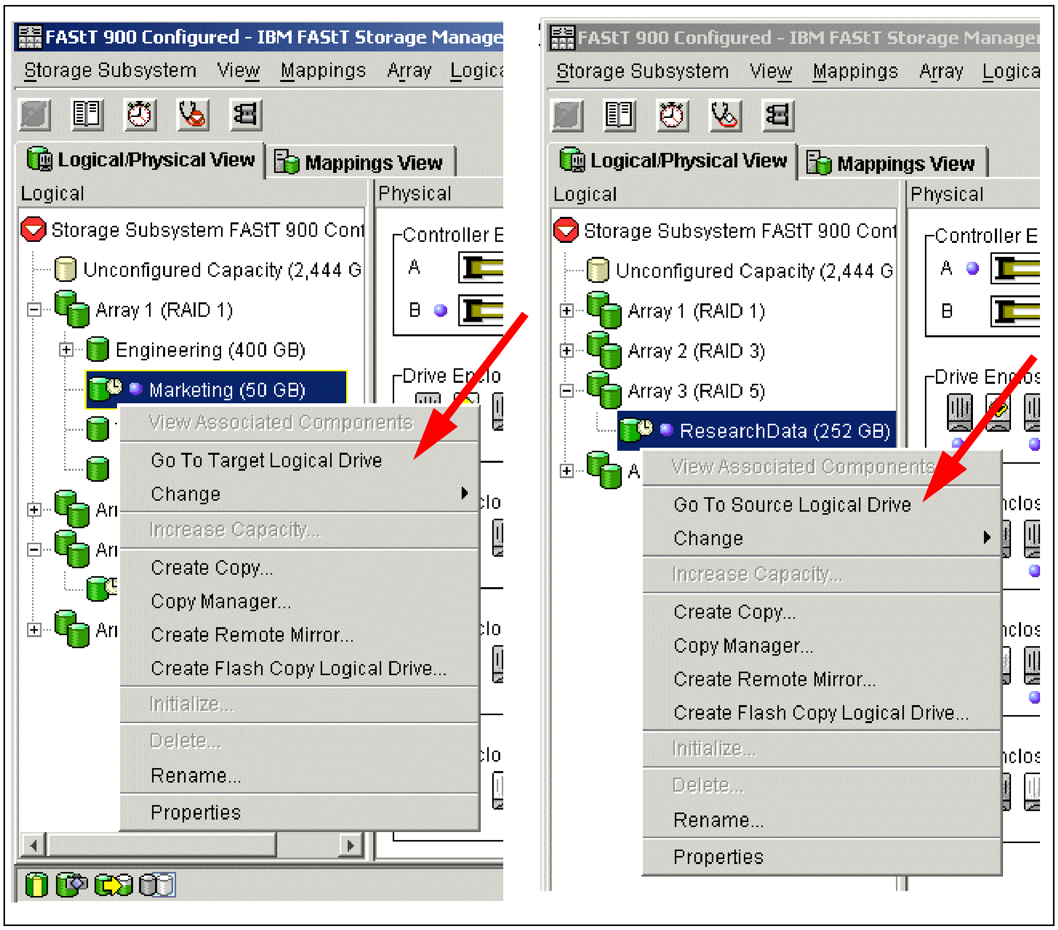

Figure 9-12 Go-to target or source logical drive

9.2.5 Using the Copy Manager

The Copy Manager dialog is used to monitor the progress of a VolumeCopy and perform the following tasks for all logical drive copies on the storage subsystem:

- Re-copy a logical drive.
- Change VolumeCopy priority.
- Set the Read-Only attribute for a target logical drive.
- Stop VolumeCopy.
- Remove copy pairs.

Information for all of the logical drive copies on the storage subsystem are displayed, including the source logical drive, the target logical drive, the status of the VolumeCopy, a timestamp for any completed copies, and the VolumeCopy priority assigned to the VolumeCopy. If a target logical drive is read-only to hosts, a lock icon is displayed in the Target Logical Drive column. You can view the progress of a VolumeCopy in the Status column.

The progress of the VolumeCopy is displayed in the status bar at the bottom-right of the Copy Manager. The status bar shows an estimate of the time remaining for the VolumeCopy to complete. You can view the progress of a VolumeCopy in the Copy Manager only when a VolumeCopy has a status of In Progress.

The procedure to launch the Copy Manager window is as follows:

1. Select either the **Logical Drive -> Copy -> Copy Manager** pull-down menu option, or **Copy Manager** from the right-mouse pop-up menu (Figure 9-13).

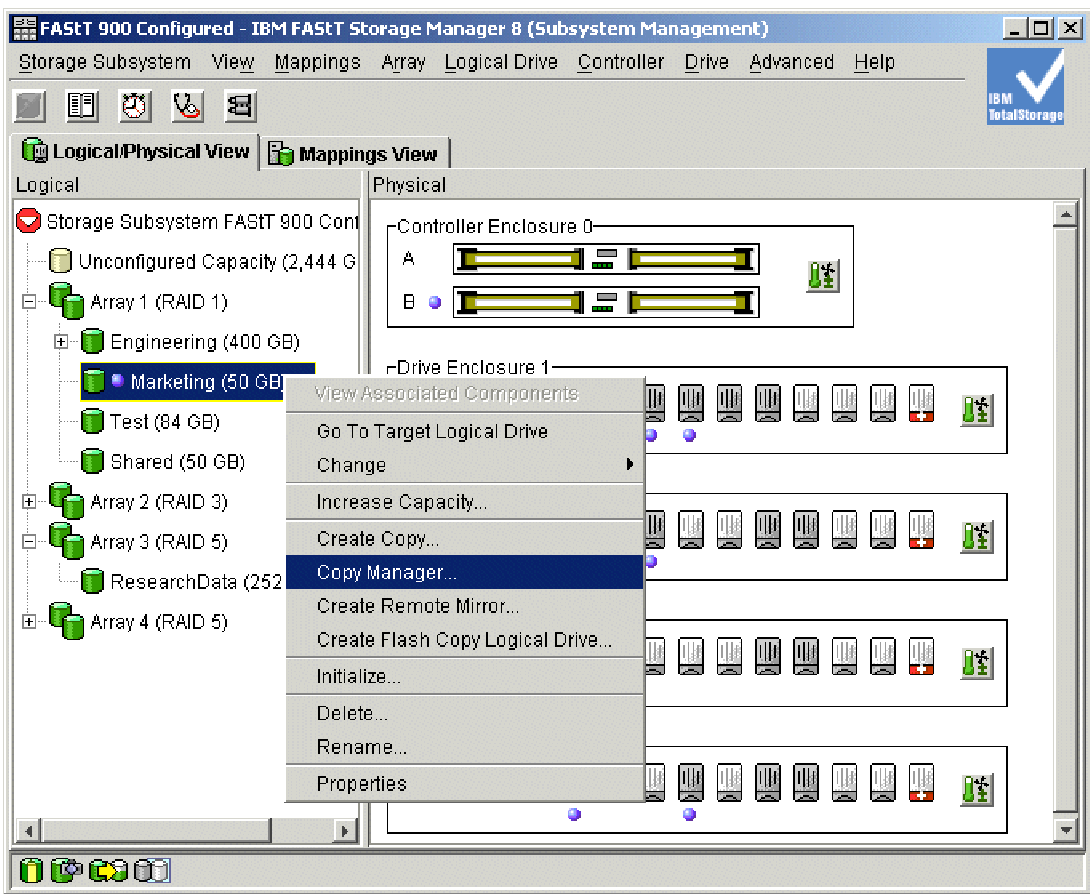

Figure 9-13 Launching the Copy Manager

2. The Copy Manger window is displayed (Figure 9-14 on page 297).

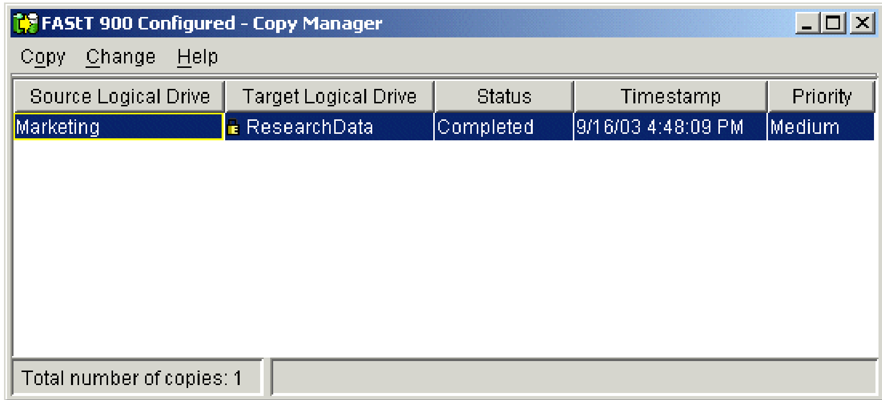

Figure 9-14 Copy Manager

3. Select the copy pair in the table by clicking the copy pair or pressing Enter.
4. Select an option from the pull-down menu bar or from the right-mouse pop-up menu.
5. Close the Copy Manager.

9.2.6 Re-Copy a logical drive

Use this option in the Copy Manager to create a new VolumeCopy for a selected source logical drive and target logical drive (copy pair). This option is used when you have stopped a VolumeCopy and want to start it again or when a VolumeCopy has failed or completed. The VolumeCopy will start over from the beginning.

Note: The Re-Copy option will:

- Overwrite existing data on the target logical drive.
- Make the target logical drive read-only to hosts until you disable the read-only attribute in the Copy Manager.

Important: The following information is important.

- If hosts have been mapped to the source logical drive, the data that will be copied to the target logical drive when you re-copy may have changed since the previous VolumeCopy was created.
- To use this option, select only one VolumeCopy in the Copy Manager.
- A Remote Mirror secondary logical drive cannot be used as a source logical drive or target logical drive.
- A logical drive currently in a modification operation cannot be used as a source logical drive or target logical drive.
- A logical drive that has a status of Degraded cannot be used as a target logical drive.
- A logical drive that has a status of Failed cannot be used as a source logical drive or target logical drive.

The procedure to re-copy a VolumeCopy is as follows:

1. Stop all I/O to the source logical drive and target logical drive.
2. Unmount any file systems on the source logical drive and target logical drive.

3. Select either the **Logical Drive -> Copy -> Copy_Manager** pull-down menu option, or **Copy Manager** from the right-mouse pop-up menu. The Copy Manager window is displayed. as shown in Figure 9-15.

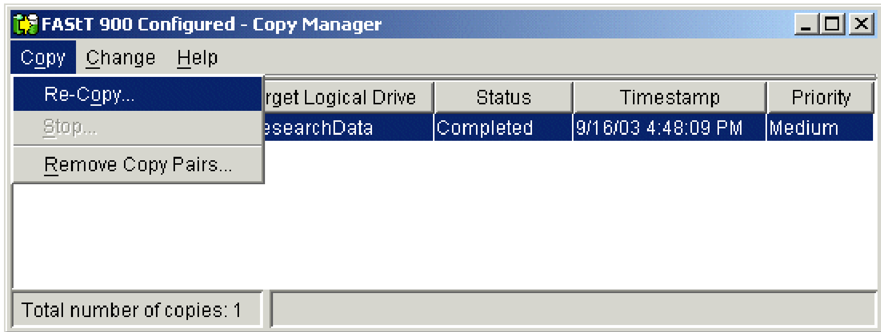

Figure 9-15 Copy Manager - Re-Copy menu

4. Select the copy pair in the table by clicking the copy pair or pressing Enter.
5. Select either the **Copy -> Re-Copy** pull-down menu option, or **Re-Copy** from the right-mouse pop-up menu. The Re-Copy dialog is displayed (Figure 9-16).

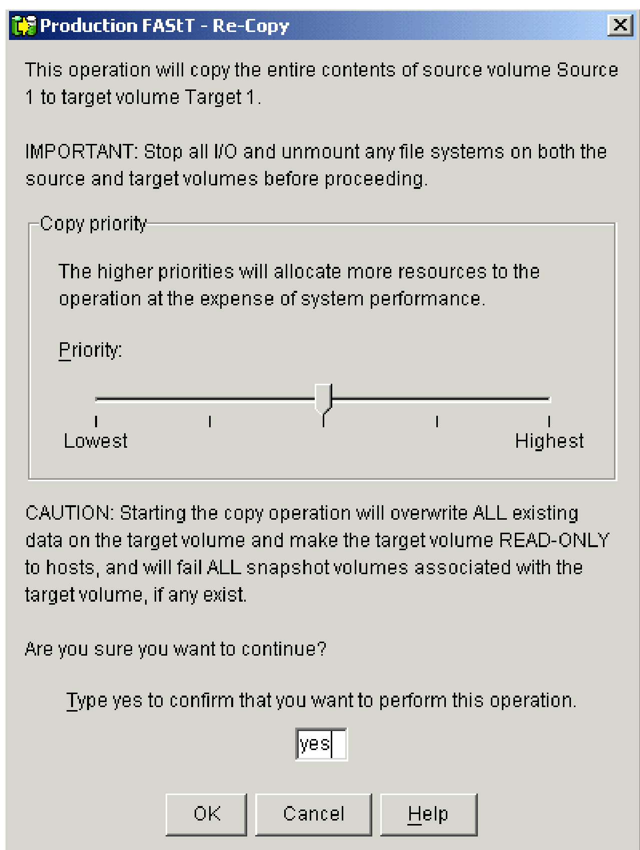

Figure 9-16 Re-copy window

6. Set the copy priority.

Note: There are five copy priority rates available: Lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized and the VolumeCopy will take longer. If the copy priority is set to the highest priority rate, the VolumeCopy is prioritized, but I/O activity for the storage subsystem may be affected.

7. Type `Yes` and select **OK**. The VolumeCopy begins.

Note: You can view the progress of a VolumeCopy in the Copy Manager only when a VolumeCopy has a status of In Progress.

8. Close the Copy Manager. Operation in progress icons are displayed on the source logical drive and target logical drive while the VolumeCopy has a status of In Progress or Pending.

9.2.7 Change VolumeCopy priority

Use this option in the Copy Manager to select the rate at which a VolumeCopy completes for a selected copy pair. You can change the copy priority for a copy pair before the VolumeCopy begins, while the VolumeCopy has a status of In Progress, or after the VolumeCopy has completed (for a re-copy). Change to VolumeCopy priority as follows:

1. Select either the **Logical Drive -> Copy -> Copy Manager** pull-down menu option, or **Copy Manager** from the right-mouse pop-up menu. The Copy Manager dialog is displayed.
2. Select one or more copy pairs in the table. You can select more than one copy pair by pressing Ctrl + the left mouse button.
3. Select **Change -> Copy Priority** from the pull-down menu option or the right-mouse pop-up menu (Figure 9-17).

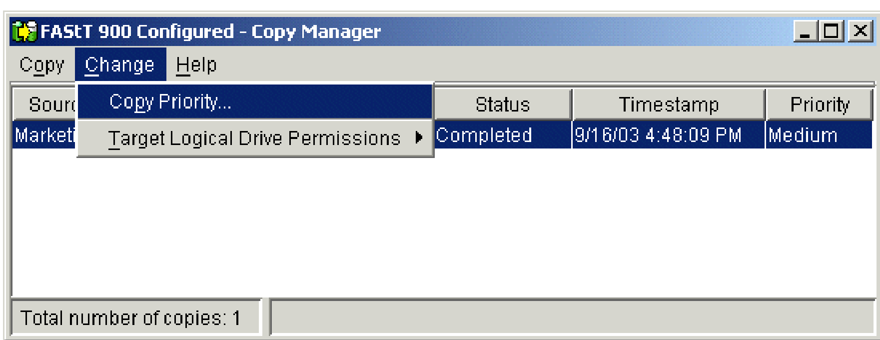

Figure 9-17 Select Copy Priority from Copy Manager

The Change Copy Priority window is displayed (Figure 9-18 on page 300).

4. In the Copy Priority section, use the slider bar to select the appropriate copy priority, depending on your system performance needs.
5. Select **OK**. The copy priority for the selected copy pair is changed.
6. Close the Copy Manager.

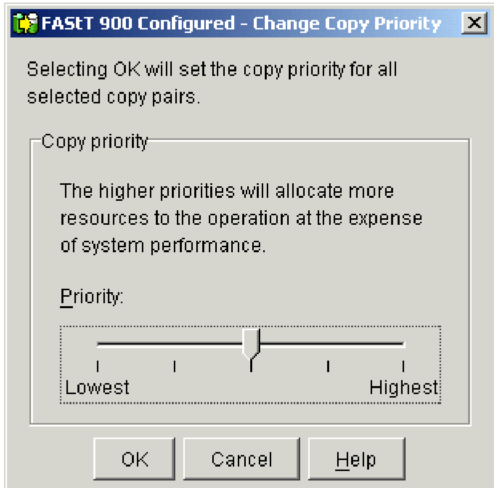

Figure 9-18 Change Copy Priority

> **Note:** There are five copy priority rates available: Lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized and the VolumeCopy will take longer. If the copy priority is set to the highest priority rate, the VolumeCopy is prioritized, but I/O activity for the storage subsystem may be affected.

9.2.8 Set the Read-Only attribute for a target logical drive

Read and write requests to the target logical drive will not take place while the VolumeCopy has a status of Pending or In Progress, or if the VolumeCopy fails prior to completing.

After the VolumeCopy has completed, the target logical drive automatically becomes read-only to hosts. Figure 9-19 shows target logical drives with the Read-Only attribute set to enabled (small lock icon next to volume name).

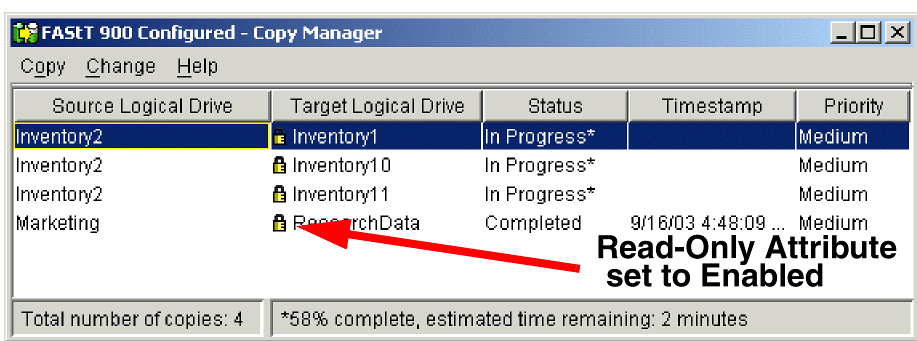

Figure 9-19 Copy Manager showing Read-Only attribute icon enabled

You may want to keep this attribute enabled in order to preserve the data on the target logical drive; for example, if you are using the target logical drive for backup purposes, if you are copying data from one array to a larger array for greater accessibility, or if you are using the data on the target logical drive to copy back to the base logical drive of a disabled or failed FlashCopy logical drive.

If you decide not to preserve the data on the target logical drive after the VolumeCopy is completed, use the Copy Manager to disable the Read-Only attribute for the target logical drive.

How to disable the Read-Only attribute

To disable the Read-Only attribute:

1. Select either the **Logical Drive -> Copy -> Copy_Manager** pull-down menu option, or select **Copy Manager** from the right-mouse pop-up menu.
2. Select one or more copy pairs in the table. You can select more than one copy pair by pressing Ctrl + the left mouse button.
3. Select either the **Change -> Target Logical Drive Permissions -> Disable Read-Only** pull-down menu option, or **Disable Read-Only** from the right-mouse pop-up menu (Figure 9-20).

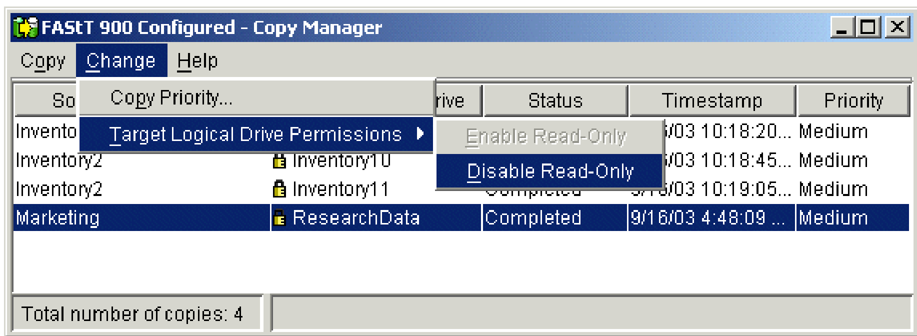

Figure 9-20 Copy Manager - Disable Read-Only menu

The Read-Only attribute is disabled on the selected target logical drive(s). Write requests to the target logical drive are permitted. Figure 9-21 shows target logical drives with the Read-Only attribute set to enabled (small lock icon next to volume name has been removed).

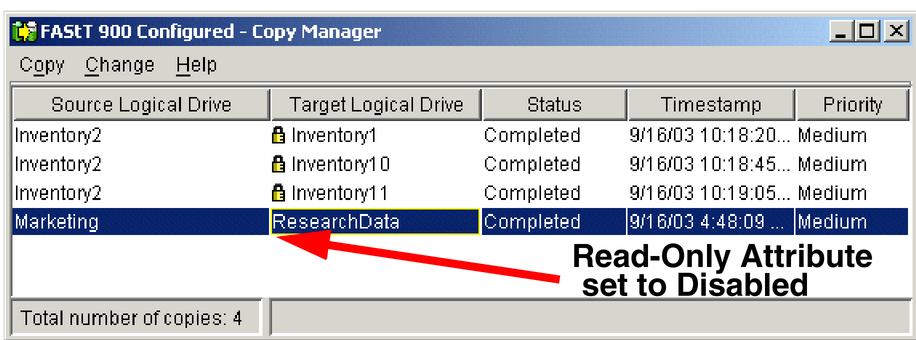

Figure 9-21 Copy Manager showing Read-Only attribute icon disabled

4. Close the Copy Manager.

9.2.9 Stopping VolumeCopy

This option in the Copy Manager is used to stop a VolumeCopy with a status of In Progress, Pending, or Failed.

Note: Using this option on a VolumeCopy with a status of Failed clears the Needs-Attention condition on a storage subsystem.

Important: The following information is important:

- To use this option, select only one copy pair in the Copy Manager.
- When the VolumeCopy is stopped, all mapped hosts will have write access to the source logical drive. If data is written to the source logical drive, the data on the target logical drive will no longer match the data on the source logical drive.

The procedure to stop a VolumeCopy is as follows:

1. 1 Select either the **Logical Drive -> Copy -> Copy Manager** pull-down menu option, or **Copy Manager** in the right-mouse pop-up menu to display the Copy Manager.
2. Highlight the copy pair in the table by clicking the copy pair.
3. Select either the **Copy -> Stop** pull-down menu option, or **Stop Copy** from the right-mouse pop-up menu (Figure 9-22).

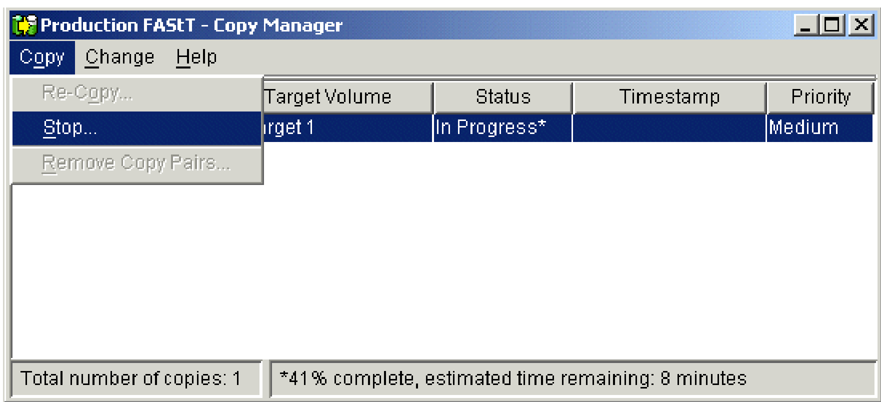

Figure 9-22 Copy Manager - Stop Copy menu

4. The Stop Copy dialog is displayed (Figure 9-23 on page 303). Read the information. Click **Yes** to proceed. The VolumeCopy is stopped.

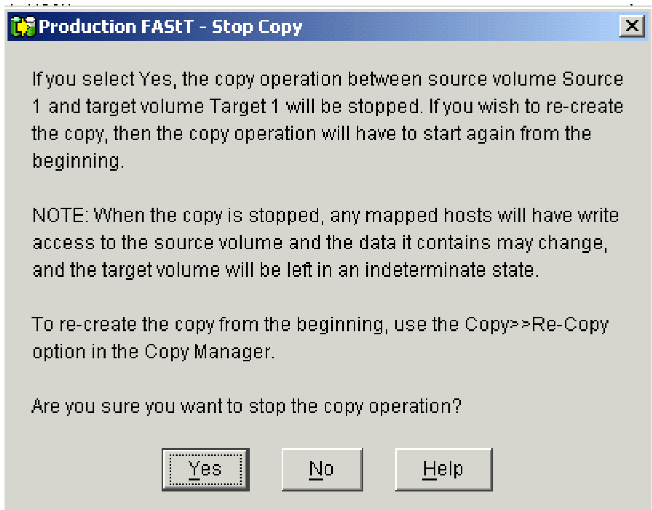

Figure 9-23 Stop Copy dialog

5. Close the Copy Manager.

9.2.10 Remove copy pairs

Use this option to remove one or more logical drive copies from the Copy Manager. Any VolumeCopy-related information for the source and target logical drive is removed from the Logical Drive Properties (no copy tab) and Storage Subsystem Profile dialogs. When you remove a VolumeCopy from the storage subsystem, the Read-Only attribute for the target logical drive is also removed.

After the VolumeCopy is removed from the Copy Manager, the target logical drive can be selected as a source logical drive or target logical drive for a new VolumeCopy.

If you remove a VolumeCopy, the source logical drive and target logical drive are no longer displayed in the Copy Manager.

Important: The following information is important:

- This option does not delete the data on the source logical drive or target logical drive.
- If the VolumeCopy has a status of In Progress, it must be stopped before you can remove the copy pair from the Copy Manager.

The procedure for removing VolumeCopy pairs is as follows:

1. Select either the **Logical Drive -> Copy -> Copy_Manager** pull-down menu option, or select **Copy Manager** from the right-mouse pop-up menu. The Copy Manager dialog is displayed.
2. Select one or more copy pairs in the table. You can select more than one copy pair by pressing Ctrl + the left mouse button.
3. Select either the **Copy -> Remove Copy Pairs** pull-down menu option, or **Remove Copy Pairs** from the right-mouse pop-up menu (Figure 9-24 on page 304).

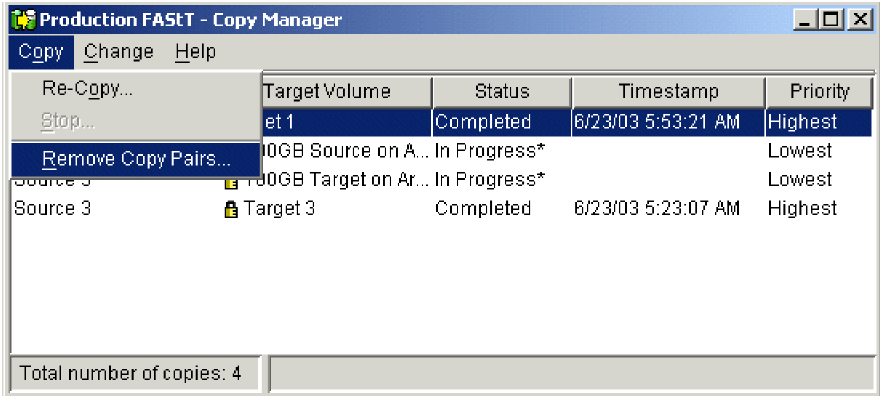

Figure 9-24 Copy Manager - Remove Copy Pairs menu

4. The Remove Copy Pairs confirmation screen is displayed. Read the information. Select **Yes** to proceed. The VolumeCopy is removed.
5. Close the Copy Manager.

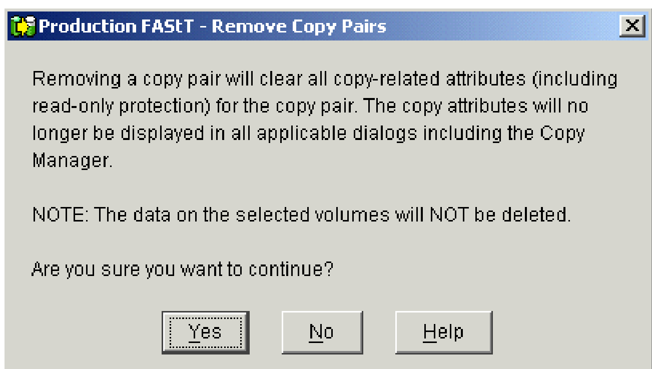

Figure 9-25 Remove Copy Pairs - Confirmation

9.2.11 Viewing the VolumeCopy in the storage subsystem profile

The storage subsystem profile provides an excellent source of information about logical drive (VolumeCopy) copies that maybe be active on the storage subsystem.

To view the storage subsystem profile:

1. Select a storage subsystem in the Subsystem Management window.
2. Select either the **View -> Storage Subsystem Profile** pull-down menu option, or **View Profile** from the right-mouse pop-up menu. See Figure 9-26 on page 305.

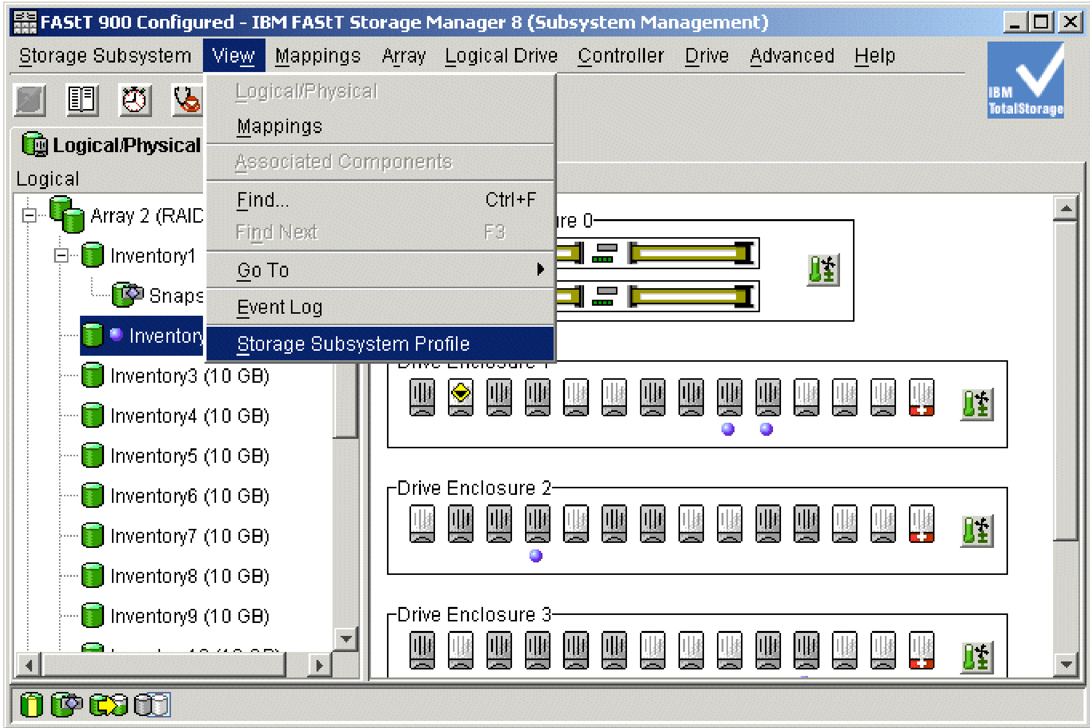

Figure 9-26 Launching the Storage Subsystem Profile

3. The Storage Subsystem Profile dialog is displayed (Figure 9-27). The Storage Subsystem Profile contains several tabs that categorize the information.

The All tab displays a summary of available storage subsystem profile data. Select a tab to view section-specific information.

Numbers of standard logical drives, mirror repository logical drives, FlashCopy repository logical drives, FlashCopy logical drives, mirror pairs, and logical drive copies are each shown separately in a series of sub-tabs contained in this tab.

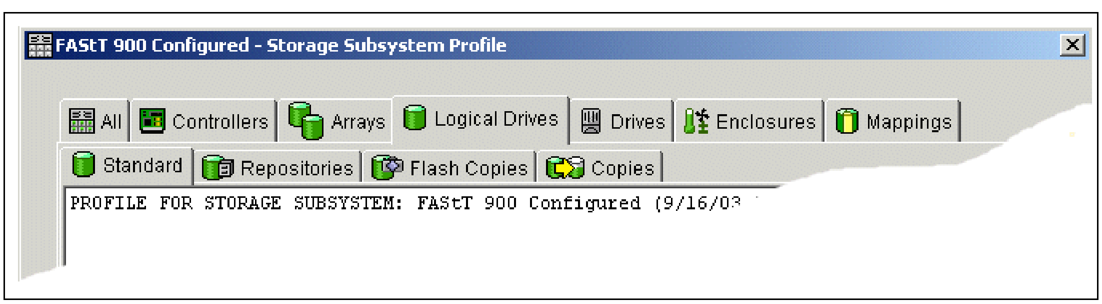

Figure 9-27 Storage subsystem profile - All tab

4. The Logical Drives tab is divided into five sections (depending on the number of Premium features enabled), with separate tabs providing access to information for each type of logical drive. A Missing Logical Drives tab is only available when missing logical drives are detected.

   Select a the Copies tab to view the VolumeCopy associated information. See Figure 9-28 on page 306. The Copies tab contains a summary section of information on all logical copy drive relationships on the Storage Subsystem, and a detailed section of each

VolumeCopy pair. The Copies tab is only available when the VolumeCopy premium feature is enabled.

.

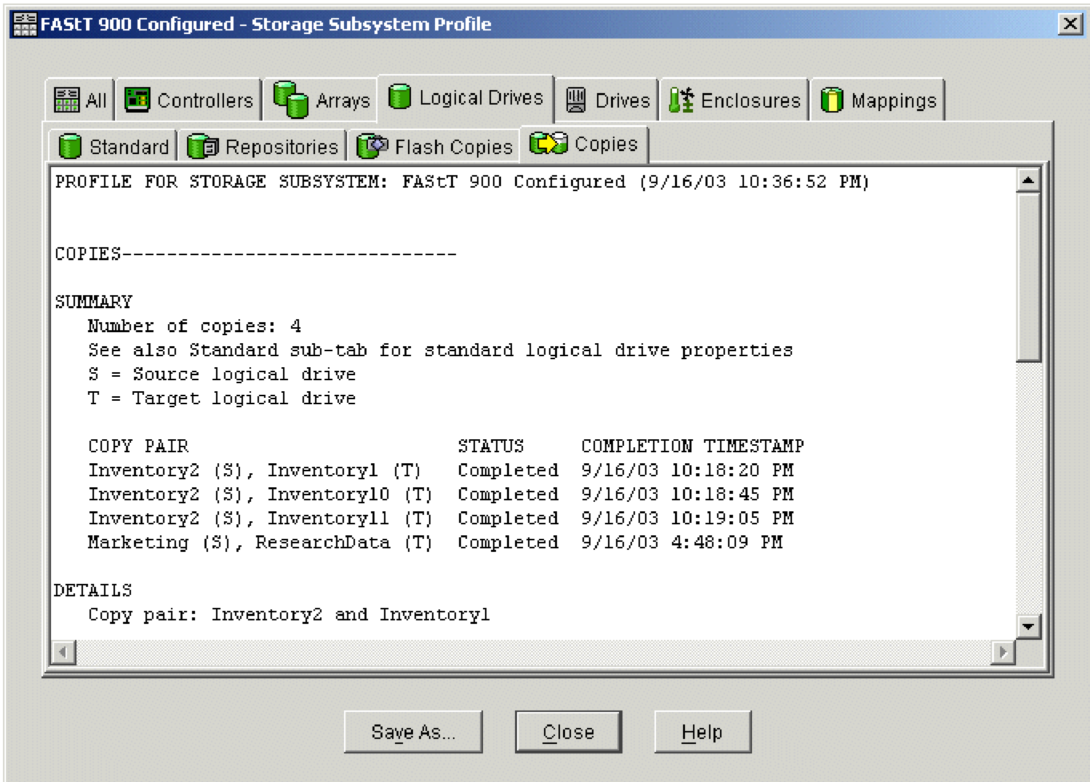

Figure 9-28 Storage Subsystem Profile - Logical Drives - Copies tab

9.3 Command Line Interface for VolumeCopy

The Command Line Interface was introduced in “Script Editor and command line interface” on page 89. For more in-depth information about the Command Line Interface and the Script Editor, refer to Chapter 10, “Command Line Interface and Script Editor” on page 311.

This section discusses commands that are specific to the VolumeCopy function. They are:

- **`start logical driveCopy`**
- **`show logical driveCopy`**
- **`set logical driveCopy`**
- **`stop logical driveCopy`**
- **`remove copyPair`**

9.3.1 start logical driveCopy - CLI

The start logical driveCopy command lets you create a VolumeCopy. You specify the source logical drive, the target logical drive, and the copy priority.

In Example 9-1 on page 307 we create a VolumeCopy, specifying logical drive “Data 2” as the source logical drive and “Data 1” as the target logical drive, with a copy priority of low. The following command will be used.

Example 9-1 start logical driveCopy

```
SMcli 10.1.1.1 -c "start volumeCopy source=\"Data 2\" target=\"Data 1\" copyPriority=low;"
```

9.3.2 show logical driveCopy - CLI

The show **logical driveCopy** (**volumeCopy**) shows data about all VolumeCopy pairs on the storage subsystem or data about a specific VolumeCopy pair if the source logical drive is specified.

In Example 9-2 the allVolumes parameter is specified and data is displayed about all VolumeCopy pairs on the subsystem.

Example 9-2 show logical driveCopy

```
SMcli 10.1.1.1 -c "show volumeCopy allVolumes;"
Performing syntax check...
Syntax check complete.
Executing script...
Copy pair: Data 2 and Data 1
      Copy status: In progress
      Completion timestamp: None
      Copy priority: Low
      Source volume: Data 2
      Target volume: Data 1
Script execution complete.
SMcli completed successfully.
```

9.3.3 stop logical driveCopy - CLI

You can stop a VolumeCopy that is in progress, pending, or failed. This was discussed in “Stopping VolumeCopy” on page 301.

To stop a copy pair you must specify the target and source logical drives names.

Example 9-3 shows how to stop the copy pair that has a source of “Data 2” and a target of “Data 1”:

Example 9-3 stop logical driveCopy

```
SMcli 10.1.1.1 -c "stop volumeCopy target[\"Data 1\"] source[\"Data 2\"];"
```

The copy is stopped as shown in Figure 9-29 on page 308.

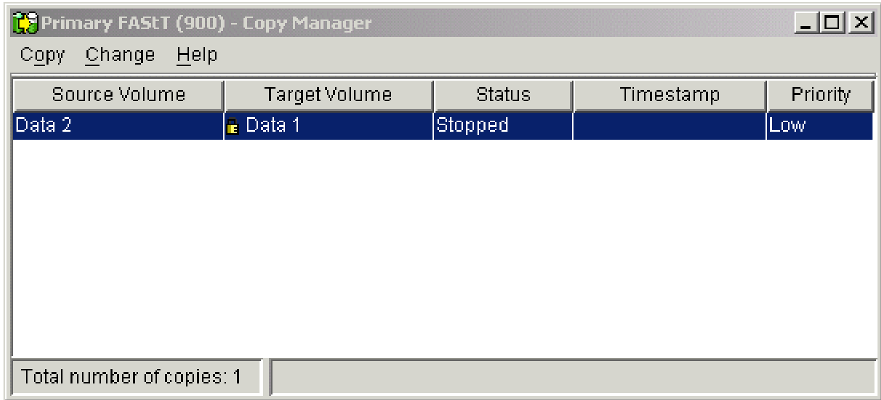

Figure 9-29 stop volumeCopy - CLI

9.3.4 set logical driveCopy - CLI

The **set logical driveCopy** command allows you to change the copy priority of a copy pair as discussed in "Change VolumeCopy priority" on page 299 and to set the read-only attribute of a copy pair as was discussed in "Set the Read-Only attribute for a target logical drive" on page 300.

The name of the target and source logical drives of the copy pair need to be specified.

In Example 9-4, we choose the copy pair consisting of target "Data 1" and source "Data 2". We change the original copy priority from low to highest and set the target read only attribute to disabled.

Example 9-4 set logical driveCopy

```
SMcli 10.1.1.1 -c "set volumeCopy target[\"Data 1\"] source[\"Data 2\"]
copyPriority=highest targetReadOnlyEnabled=false;"
```

The result can be seen in Figure 9-30.

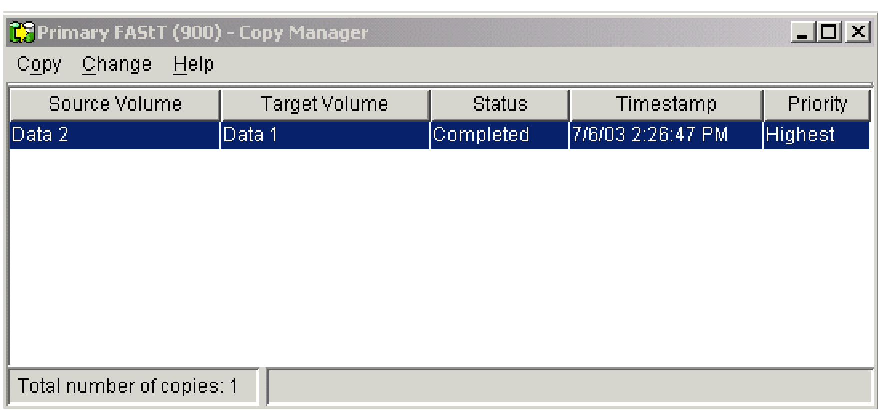

Figure 9-30 set Logical driveCopy - CLI

9.3.5 remove copyPair - CLI

Removing a copy pair relationship was discussed in "Remove copy pairs" on page 303.

In Example 9-5, we remove the copy pair with the target of “Data 1” and the source of “Data 2”:

Example 9-5 remove copyPair

```
SMcli 10.1.1.1 -c "remove volumeCopy target[\"Data 1\"] source[\"Data 2\"];"
```

As a result, the copy pair consisting of Data 1 and Data 2 is removed and these two logical drives no longer have any association with each other.

10

Command Line Interface and Script Editor

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.

This chapter explains how to work with the Command Line Interface and the Script Editor

10.1 Command Line Interface (CLI)

The CLI provides an efficient way to edit, send, and execute storage management commands on multiple network storage subsystems.

The CLI helps in automating many tasks such as backup procedures. For example, you can write a script using FlashCopy or VolumeCopy, integrate it with your backup software, and thus automate your backup.

The CLI is part of the SMclient package: No additional software is required. For installation instructions for SMclient on specific operating systems refer to “Installing the HBA device drivers and host software” on page 110.

Commands are invoked through the SMcli program. The commands must be launched from the installation directory (the default folder on Windows is c:\Program Files\IBM_FAStT\client) or include the path to the SMcli program in the system environment settings.

Example 10-1 shows the help information displayed as a result of executing **SMcli -?**.

Example 10-1 Help on SMclient

```
C:\Program Files\IBM_FAStT\client>SMcli -?

SMcli <hostname or IP address>
        [<hostname or IP address>]
        [-c "<command>;[<command2>;...]"]
        [-n <SAName> | -w <WWName>]
        [-o <outputfile>] [-p <password>] [-e]
SMcli <hostname or IP address>
        [<hostname or IP address>]
        [-f <scriptfile>]
        [-n <SAName> | -w <WWName>]
        [-o <outputfile>] [-p <password>] [-e]
SMcli {-n <SAName> | -w <WWName>}
        [-c "<command>;[<command2>;...]"]
        [-o <outputfile>] [-p <password>] [-e]
SMcli {-n <SAName> | -w <WWName>}
        [-f <scriptfile>]
        [-o <outputfile>] [-p <password>] [-e]
SMcli -d [-i] [-s] [-w]
SMcli -A [<ip address1> [ip address2]]
SMcli -m <ip address> -F <email address>
SMcli {-a | -x} email:<email address>
        [{<ip address1> [ip address2] [-n <SAName>
        | -h <hostName>] | -n <SAName>
        | -h <hostName>}]
SMcli {-a | -x} trap:<community>,<ip address>
        [{<ip address1> [ip address2] [-n <SAName>
        | -w <WWName> | -h <hostName>]
        | -w <WWName> | -h <hostName>}]
SMcli -?

For additional information, please see the "Command Reference" topic
in the Enterprise Management Window online help or use man SMcli
or help SMcli depending on your operating system.

SMcli completed successfully.

C:\Program Files\IBM_FAStT\client>
```

Note: Getting help on a CLI interface can vary depending on the operating systems. Enter `SMcli` without any parameters to display a short help message.

10.1.1 Using CLI commands

Sending the commands to your system depends on your storage subsystem management method:

- Directly managed (out-of-band): Use the *host name* or *IP address* the controller(s).

  ```
  SMcli <IP address either of the controllers>
  ```

  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. The name must be unique to the Enterprise Management window, for example:

  ```
  SMcli -n <name of the FAStT>
  ```

  If the name of the FAStT consists of spaces and special characters you have to put it into double quotes ("), for example:

  ```
  SMcli -n "Remote FAStT (700)"
  ```

  Note: When executing CLI using an IP address, you can use the IP address of either controller.

- Host-agent managed (in-band): Use the *host name* or *IP address* of the host.

  ```
  SMcli <hostname of managing station>
  ```

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

  ```
  SMcli <hostname of managing station> -n <name of the FAStT>
  ```

  If you specify the world-wide name of the storage subsystem, use the -w option instead of the -n option, for example:

  ```
  SMcli -w 600a0b800006602d000000003beb684b
  ```

After executing one of the above commands, you are in interactive mode and can enter one or more commands. 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 Example 10-2.

Example 10-2 SMcli command in interactive mode

```
C:\Program Files\IBM_FAStT\client> SMcli 10.1.1.3
Executing script...

show storageSubsystem batteryAge; (first command)
Battery status: Optimal
   Age: 480 day(s)
   Days until replacement: 689 day(s)
show storageSubsystem healthStatus; (second command)

The controller clocks in the storage subsystem are out of synchronization with the storage
management station.

Controller in Slot A: Mon Jul 07 17:02:46 MDT 2003
Controller in Slot B: Mon Jul 07 17:02:49 MDT 2003
Storage Management Station: Mon Jul 07 16:06:27 MDT 2003
```

```
Storage Subsystem health status = optimal.
Script execution complete.

SMcli completed successfully. (invoking CTRL-D)

C:\Program Files\IBM_FAStT\client>
#
```

Invoking SMcli specifying a storage subsystem, but not specifying the commands or script file to execute will cause SMcli to run in interactive mode, allowing you to interactively specify the commands. Use Ctrl + D to stop SMcli execution.

You do not have to run the CLI in interactive mode. You can pass parameters to CLI as you run it by specifying the -c option with the SMcli program.

```
SMcli <IP address either of the controllers> -c "<command>;[<command2>;...]"
```

This executes the command that is found after the -c parameter. See the following example for the output of the command.

Example 10-3 Non-interactive commands

```
C:\Program Files\IBM_FAStT\client> SMcli 10.1.1.3 -c “show storageSubsystem batteryAge;
show storageSubsystem healthStatus;"
Performing syntax check...

Syntax check complete.

Executing script...

Battery status: Optimal
   Age: 480 day(s)
   Days until replacement: 689 day(s)

The controller clocks in the storage subsystem are out of synchronization with t
he storage management station.

Controller in Slot A: Mon Jul 07 17:12:40 MDT 2003
Controller in Slot B: Mon Jul 07 17:12:42 MDT 2003
Storage Management Station: Mon Jul 07 16:16:21 MDT 2003

Storage Subsystem health status = optimal.
Script execution complete.

SMcli completed successfully.

C:\Program Files\IBM_FAStT\client>
```

You can also call complete pre-written scripts. You do this by specifying the -f option with the SMcli program:

```
SMcli <IP address either of the controllers> -f <script file name>
```

This executes the commands contained in the script file. If the script file is not in SMclient directory or you do not have the path to the SMclient in your system environment settings, you have to specify the full path to the script. See Example 10-4 on page 315, where the command (**`show StorageSubsystem batteryage`**) is imbedded in the prewritten script (script.scr).

Example 10-4 Sending scripts to SMclient

```
C:\Program Files\IBM_FAStT\client>SMcli 10.1.1.3 -f script.scr

Performing syntax check...

Syntax check complete.

Executing script...

Battery status: Optimal
   Age: 477 day(s)
   Days until replacement: 692 day(s)
Script execution complete.

SMcli completed successfully.

C:\Program Files\IBM_FAStT\client>
```

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

For detailed information on the CLI parameters, consult the Command Line reference included in the SMclient online help.

10.1.2 Syntax requirements

SMcli has the following usage and formatting requirements:

- Usage requirements that apply to all operating systems:
  - All statements must end with a semi-colon (;).
  - Separate each base command and any parameters with a space.
  - Separate each parameter and its parameter value with an equal sign.
  - The Script Editor and Command Line Interface are not case-sensitive. You can enter any combination of upper- and lowercase letters. Following the convention of having a capital letter start the second word of a parameter could make scripts easier to read.
  - Invoking SMcli with no arguments or with an unrecognized parameter will cause usage information to be displayed.
  - Arguments used after the -n, -o, -f, and -p options that contain a space, a number, or a special character (<, >, ', !, *, for example) need to be enclosed in single quotes (') or double quotes ("), depending on your operating system.
  - Arguments used after the -n, -o, -f, and -p options that contain a single quote character (') need to be enclosed in double quotes (").
  - Invoking SMcli specifying a storage subsystem, but not specifying the commands or script file to execute, will cause SMcli to run in interactive mode, allowing you to interactively specify the commands. Use Ctrl + C to stop SMcli execution.
- Usage requirements that apply to Windows operating systems only:
  - Insert a backslash (\) before each double quote character (") when the double quotes are used as part of a name or command syntax (for example, -c "set storageSubsystem userLabel=\"string\";").

- Insert a backslash (\) before each quote around a user label that contains spaces (for example, -c "start logical driveCopy source=\"Mirror Repository 1\" target=trg9 priority=high;").
- Insert three backslashes (\\\) in front of the (") to display the backslash when used with the -n, -o, -f, or -p option (for example, -n "Jason\\\" to specify storage subsystem named Jason\).
- Insert five backslashes (\\\\\) in front of the (") to use the backslash character as part of the literal command string (for example, -c "set storageSubsystem userLabel=\"Jason\\\\\";" to change the name of the storage subsystem to Jason\\).
- Insert a caret (^) before each special script character (^, &, |, <, >) when that character is used with the -n, -o, -f, and -p options (for example, -n "CLI^&CLIENT" to specify storage subsystem "CLI&CLIENT". (See the appropriate operating system scripting documentation for a list of special script characters.)
- Insert three carets (^^^) before each special script character when used within a literal script command string (for example, -c "set storageSubsystem userLabel=\"Finance^^^&payroll\";" to change the name of the storage subsystem to Finance&Payroll).

- Usage requirements that apply to UNIX operating systems only:

  The entire command string must be enclosed in single quotes ('), although some simple commands may also work with double quotes (“).

10.1.3 Command reference

This section contains commands that are current with the latest storage management software (Version 8.4). If you are managing storage subsystems running firmware for previous releases, not all of these commands may be supported.

> **Tip:** While the Script Editor and Command Line Interface syntax have undergone some revisions, the former syntax is still supported. Any scripts that adhere to previous syntax rules will still pass the Script Editor syntax check and will execute.

create

The `create` command can do the following:

- 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

The `delete` command can do the following:

- 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

The **disableFlashCopy** command can disable a FlashCopy logical drive.

download

The **download** command can do the following:

- 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

The **on error** command can 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

The **recreateFlashCopy** command can do the following:

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

remove

The **remove** command can do the following:

- Remove logical drive copies.
- Remove logical drive reservations.

reset

The **reset** command can do the following:

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

set

The **set** command can do the following:

- 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

The **show** command can do the following:

- Display the properties of the different logical and physical components comprising the storage subsystem.

- Embed text strings (comments) in your command line that will display in the Output View during command execution.
- Show storage partition topology and mappings information.

start

The `start` command can do the following:

- Start a logical drive expansion.
- Start a VolumeCopy.

stop

The `stop` command can stop a VolumeCopy.

upload

The `upload` command can do the following:

- 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

The `use` command can 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.

10.1.4 CLI parameters

The command line interface supports the command line parameters shown in Table 10-1.

Table 10-1 CLI parameters

Command line parameter	Lets you...
<IP address> or <hostname>	Specify an IP address (xx.xx.xx.xx) or host name (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	Specifies 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.

10.1.5 CLI examples

Following are examples of how you can use the CLI to access and execute script engine commands. Note that the usage for the -c parameter varies depending on your operating system (enclosed in single quotation marks (') in UNIX or double quotation marks (") in Windows.

1. Rename "Remote FAStT (700)" to "Test FAStT 700" using the host IP10.1.1.3.

Example 10-5 Windows system

```
SMcli 10.1.1.3 -c “set storageSubsystem userlabel=\"Test FAStT 700\";"
Performing syntax check...
Syntax check complete.
Executing script...
Script execution complete.
SMcli completed successfully.
```

Example 10-6 UNIX system

```
SMcli 10.1.1.3 -c ‘set storageSubsystem userlabel="Test FAStT 700";’
Performing syntax check...

Syntax check complete.

Executing script...

Script execution complete.

SMcli completed successfully.
```

2. On the storage subsystem with controller names "Test FAStT 700" use the password 700 to do the following:
   - Delete the logical drive named "test".
   - Create a new logical drive named "Development".
   - Show the health status of the storage subsystem, which is managed via the direct management method.

Example 10-7 Windows system

```
SMcli -n "Test FAStT 700" -c "use password \"700\"; delete logicalDrive[\"test\"];create
logicalDrive array=1 capacity=10GB UserLabel=\"Development\";show storageSubsystem
healthStatus;"

Performing syntax check...

Syntax check complete.

Executing script...

Storage Subsystem health status = optimal.
Script execution complete.

SMcli completed successfully.
```

Example 10-8 UNIX system

```
SMcli -n "Test FAStT 700" -c 'use password "700"; delete logicalDrive[test];create
logicalDrive array=1 capacity=10GB UserLabel="Development";show storageSubsystem
healthStatus;'

Performing syntax check...
```

```
Syntax check complete.

Executing script...

Storage Subsystem health status = optimal.
Script execution complete.

SMcli completed successfully.
```

10.2 Script editor

The script editor is a powerful tool to create and edit the scripts. While creating the script you can verify its syntax or even execute it to see if it meets your requirements. Scripts obviously can be saved and later executed to automate storage management procedures.

10.2.1 Using the script editor

To open the script editor, follow these steps:

1. Select a storage subsystem in the Device Tree View or Device Table from the Enterprise Management Window in the Storage Manager Client.
2. Click **Tools -> Execute Script** as shown in Figure 10-1, or right-click and select **Execute Script**.

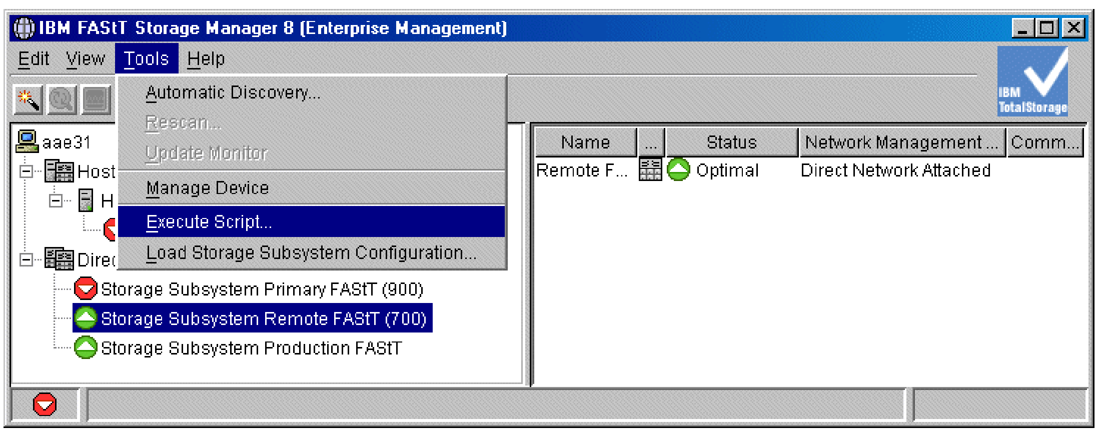

Figure 10-1 Starting the Script Editor

The Script Editor opens as shown in Figure 10-2 on page 323. 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.

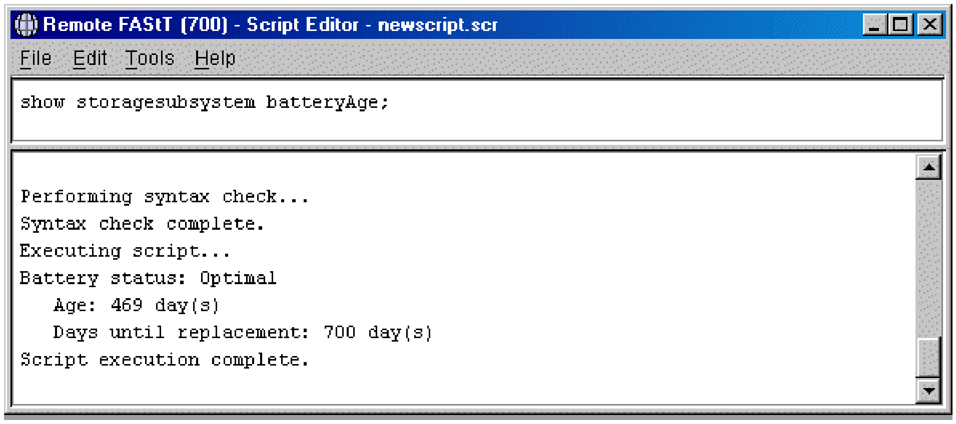

Figure 10-2 The Script Editor

Usage guidelines

Follow 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 command

Use the **show** command with string (enclosed in a double quotes) and no options to embed in your script comments that display in the output view during script execution. For example, including a Show **"START of the script"** command in your script results in the display of START of the script in the output view when this line is processed during script execution. Including show statements can help you when writing longer scripts where intermediate steps can make troubleshooting easier. See Figure 10-3.

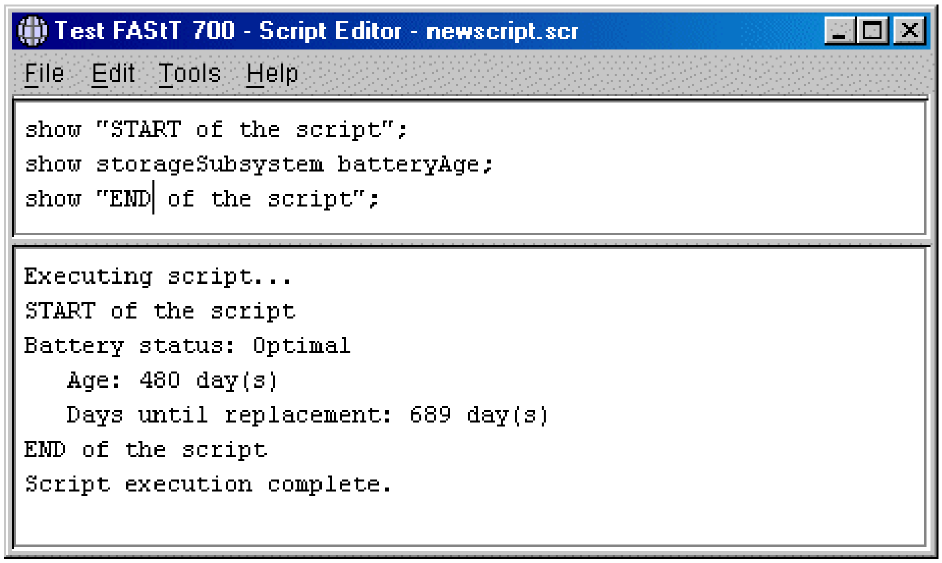

Figure 10-3 show command in script editor

Script editor tools

The script editor offers the following tools to help you when writing scripts:

- Verify Syntax

  To run this option select the **Tools -> Verify Syntax** pull-down menu option. The script editor engine parses the statements in the script file one line at a time and verifies that they have the correct syntax. Any syntax errors are displayed in the output view, reporting the line number of the error and a description of the error. If the script editor encounters a syntax error, no further syntax verification is performed on the script. Fix the syntax error and rerun the **Verify Syntax** command to validate the error correction and check the remainder of the statements in the script.

- Verify and Execute

  To run this option select the **Tools -> Verify and Execute** option. The script editor engine parses the command statements in the script, interprets and converts the statements to the appropriate commands, and sends the commands to the storage subsystem.

  If a *syntax error* is encountered, the execution stops and an error message is displayed. Fix the error, then use the Verify Syntax or Verify and Execute options to validate the error correction.

  If an *execution error* occurs, the script may or may not continue to execute depending on the included On Error script statement. The On Error Stop statement stops the script if an execution error is encountered. The On Error Continue statement allows the script to continue even after an execution error is encountered. (This is the default.)

- Execute Only

  To run this option select the **Tools -> Execute Only** option. The Script Editor engine executes a script. It displays an error message if a syntax error is encountered.

  If an execution error occurs, the script may or may not continue to execute depending on the included On Error script statement. The On Error Stop statement stops the script if an execution error is encountered. The On Error Continue statement allows the script to continue even after an execution error is encountered. (This is the default.)

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

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 indication 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.

10.2.2 Embedding commands in batch files

Due to the business demand for higher availability of data, the time window allocated to take backups is shrinking. Customers can no longer afford the daily downtime on their production server to perform backups. This becomes especially true as databases become larger and larger. Online backups improve the availability of the database, but cost CPU, disk, and network resources. Even with the use of incremental and differential backups, the time necessary to perform a backup can be significant.

On the other hand, a feature like FlashCopy enables customers to perform offline backups. As an example we take the reader through the creation of FlashCopy on a Windows system, and show how command line scripts can be embedded in a batch file; the batch file can in turn be integrated with the backup software for automation purposes.

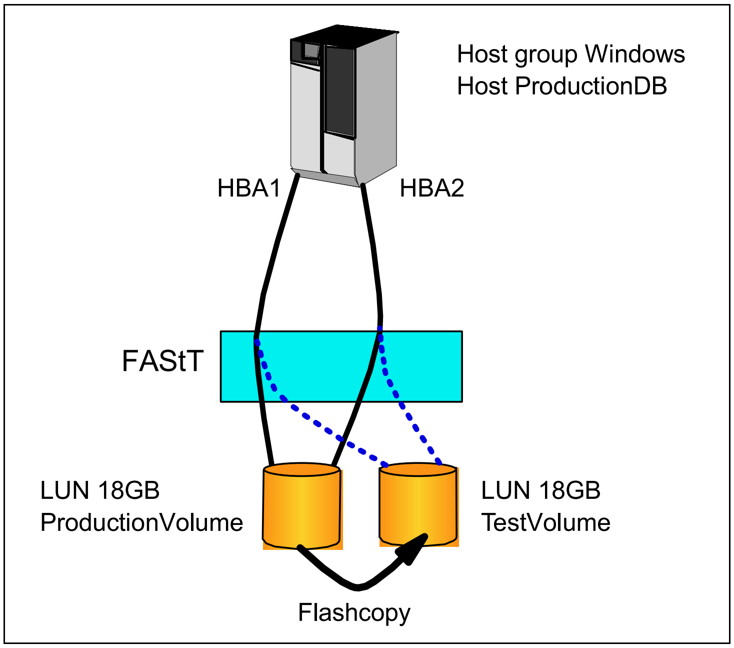

Figure 10-4 FlashCopy

The needed steps are:

1. Flushing all file systems for the logical drive that will have FlashCopy performed.
2. Creating a FlashCopy logical drive on the storage subsystem.
3. Mapping the new FlashCopy logical drive to the host with a unique LUN.
4. Using the hot_add utility to cause a bus rescan. New devices will be found.

State of the system before FlashCopy

Before FlashCopy the state of the system and volumes were as follows:

1. Volume with name ProductionVolume present in Array 4. The capacity of ProductionVolume is 18GB. See Figure 10-5 on page 327.
2. Host group Windows has host ProductionDB defined with two host ports. Mapping exists for host group Windows only for volume ProductionVolume. See Figure 10-6 on page 327.
3. Because the FAStT is connected to an Intel server with two host bus adapters; Windows computer management utility shows (Device Managemer) two 1742 SCSI Disk devices and one RDAC Virtual disk. See Figure 10-7 on page 327.
4. The Windows system has two drives defined. One is the C drive, which is System disk; and the other is the 18GB disk on the FAStT Storage Server. See Figure 10-8 on page 328.

We want to perform a FlashCopy of the ProductionVolume and mount it back to the original system (we have to use Windows basics disks for this). The name of the new volume will be TestVolume and we want to use it for testing purposes. A FlashCopy script will enable testing and backup of the production data from TestVolume with almost no downtime.

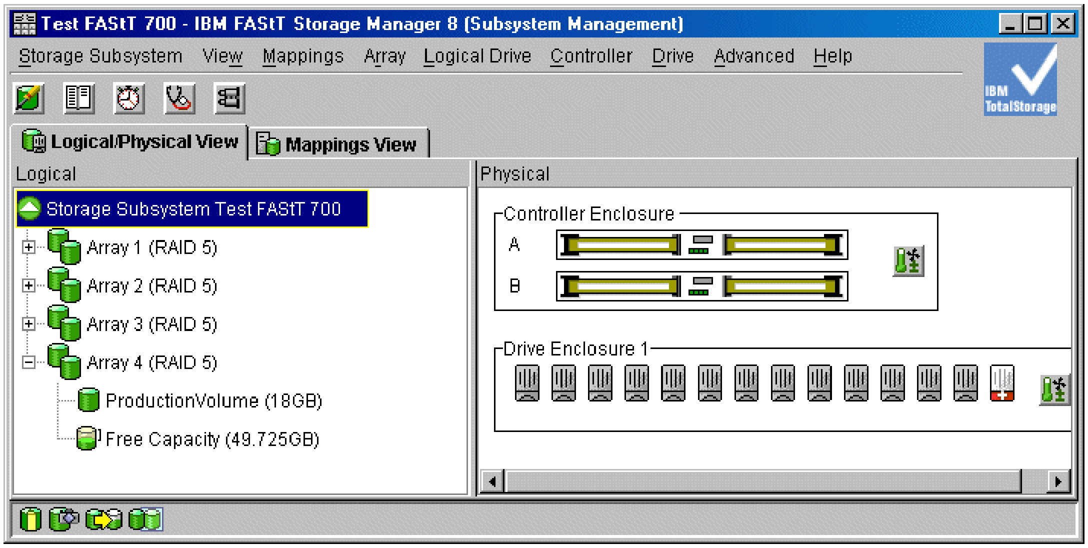

Figure 10-5 State of the volumes before FlashCopy

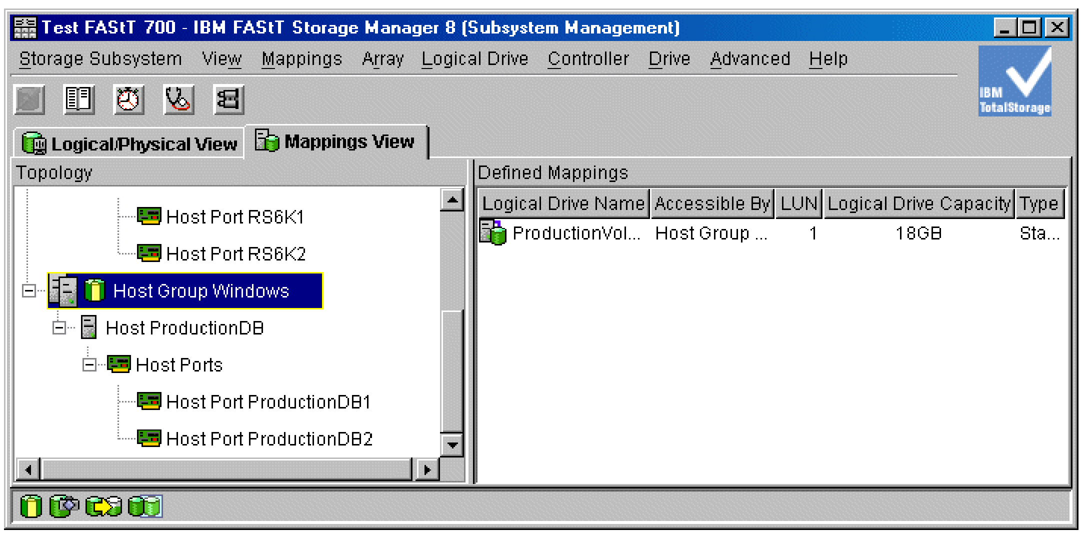

Figure 10-6 State of host mapping before FlashCopy

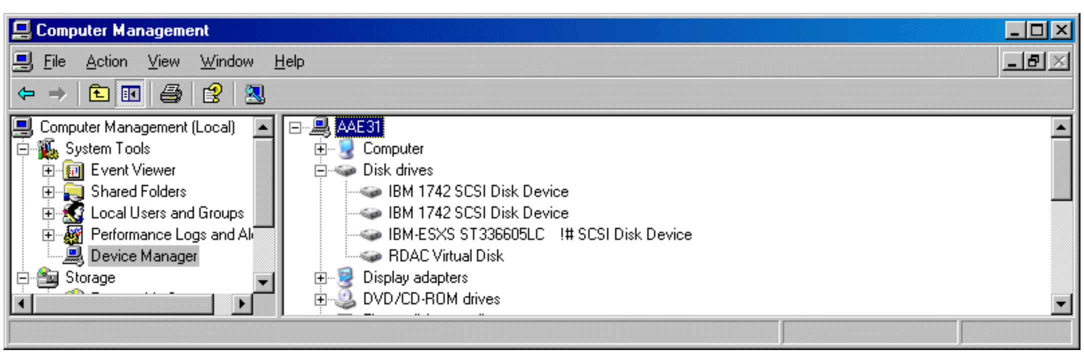

Figure 10-7 Windows device status before FlashCopy

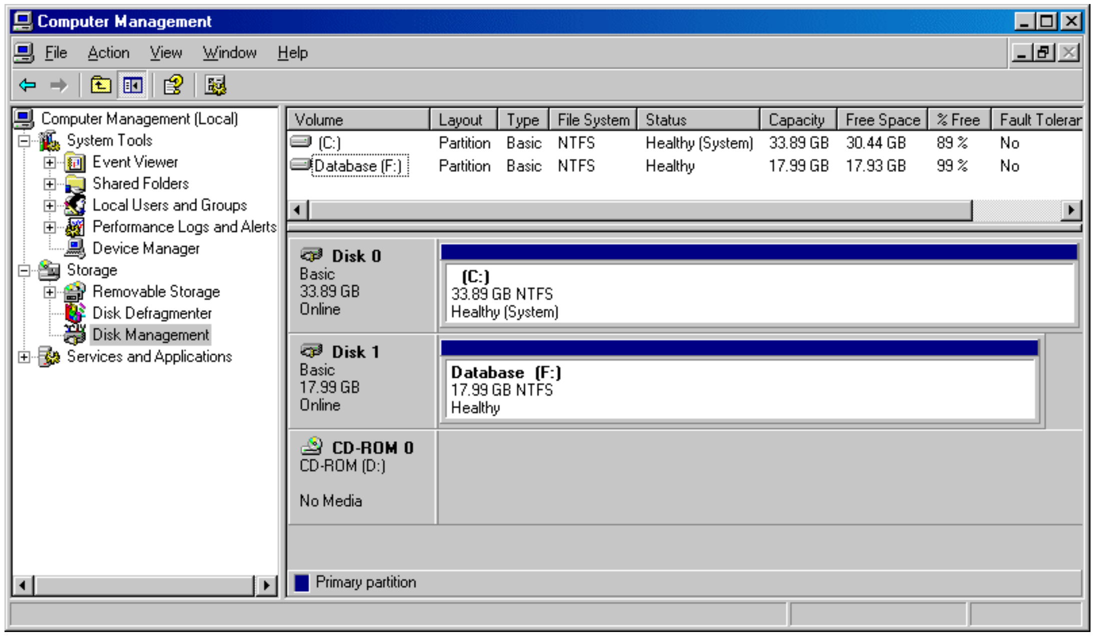

Figure 10-8 Windows drive status before FlashCopy

FlashCopy batch file

Example 10-9 shows the batch file used to perform a FlashCopy in a Windows system.

- The first step is to flush the buffers (forces all the data in the cache to disk storage).
- Next we perform a FlashCopy—on the same array—with new names for the target drive and repository drive.
- After completion of the FlashCopy we define the LUN mapping to the system (for Windows to see the new drive).
- Finally, the hot_add utility helps us rescan the bus of the system and recognize the new drive without restarting the system.

Note also that the batch file includes the complete paths to the executable files. This guarantees that the batch file can be run from any directory.

Example 10-9 FlashCopy script

```
@echo off
rem
rem Example of automating the creation of a flashcopy logical drive under Windows
rem The steps that will be performed are:
rem 1) Flushing file systems for the logical drive that will have flashcopy performed.
rem 2) Creating a flashcopy logical drive on the storage array.
rem 3) Mapping the new flashcopy logical drive to the host with a unique LUN.
rem 4) Using the hot_add utility to cause a bus rescan.  New devices will be found.

rem Set the path to the hot_add utility (in the SMutil package).
set HOTADD_PATH="c:\Program Files\IBM_FAStT\Util\hot_add"

rem Set the path to the SMflashcopyassist utility.
set FLASHASSIST_PATH="c:\Program Files\IBM_FAStT\Util\SMflashcopyassist"

rem A drive letter residing on a logical drive that will have flashcopy performed.
set DRIVE_TO_SNAPSHOT="f:"
```

```
rem Set the path to the SMClient command line script executable.
set SM_CLI_PATH="C:\Program Files\IBM_FAStT\client\SMcli"

rem Address of the FAStT
set FAStT_ADDRESS="10.1.1.3"

rem Constants for the create flashcopy logical drive script.
rem Name of the base logical drive to take a flashcopy of.
set BASE_VOLUME_NAME=\"ProductionVolume\"
rem Label to give the new flashcopy logical drive.
set USER_LABEL=\"TestVolume\"
rem Label to give the new repository logical drive.
set REPOSITORY_USER_LABEL=\"TestRepository\"

rem The example script uses an existing RAID array to create the repository logical drive.
set CREATE_SNAPSHOT_SCRIPT="on error stop; create FlashCopyLogicalDrive
baseLogicalDrive=%BASE_VOLUME_NAME% userLabel=%USER_LABEL%
repositoryUserLabel=%REPOSITORY_USER_LABEL%;"

rem This script performs the mapping of the flashcopy logical drive to the host computer.
rem Logical unit number for the logical drive/host mapping.
set LUN=13
rem Name of the host to map the logical drive to.
set MAP_TO_HOSTGROUP=\"Windows\"
set CREATE_MAPPING_SCRIPT="on error stop; create mapping logicaldrive=%USER_LABEL%
logicalUnitNumber=%LUN% hostgroup=%MAP_TO_HOSTGROUP%;"


rem Flush all the file systems buffers for the logical drive the drive is on
echo Flushing drive %DRIVE_TO_SNAPSHOT%.
%FLASHASSIST_PATH% -f %DRIVE_TO_SNAPSHOT%
if errorlevel 1 goto flush_fail_exit

rem Create a new flashcopy logical drive
echo Creating new flashcopy logical drive.
call %SM_CLI_PATH% %FAStT_ADDRESS% -c %CREATE_SNAPSHOT_SCRIPT%
if errorlevel 1 goto create_flashcopy_fail_exit

rem Create mapping for the new flashcopy logical drive
echo Creating host mapping for flashcopy logical drive.
call %SM_CLI_PATH% %FAStT_ADDRESS% -c %CREATE_MAPPING_SCRIPT%
if errorlevel 1 goto create_mapping_fail_exit

rem Use hot_add to cause the system to find all new devices
echo Running hot_add....
%HOTADD_PATH%

echo The flashcopy creation is complete
goto done

:flush_fail_exit
echo Flushing file systems for logical drive containing drive %DRIVE_TO_SNAPSHOT% failed
goto done

:create_flashcopy_fail_exit
echo Creating the flashcopy logical drive failed
goto done

:create_mapping_fail_exit
echo Creating a mapping for the flashcopy logical drive failed
```

```
goto done

:hotadd_fail_exit
echo Bus rescan using hot_add failed
goto done


:done


rem Unset all of the environment variables that were used.
set HOTADD_PATH=
set FLASHASSIST_PATH=
set DRIVE_TO_SNAPSHOT=
set SM_CLI_PATH=
set FAStT_ADDRESS=
set BASE_VOLUME_NAME=
set USER_LABEL=
set REPOSITORY_USER_LABEL=
set LUN=
set MAP_TO_HOST=
set CREATE_SNAPSHOT_SCRIPT=
set CREATE_MAPPING_SCRIPT=
```

Example 10-10 Output of the batch file

```
C:\temp>flashcopy
Flushing drive "f:".
Creating new flashcopy logical drive.
Performing syntax check...

Syntax check complete.

Executing script...
Script execution complete.

SMcli completed successfully.

Creating host mapping for flashcopy logical drive.
Performing syntax check...
Syntax check complete.

Executing script...
Script execution complete.

SMcli completed successfully.

Running hot_add....
C:\temp>
```

State of the system after FlashCopy

After execution of the FlashCopy script, the state of the system and volumes is as follows:

- Volume with name ProductionVolume present in Array 4. Two new volumes appear in the Array 4: FlashCopy volume TestVolume and repository volume TestRepository.

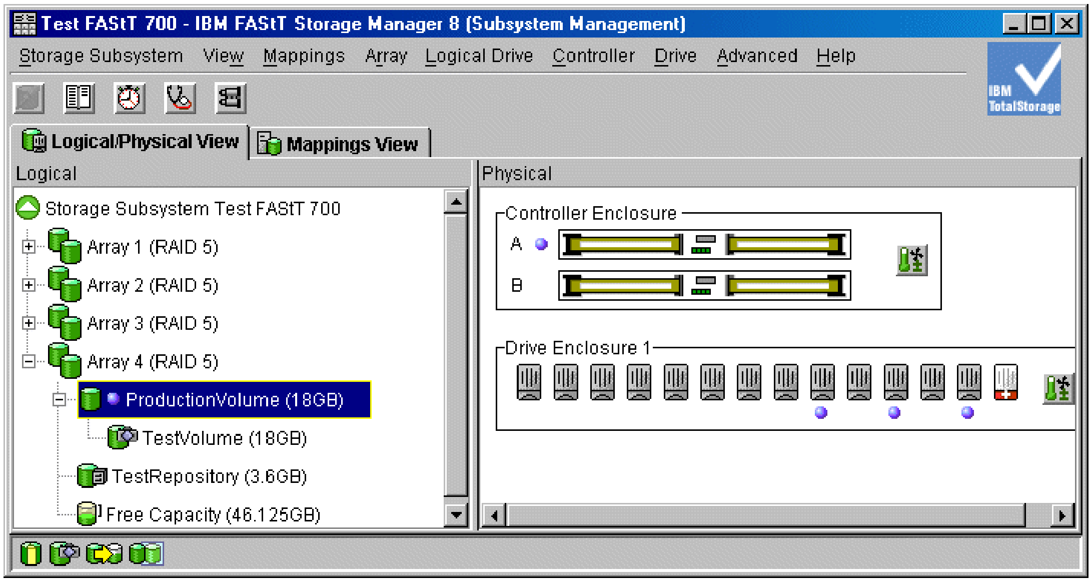

Figure 10-9 State of the volumes after FlashCopy

- Host group Windows has host ProductionDB defined with two host ports. Mapping exists for host group Windows for volume ProductionVolume. New mapping to Windows host group appears for the new volume TestVolume.

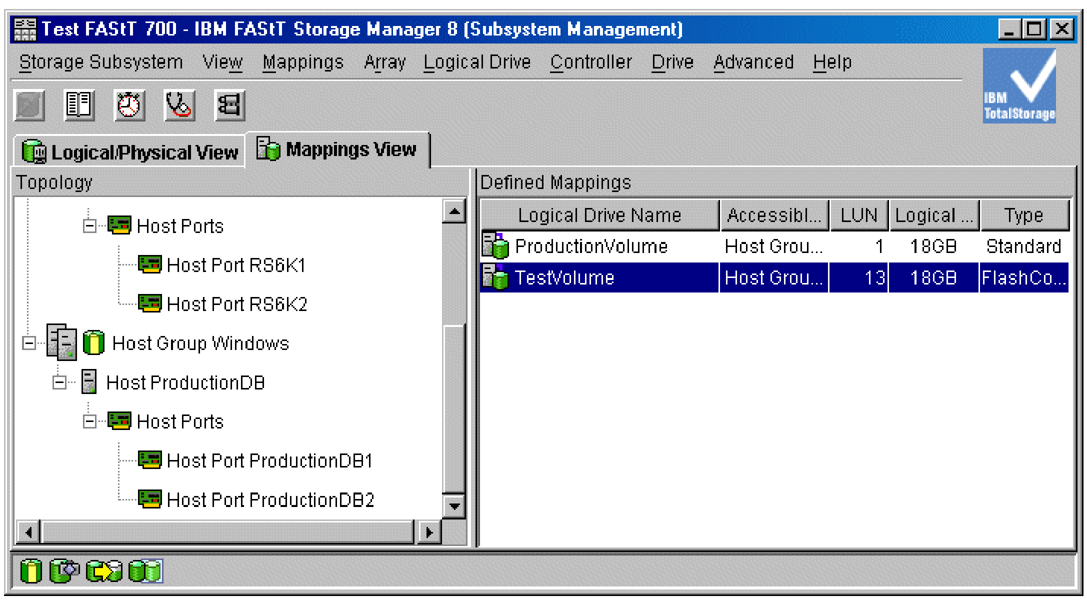

Figure 10-10 State of host mapping after FlashCopy

- The FAStT is connected to an Intel server with two host bus adapters and the Windows device status now shows now four 1742 SCSI Disk devices and two RDAC Virtual disks, because the FlashCopy volume is now accounted for in the system. See Figure 10-11 on page 332.

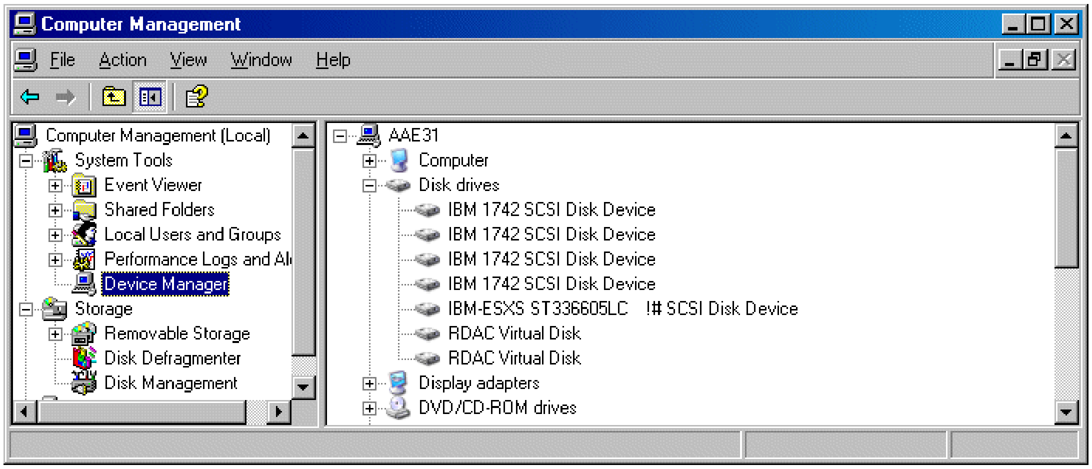

Figure 10-11 Windows device status after FlashCopy

- Windows "discovers" a new drive with the same name and capacity as the source drive (drives F: and E: in Figure 10-12).

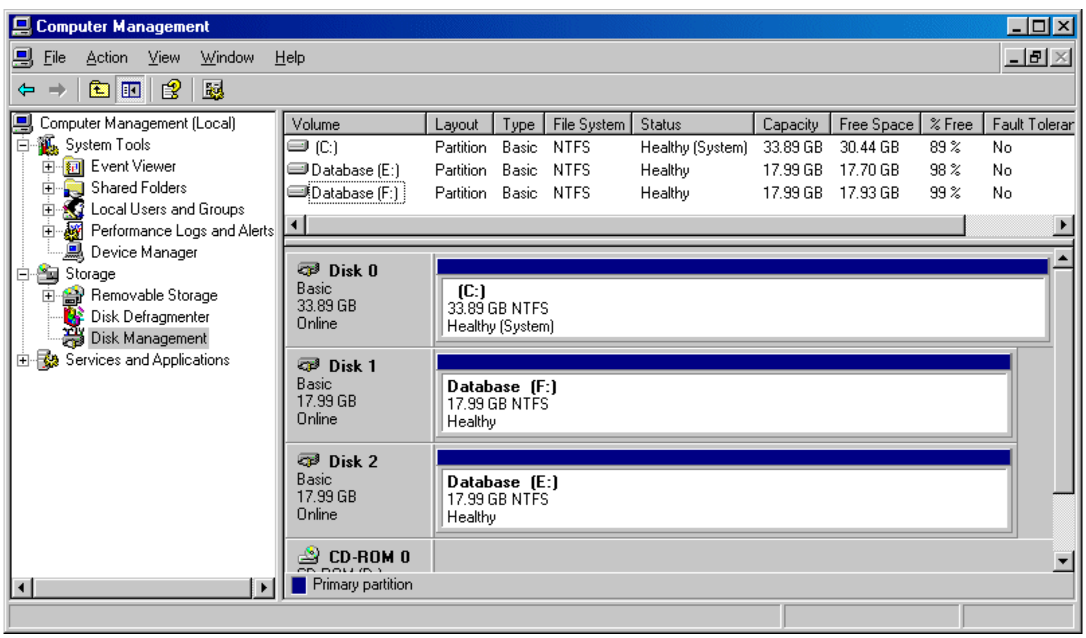

Figure 10-12 Windows drives after FlashCopy

.

11

Advanced maintenance

This chapter shows how to upgrade the firmware and associated components of the FAStT Storage Server; it also explains how to handle premium features.

11.1 Upgrades and maintenance

This section describes the steps that are required to upgrade your IBM FAStT Storage Server when firmware updates and/or fixes become available. We also describe how to enable and disable premium features, view the MEL, and save/import a configuration.

11.1.1 Prerequesites for upgrades

Upgrading the firmware and management software for the FAStT Storage Server is a relatively simple procedure. Before you start you should make sure that you have an adequate maintenance window to do the procedure, as on large configurations it can be a little time consuming. Time estimates for upgrading all the associated firmware and software are listed in Table 11-1. Figures are approximate and may vary from system to system.

Table 11-1 Upgrade Times

Element being upgraded	Approximate time of upgrade
Storage Manager Software and associated drivers and software	35 minutes
FAStT Storage Server Firmware	5 minutes
FAStT ESM Firmware	5 minutes
Hard drives	3 minutes per drive

It is critical that if you update one part of the firmware/software, you update all the firmware/software to the same level. You must *not* run a mismatched set.

All the nessessary files for doing this upgrade can be found at the following URL

http://www.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-52951

11.1.2 Updating FAStT host software

The host software must be updated first.

Updating for the Windows environment

To update the Window environment:

1. Uninstall the storage management components in the following order:
   - SMagent
   - SMutil
   - RDAC
   - SMclient
2. Verify that the IBM host adapter device driver versions are current. If they are not current, refer to the readme file located with the device driver and then upgrade the device drivers.
3. Install the SMclient. Refer to “Starting the FAStT Storage Manager client” on page 137.
4. Install the RDAC. Refer to “Install FAStT software on a Windows host server” on page 111.
5. Install the SMagent. Refer to “Install FAStT software on a Windows host server” on page 111.
6. Install the SMutil. Refer to “Install FAStT software on a Windows host server” on page 111.

Updating for the Linux environment

To update the Linux environment:

1. Uninstall the storage manager components in the following order:
   - FAStT Runtime enviroment
   - SMutil
   - SMclient
2. Verify that the IBM host adapter device driver versions are current. If they are not current refer to the readme file located with the device driver and then upgrade the device drivers.
3. Install the FAStT Runtime enviroment. Refer to “Installing SMruntime” on page 128.
4. Install SMutil. Refer to “Installing SMutil” on page 128.
5. Install SMclient. Refer to “Installing SMclient” on page 128.

11.1.3 Updating 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:** The following information is important:
   >
   > - Ensure that all hosts attached to FAStT have a multipath I/O driver installed.
   > - Update the controller firmware and then the NVRAM.
   > - 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 11-1 on page 336).

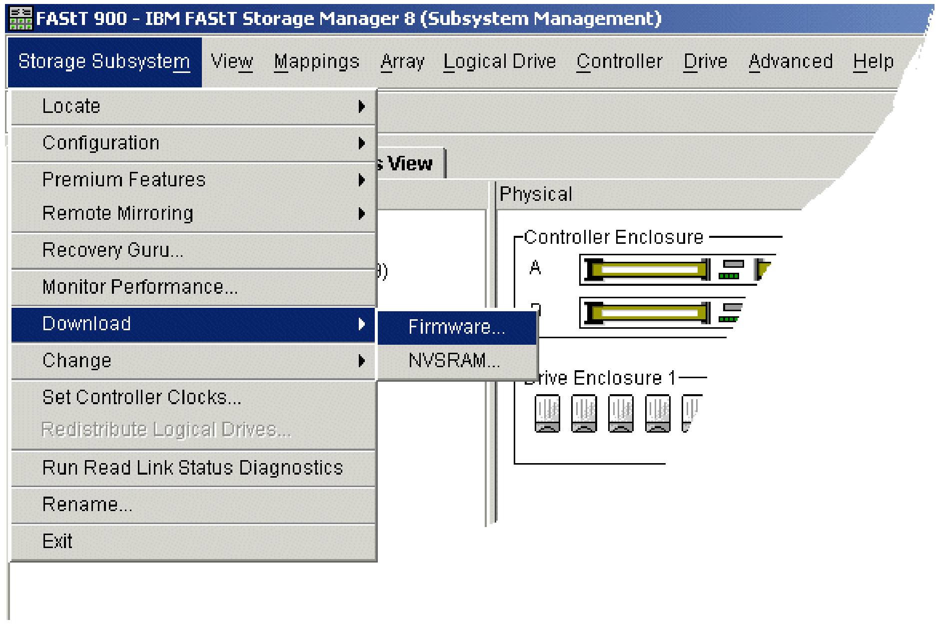

Figure 11-1 Subsystem Management window: Selecting Download -> Firmware

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 11-2.

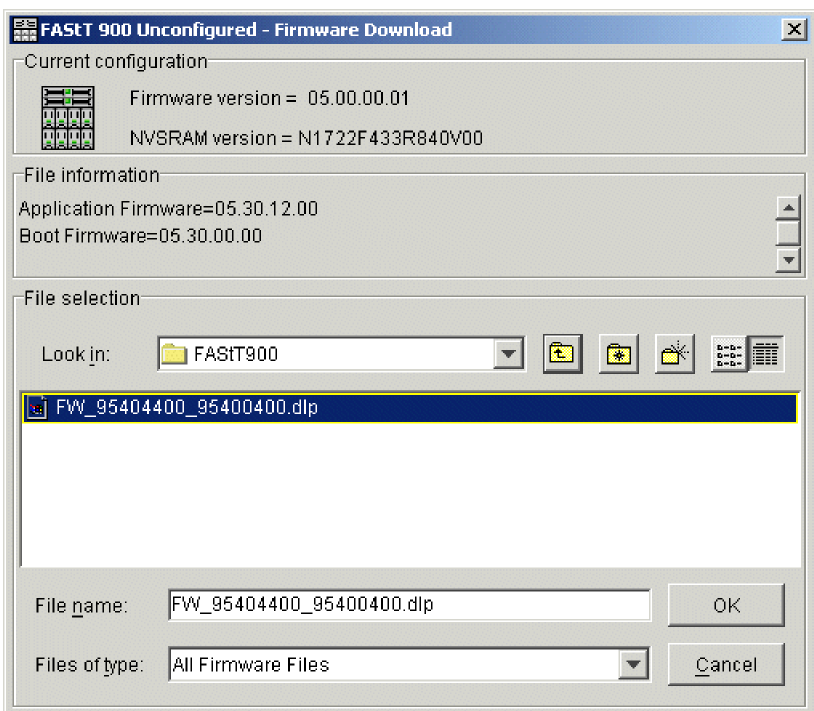

Figure 11-2 Firmware Download window

Click **OK**.

c. A new window (Figure 11-3) appears where you have to confirm that all the requirements are met. Click **Yes** to confirm the download.

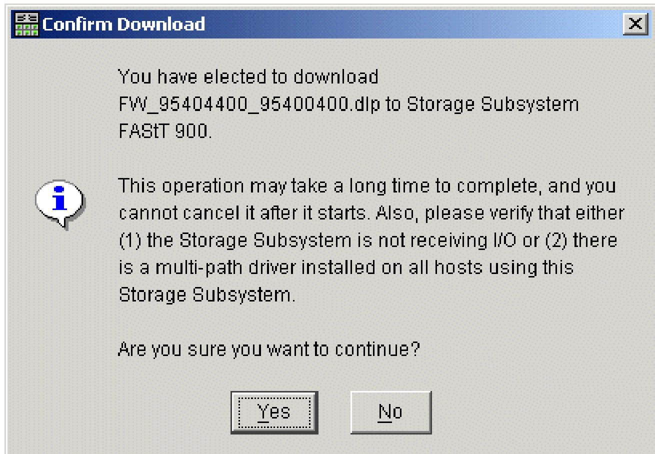

Figure 11-3 Confirming the download of the new firmware

3. The firmware upgrade continues. You can follow the process as shown in the Downloading window (Figure 11-4). The entire download process can take several minutes to finish.

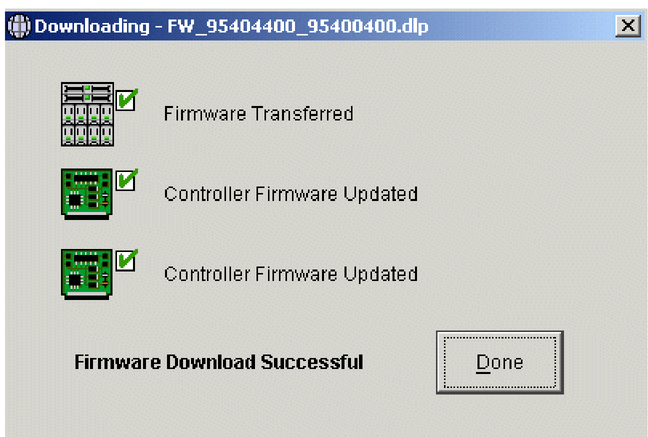

Figure 11-4 Download of firmware completed

You are then presented with the following window (Figure 11-5 on page 338), which tells you that you must launch the compatable version of the Management window.

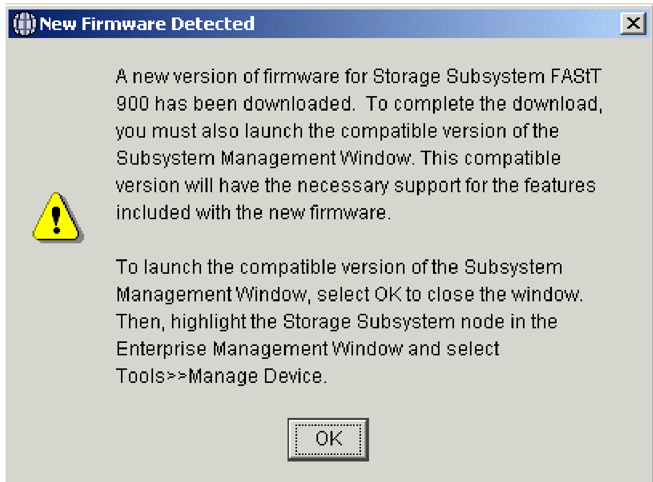

Figure 11-5 New version of Management window needed

You may also be asked to syncronise the clocks on the FAStT Storage Server with the host that you are using (see Figure 11-6).

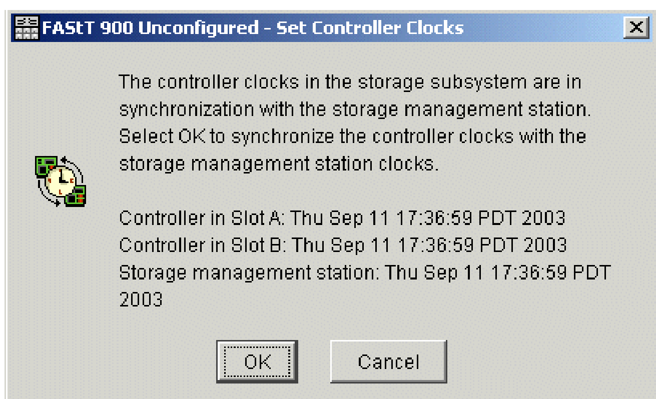

Figure 11-6 Syncronise clocks

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 11-7 on page 339.

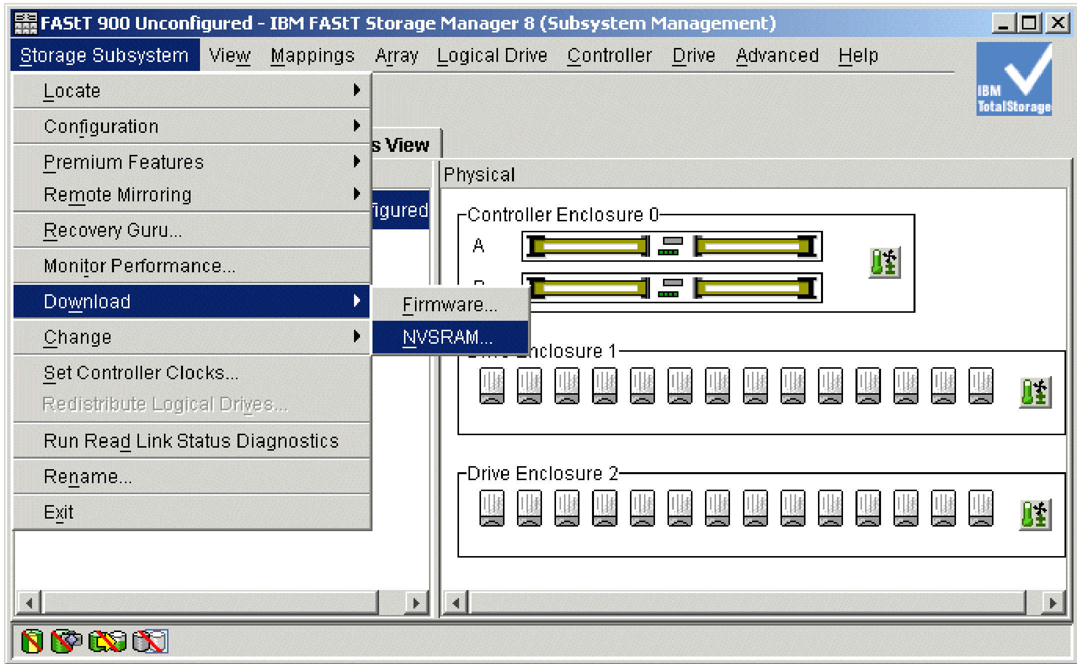

Figure 11-7 Subsystem Management window: Download -> NVSRAM

b. As for the firmware download, choose the correct file you want to download to the controllers (Figure 11-8). The window has the same structure as the Firmware Download dialog (Figure 11-2 on page 336). Therefore, you see the new firmware revision and the NVSRAM revision installed on the storage subsystem.

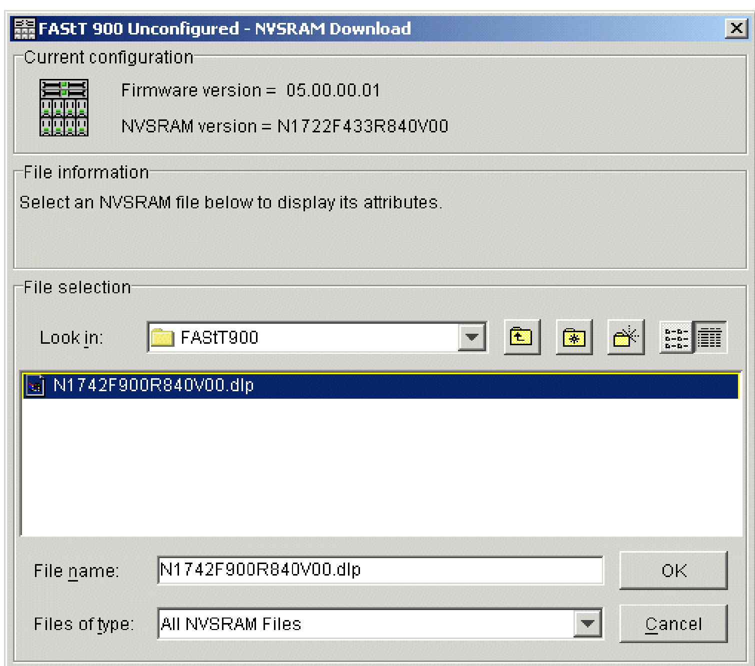

Figure 11-8 NVSRAM Download window

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 11-9 on page 340).

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.

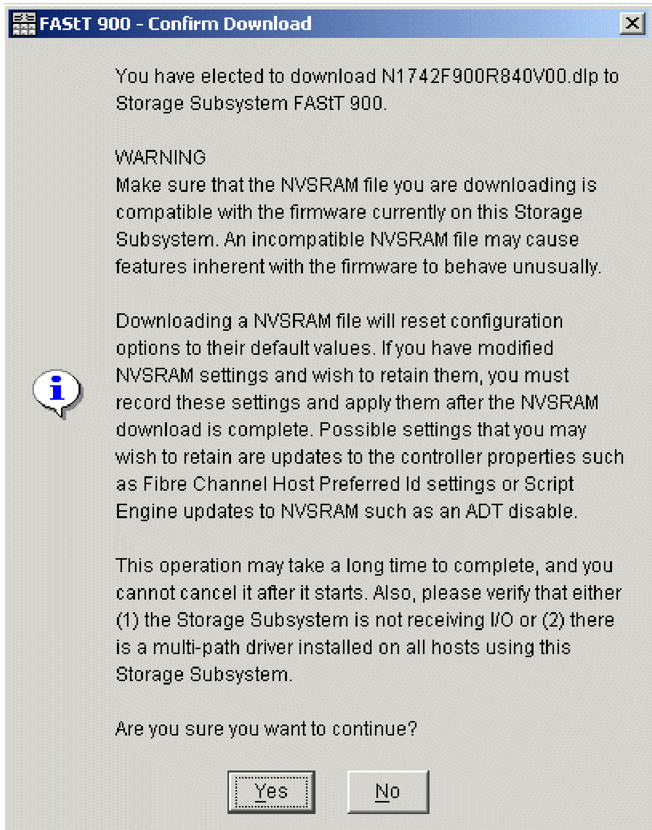

Figure 11-9 Confirm Download of the new NVSRAM

Since the NVSRAM is much smaller than the firmware package, it does not take as long as the firmware download. The window in Figure 11-10 shows the status of the process.

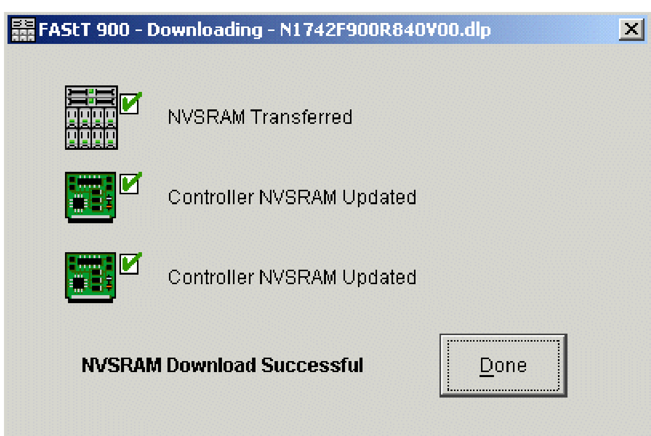

Figure 11-10 Download of NVSRAM completed

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 properly recognizing the updated FAStT.

11.1.4 Updating hard disk drives firmware

After you have updated the controller firmware/NVSRAM, you need to update the hard drives firmware.

The firmware and full instructions can be found at :

http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/disk

Then click your FAStT model and click the approriate version of the device driver matrix for Storage Manager.

Follow these steps:

1. To do this you click the **Advanced** tab, then select **Download**, then select **Drive Firmware**.

> **Attention:** The following information is important:
>
> - Stop all I/O to the storage subsystem prior to the download to prevent application errors.
> - Do not make any configuration changes to the storage subsystem while downloading drive firmware. Doing so could cause the firmware download to fail and make the selected drives unusable.

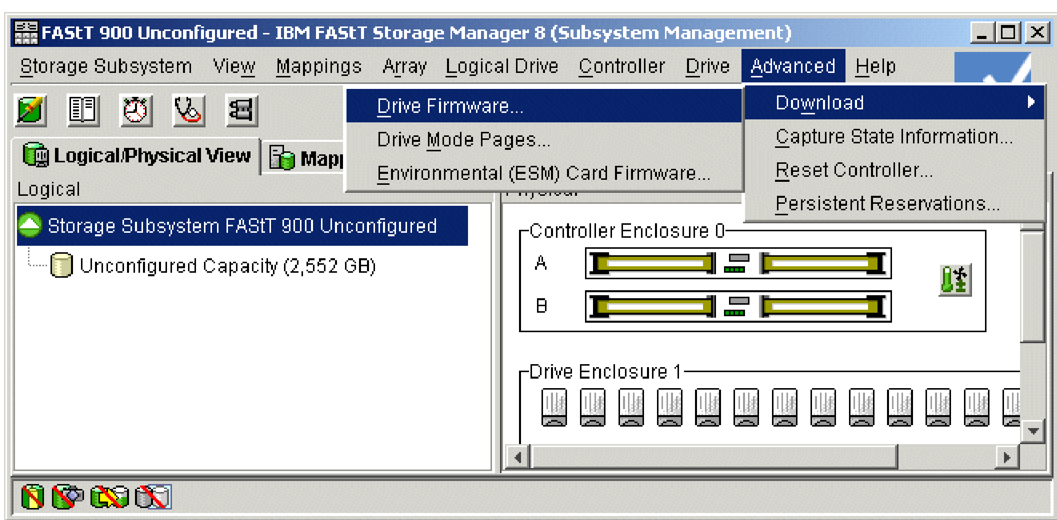

Figure 11-11 Selecting hard drive firmware update

2. The Downlaod Drive Firmware dialog is displayed. In the Select drives area, highlight each drive that you wish to download firmware to or select the **Select All** button to highlight all drives in the storage subsystem.
3. Select the firmware (file) to be applied (enter the name in the Select field or click the Browse buton to select a local or network file).

Attention: It is your responsibility to ensure that the firmware that you download to the drives is compatible with the drives you select.

If downloading firmware to more than one drive, all drives selected must be from the same manufacturer or the Start button will be unavailable. Also, each drive selected should have the same product ID.

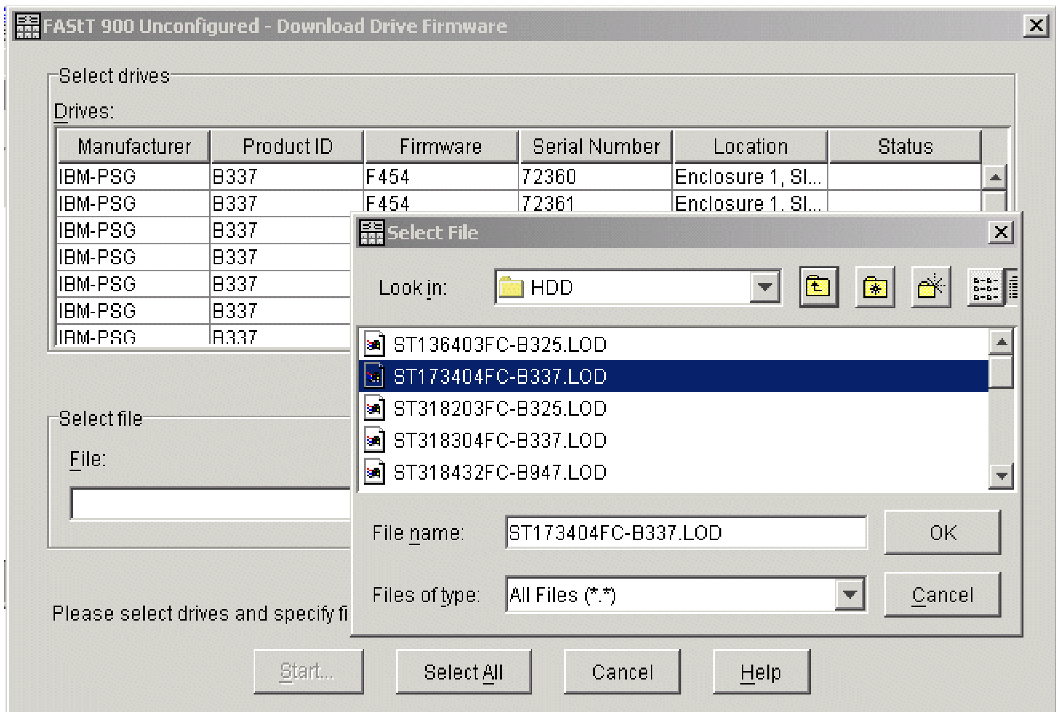

Figure 11-12 Selecting firmware to be applied to hard drives

4. Once you have selected the correct firmware to apply to the hard drives, press the **Start** button.
5. You have to confirm that you want to download the firmware. You must type in `Yes` to proceed and click **OK**.
6. As the firmware is updated, you will see a screen similar to Figure 11-13 on page 343 indicating the status of the update.

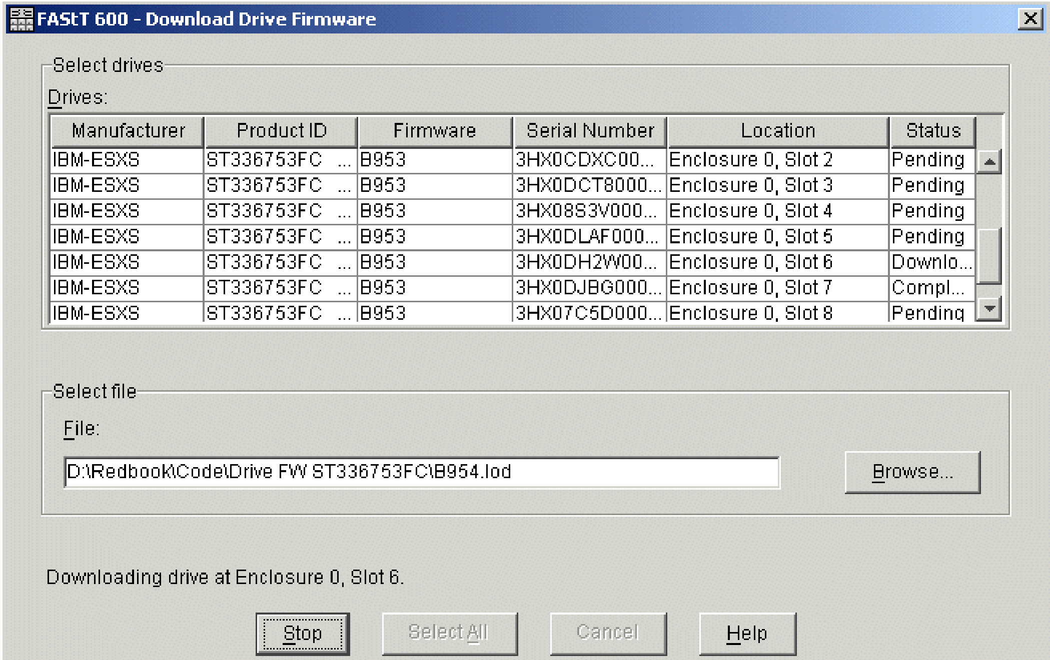

Figure 11-13 Progress of hard drive firmware update

7. A Download complete indication will appear at the bottom left of the dialog once all the selected disks have been processed.

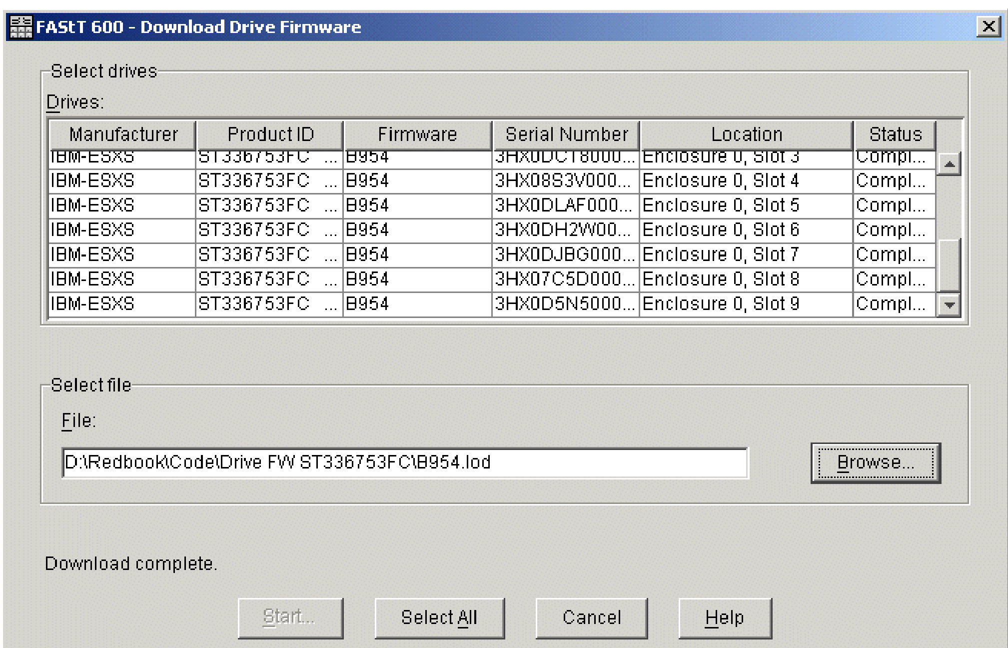

Figure 11-14 Hard drive firmware complete

11.1.5 Updating the ESM board firmware

If you have expansion enclosures, use the following instructions to transfer a downloadable firmware file to Environmental (ESM) Cards.

1. From the Subsystem Management Window, select the **Advanced -> Download -> Environmental (ESM) Card Firmware** pull-down menu option.

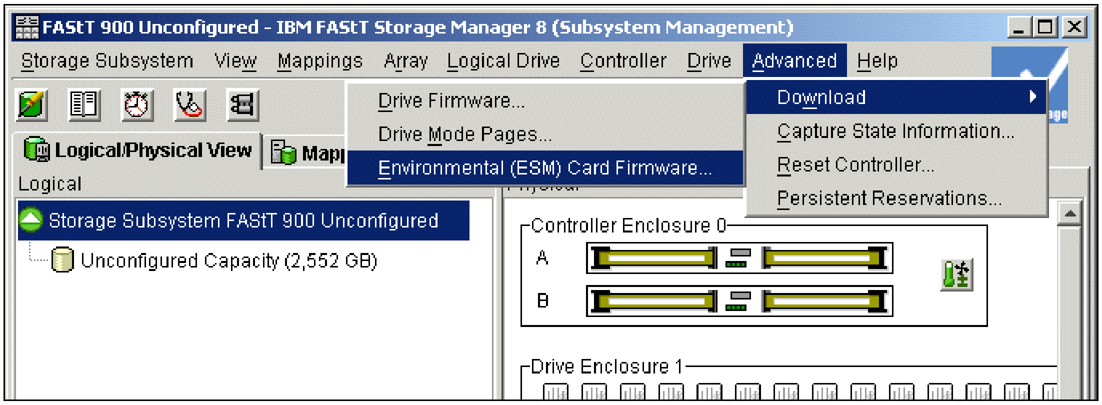

Figure 11-15 Invoke download ESM firmware

2. The Download Environmental (ESM) Card Firmware main screen is displayed. The Select enclosures table lists the details of all enclosures in the storage subsystem that contain ESM cards.

   The Select file area allows you to select an ESM card firmware file for download.

   Note: If an ESM card does not show up in the list (because of a loss of redundancy or some other problem) run the Recovery Guru to diagnose and fix the problem before continuing with the download.

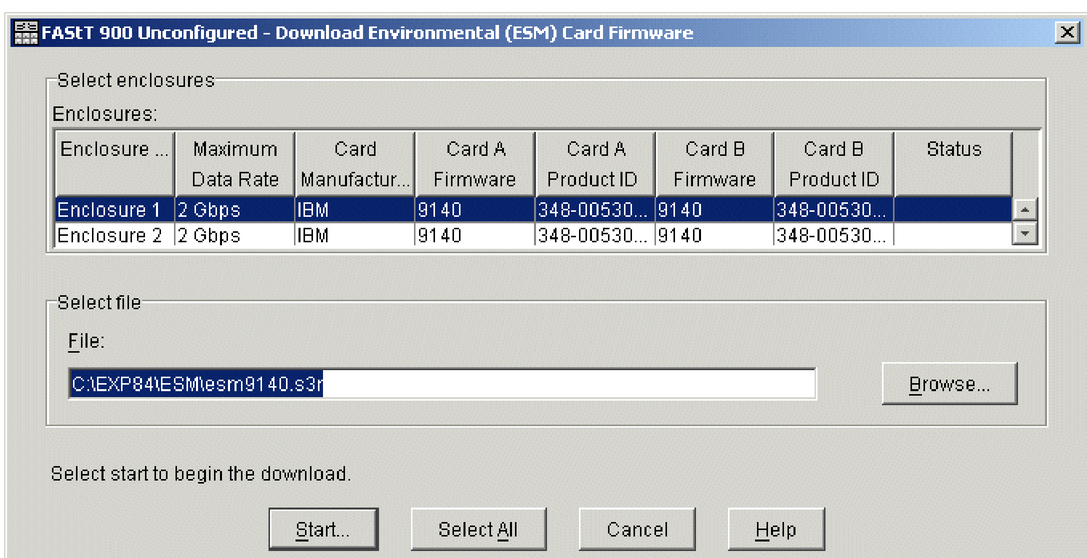

Figure 11-16 Downlaod ESM firmware dialog

3. In the Select enclosures area, highlight each enclosure to which you wish to download firmware or select the Select All button to highlight all drive enclosures in the storage subsystem (each drive enclosure selected should have the same product ID) .
4. Enter the firmware file to download in the Select file area by either entering the location and name of the file in the Select file text box, or selecting **Browse** and getting the firmware file from a local or network drive. (The Browse button will be unavailable until an enclosure has been selected.)
5. Select **Start**. Confirm your selections and then select **Yes** to continue with the firmware download or **No** to quit.
6. The Status field in the Select enclosures table changes from Pending to Downloading for the ESM card firmware operation in progress.

   Monitor the progress and completion status of the download to the enclosures. The progress and status of each drive enclosure participating in the download is displayed in the status field of the Select enclosures table.

11.1.6 Updating the QLogic HBA BIOS

The Utility tab contains options for downloading BIOS flash and the NVSRAM for a single adapter or for both ports on a dual port adapter.

The Flash BIOS and NVSRAM files are part of the IBM FAStT FC-2 adapter files. You can find the latest version on the IBM support side (`http://www.storage.ibm.com/support`. Follow the link for your IBM FAStT Storage Server product and select Storage Manager 8.4 for download.

In order to update the Flash-BIOS click the **Utilities** tab and you will be presented with a screen similar to Figure 11-17.

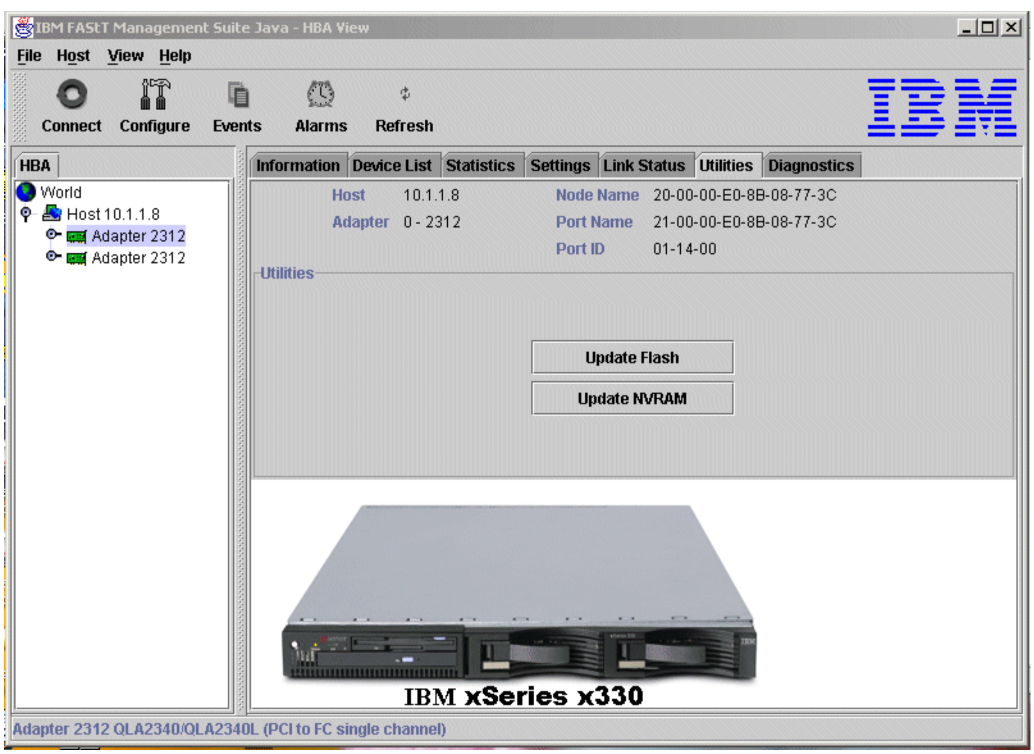

Figure 11-17 Utilities tab of FAStt MSJ

Click **Update Flash**. You will then be presented with a warning, as shown in Figure 11-18. Make sure you have suspended all I/O to the adapter before you proceed.

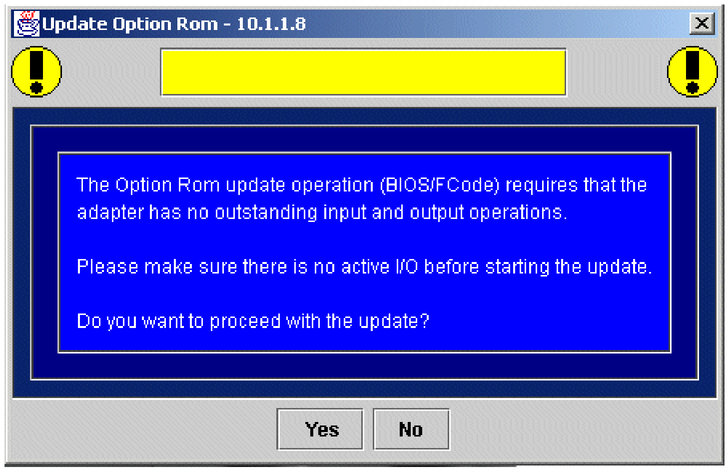

Figure 11-18 IBM FASt MSJ warning before upgrading Qlogic Flash BIOS

Acknowledge the warning and click **Yes**. The following files are needed:

- BIOS update file
- NVSRAM update file

Select the corresponding file for your adapter. In Figure 11-19 we show the download for a Qlogic 231x adapter.

The process for updating the NVSRAM is similar.

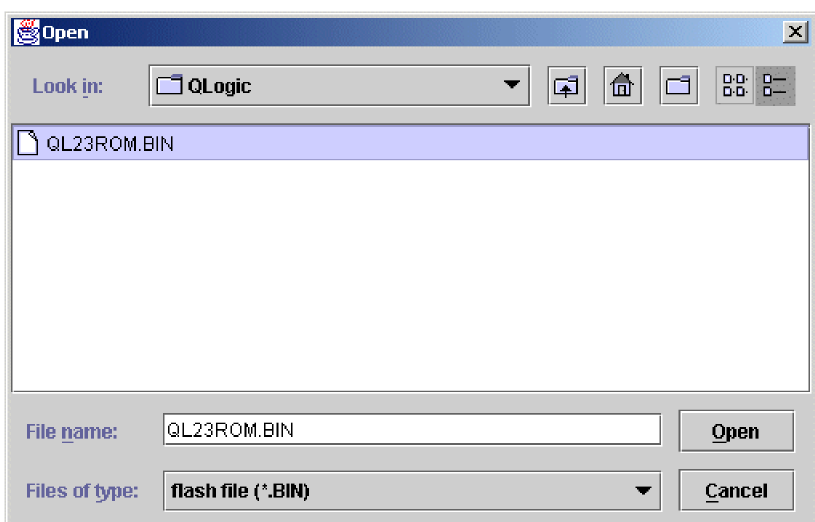

Figure 11-19 IBM FAStT HBA FC-2 BIOS update file

Once you have selected the file you will be prompted for the password. The default password for all adapters is *config*. Once your upload is finished you will be prompted to reboot the server so the changes can take affect. You can upload BIOS and NVSRAM together and then reboot the server. If you do not reboot the server the changes will not take affect.

11.2 Handling premium features

Some features available for the FAStT Storage Server, although included in the firmware, are not enabled.

The premium features are Storage Partitioning, FlashCopy, Remote Volume Mirroring and VolumeCopy.

The storage partitioning premium feature is activated by default. To enable any other premium features, you must obtain (from your storage supplier, that is, IBM) a Feature Key File that is specific to the premium feature you want to enable.

Obtaining a Feature Key File varies depending upon FAStT packaging procedures for the country where the box was purchased and time of order:

- If you bought any premium feature together with the FAStT, the feature key file might be included in the installation package.
- If no Feature Key File has been supplied on the installation media or you are purchasing a new premium feature, you can generate a key on the Web (currently only for FAStT600 and 900) at:

  `http://www.storage.ibm.com/pfeatures.html`
- If no Feature Key File has been supplied on the installation media or you are purchasing a new premium feature, you can get the key from your local storage support (IBM or Business Partner).

  Use the following procedure to obtain a key when contacting your local support:

  a. From the Subsystem Management Window, select either the **Storage Subsystem -> Premium Features -> List** pull-down menu option, or **Premium Features -> List** from the right-mouse pop-up menu (see Figure 11-20).

  b. Write down the number next to Feature Enable Identifier.

  c. Type or copy the Feature Enable Identifier into an email or a text document together with the serial number of the box and version of Storage Manager you are using, and send it to your storage supplier. Be sure to tell them which feature you need enabled.

11.2.1 Listing premium features

You can always check for a list of the premium features that are applied on the FAStT Storage Server. From the Subsystem Management Window, select either the **Storage Subsystem -> Premium Features -> List** pull-down menu option, or **Premium Features -> List** from the right-mouse pop-up menu. As a result the Premium Features dialog is displayed with a list of enabled premium features (see Figure 11-20).

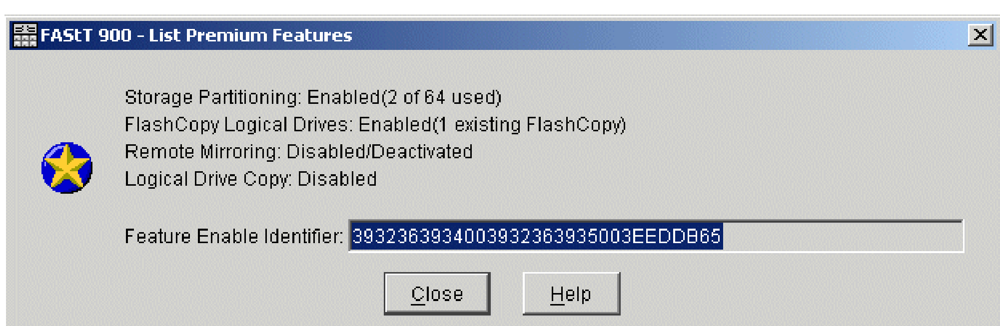

Figure 11-20 Premium features listing and feature enabler key

If you receive a `Premium Features - Out of Compliance` error message during a management session, use the Recovery Guru to resolve the problem.

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. .

11.2.2 Enabling a premium feature

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 11-21.

If the Feature Enable Identifier does not match the FAStT Storage Server, the key is not installed.

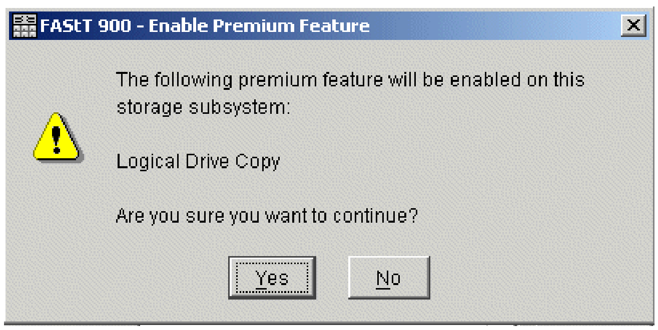

Figure 11-21 Enabling a premium feature

The change happens immediately, and you can use the new functions directly.

11.2.3 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, as shown in Figure 11-22, and confirm.

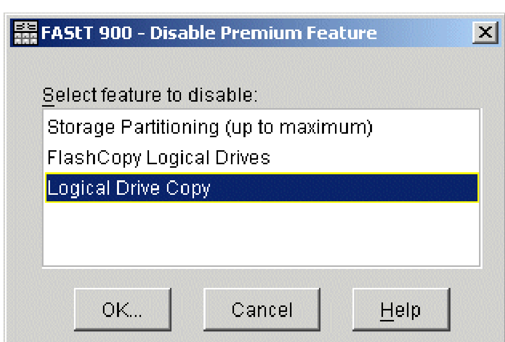

Figure 11-22 Disable Premium Feature

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.

11.3 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

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:** . 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 FAStT 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 and Windows Server 2003- Basic Disks
  - Additional Instructions for Windows NT - Regular Disks
- Windows - Dynamic Disks
  - Process Overview
  - Additional Instructions for Windows 2000 and Windows Server 2003- 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 “Creating arrays and logical drives” on page 144.

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, Windows Server 2003 or Windows NT (using basic or regular disks). See Figure A-1.
- Reuse a disabled FlashCopy logical drive on Windows 2000, Windows server 2003 or Windows NT. See Figure A-2 on page 354.

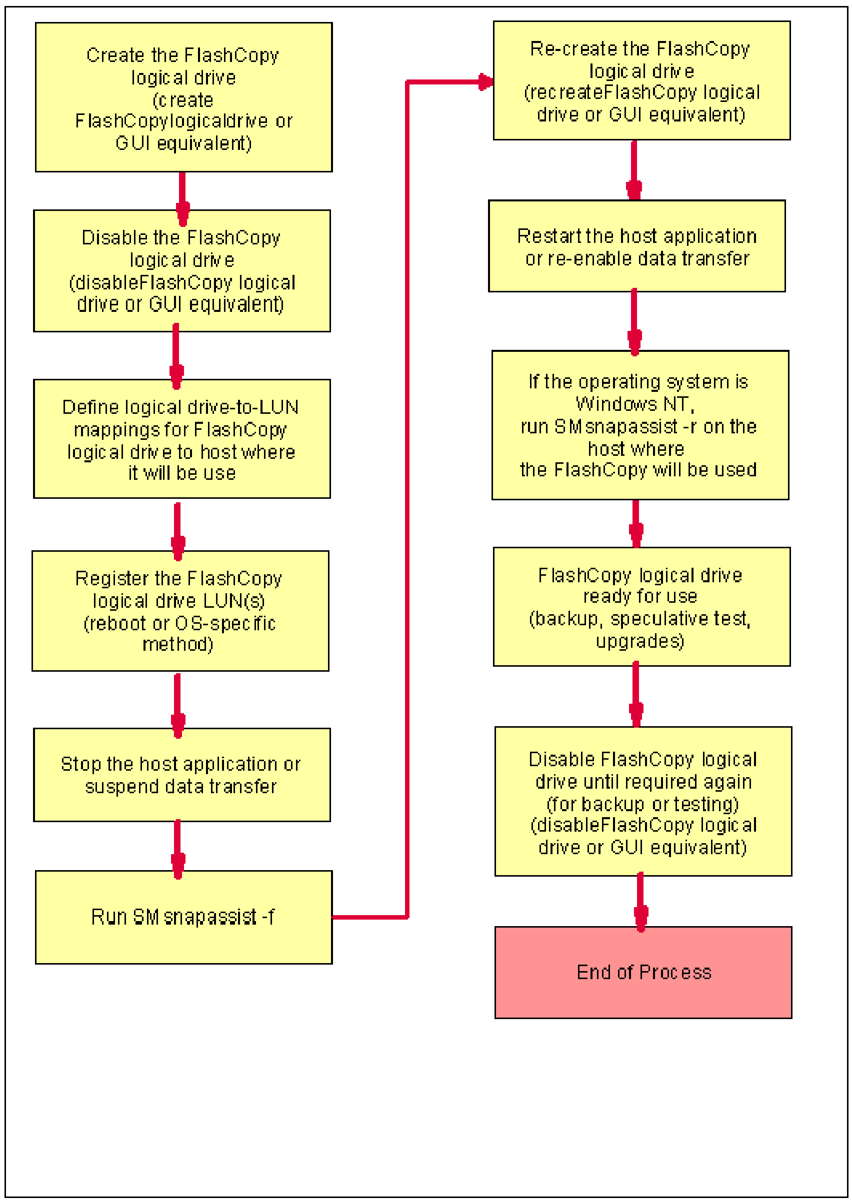

Figure A-1 Creating a FlashCopy logical drive: Windows basic/regular disks

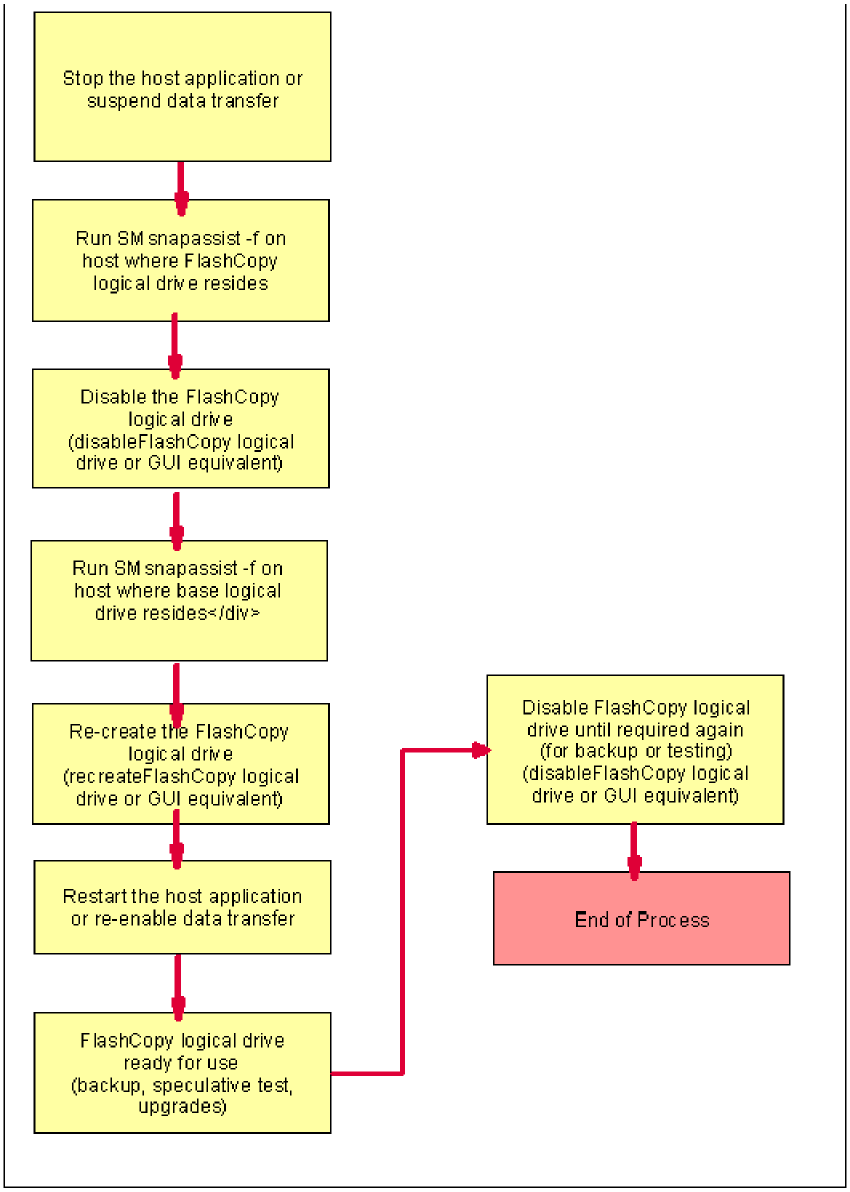

Figure A-2 Re-using disabled FlashCopy logical drives: Windows basic/regular disks

Additional instructions for Windows 2000/2003 basic disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows 2000 or Windows Server 2003, 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 357.

Tip: For command references and information on the use of the script editor and CLI, refer to Chapter 10, "Command Line Interface and Script Editor" on page 311.

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 a 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 that 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 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 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 a 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 that 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 2000/2003 (using dynamic disks). See Figure A-3.
- Reuse a disabled FlashCopy logical drive on Windows 2000/2003 (using dynamic disks). See Figure A-4 on page 363.

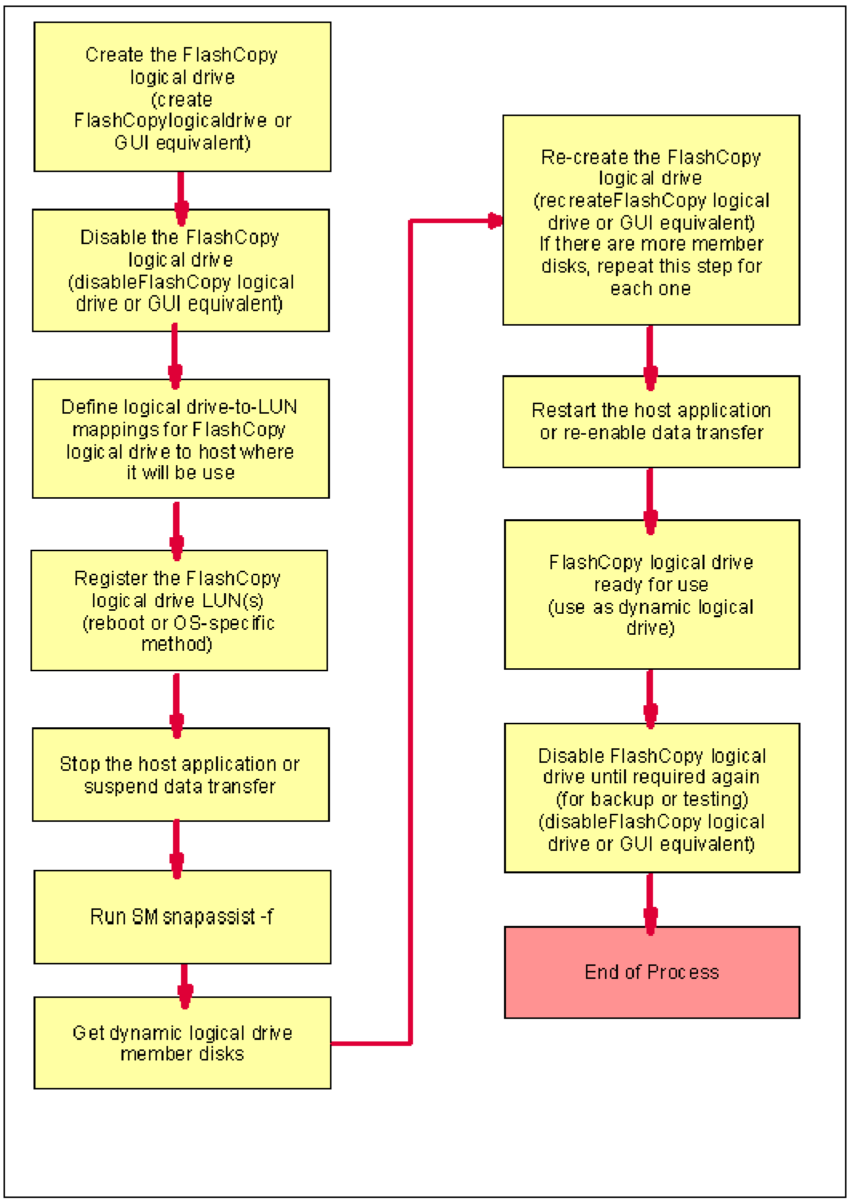

Figure A-3 Creating a FlashCopy logical drive: Windows 2000/2003 dynamic disks

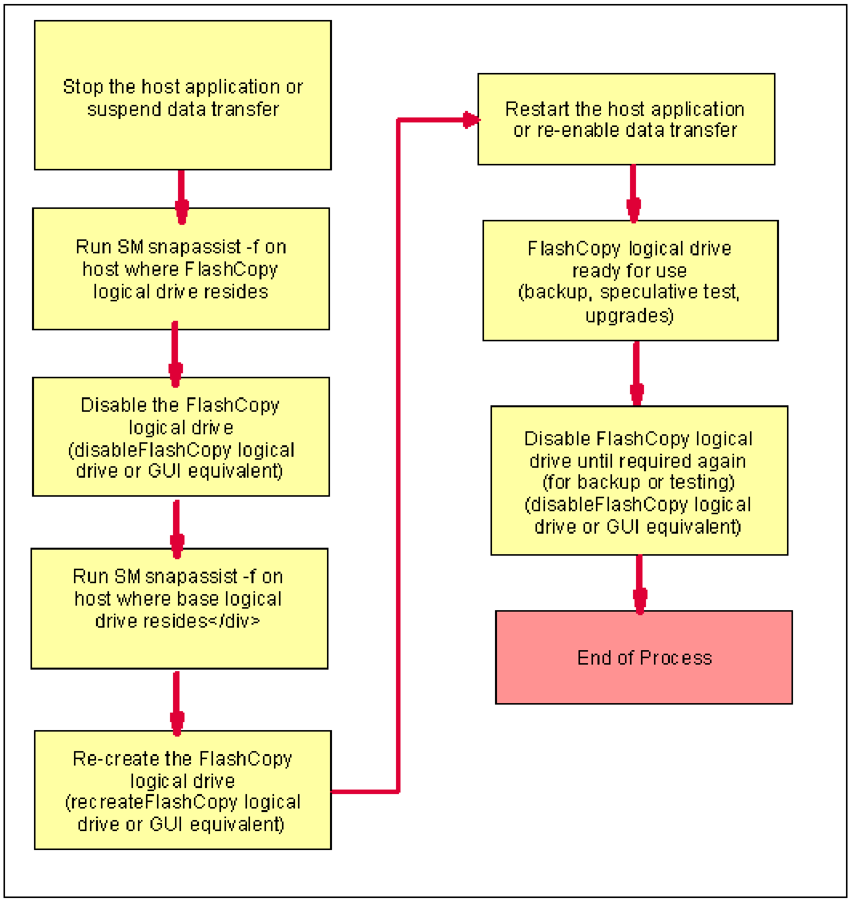

Figure A-4 Reusing a disabled FlashCopy logical drive: Windows 2000/2003 dynamic disks

Additional instructions for Windows 2000/2003 dynamic disks

Use the following procedure when creating FlashCopy logical drives on a host running Windows 2000 or Windows Server 2003, 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 a 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 a 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 window 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)

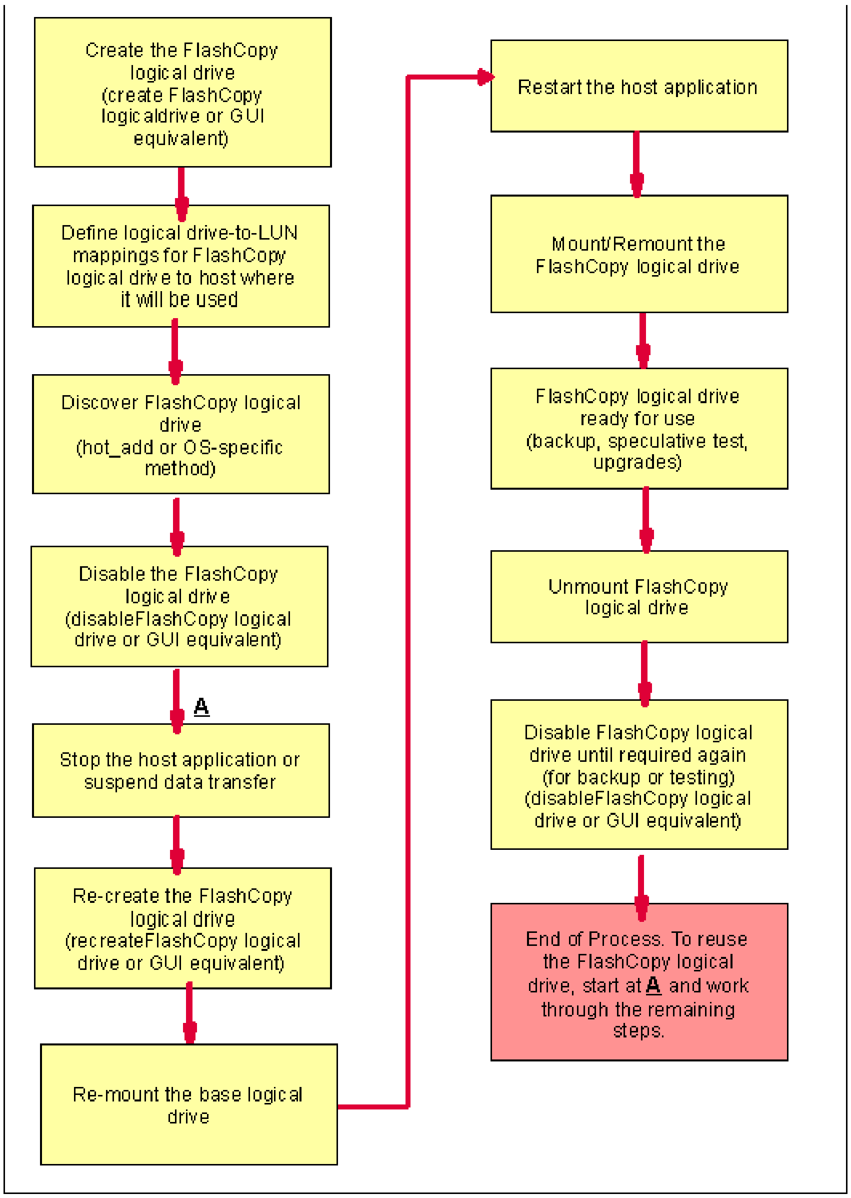

Figure A-5 FlashCopy logical drive on UNIX-based operating systems

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 a 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 a 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/fastt900

   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

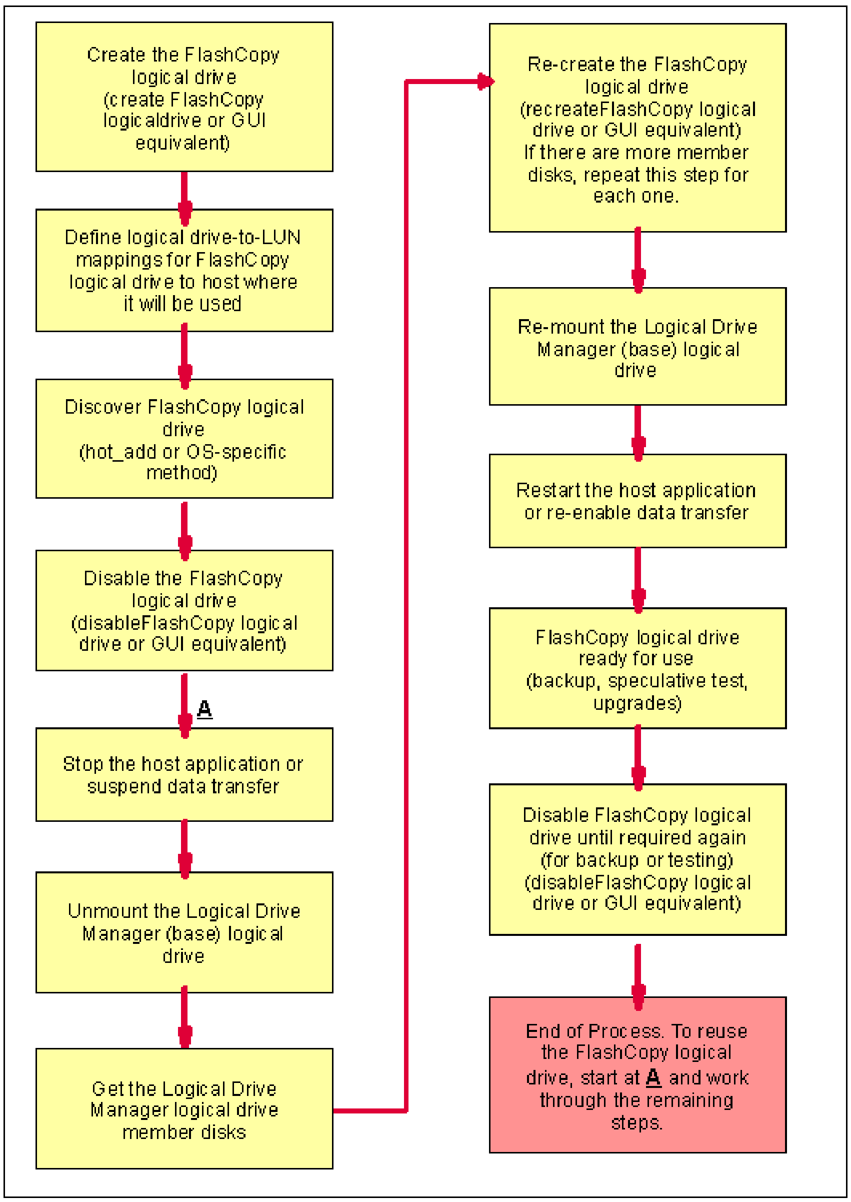

Figure A-6 FlashCopy logical drive UNIX operating systems

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 file systems.

   At the host prompt, type the following command and press Enter:

   ```
   umount mount-point
   ```

   Here *mount-point* is the name of the file system 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 a 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. Log into 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 appear.
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 file system 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 directory name 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 file system. At the host prompt, type the following command and press Enter:

    ```
    umount mount-point
    ```

    Here *mount-point* is the name of the file system 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 file systems 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 file system 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 file system 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 a 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 the state from Disabled to Enabled, and is now ready for use.

12. In the Veritas Logical Drive Manager System Administrator, mount the file system associated with the disk group.

    Select the file system associated with the disk group, right-click, and select **Filesystem -> Mount**. The disk groups associated with the file system are mounted.

13. On the host where the FlashCopy logical drive resides, ensure that the file system 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 file system associated with the disk group:

   Select the file system associated with the disk group. Then right-click and select **Filesystem -> Mount**. The disk groups associated with the file system are mounted.

5. On the host where the FlashCopy logical drive resides, ensure that the file system 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 later) 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 a 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 file system, always perform a `sync` to flush the file system cache immediately prior to creating a FlashCopy logical drive.

   At the host prompt, type the following command and press Enter:

   ```
   sync
   ```

   All unwritten file system 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 file system 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 file system check on it. A file system check is performed to ensure that the file system 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 file system 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 file system check on it.

   A file system check is performed to ensure that the file system 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 on page 386 shows the process flow.

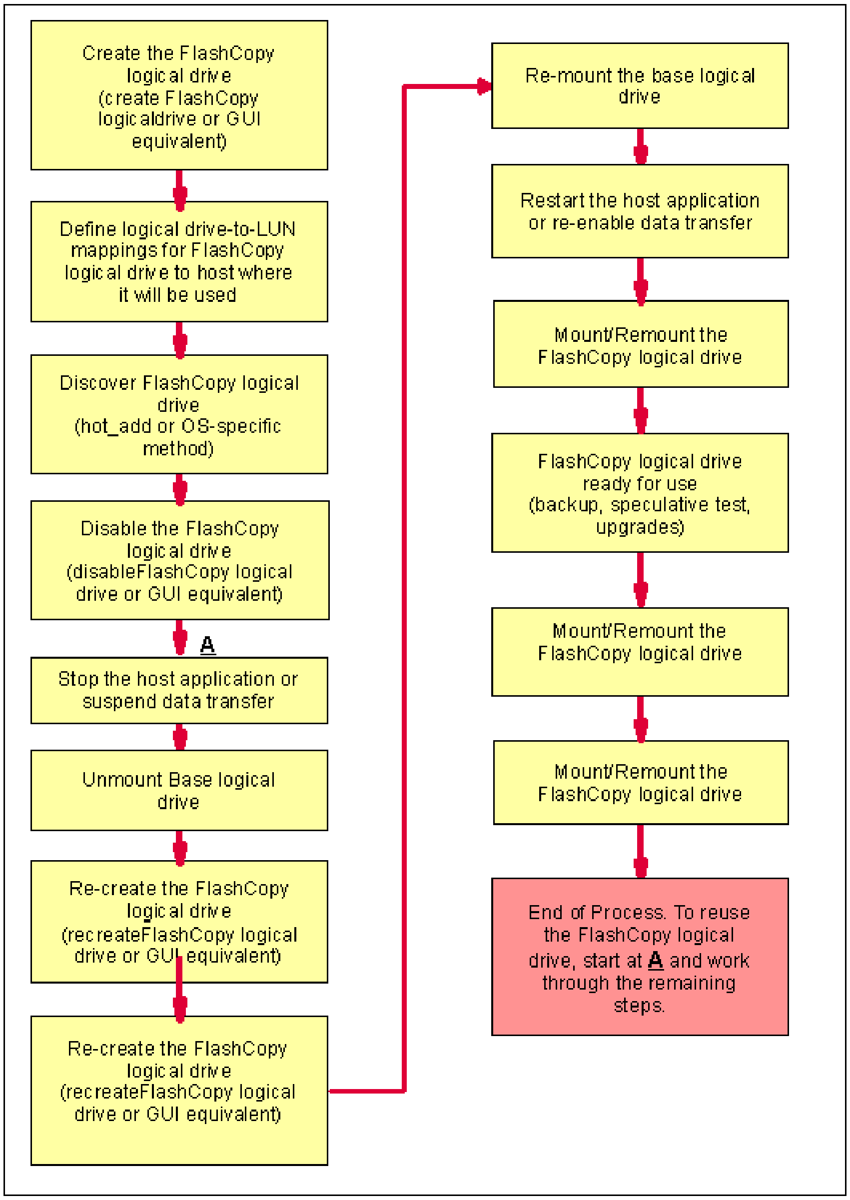

Figure A-7 Creating FlashCopy logical drives on NetWare

Additional instructions for NetWare

Use the following procedure when creating FlashCopy logical drives on a host running NetWare 5.1 (or later). 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 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 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

IBM FAStT Service Alert

FAStT Service Alert is a feature of the IBM TotalStorage FAStT Storage Manager that monitors system health and automatically notifies the IBM Support Center when problems occur. Service Alert sends an e-mail to a call management center that identifies your system and captures any error information that can identify the problem. The IBM Support Center analyzes the contents of the e-mail alert and contacts you with the appropriate service action.

Service offering contract

Following is an overview of the process for obtaining a service offering contract:

1. The account team submits an RPQ requesting Service Alert, using the designated country process.
2. The TotalStorage Hub receives the request and ensures that the following prerequisites are met:
   - The machine type, model, and serial number are provided.
   - The FAStT Storage Server management station is running Storage Manager client Version 8.4 (or at least SM 8.3).
   - The FAStT Storage Server firmware level is at 05.3x.xx.xx or 04.01.xx.xx.
   - The FAStT Storage Server management station has Internet access and e-mail capability.
   - Willingness to sign the contract with the annual fee.
3. After the prerequisites are confirmed, the service offering contract is sent.
4. When the contract has been signed, the approval is sent from the TotalStorage Hub, with the support team copied.
5. Billing is sent at the start of the contract.

Activating FAStT Service Alert

To activate Service Alert, you must do the following tasks:

1. Create a user profile (userdata.txt).
2. Rename each storage subsystem and synchronize the controller clock.
3. Configure the e-mail server.
4. Configure the alert destination.
5. Validate the installation.
6. Test the system.

Prerequisites

Before you use Service Alert, you must install Storage Manager Client Version 8.3 (or later) in the FAStT Storage Server management station. In addition, the FAStT Storage Server firmware levels must be at 05.3x.xx.xx or 04.01.xx.xx.

If your FAStT Storage Server is running with firmware versions 05.0, 05.20, or 05.21, you must download Storage Manager Client 8.3 and firmware version 05.3x.xx.xx. This Storage Manager client and firmware upgrade is provided at no charge. For upgrading Storage manager firmware refer to “Updating the controller microcode” on page 140

Creating a user profile

The user profile (userdata.txt) is a text file that contains your individual contact information. It is placed at the top of the e-mail that Service Alert generates. A template is provided, which you can download and edit using any text editor.

Important: The user profile file name must be userdata.txt. The file content must be in the format as described in step 2. In addition, the file must be placed in the appropriate directory in the FAStT Storage Server management station as indicated in step 4.

Perform the following steps to create the user profile:

1. Download the userdata.txt template file from one of the following Web sites:
   - http://www.ibm.com/storage/fast900
   - http://www.ibm.com/storage/fast700
   - http://www.ibm.com/storage/fast500
   - http://www.ibm.com/storage/fast200

   The userdata.txt template is named "userdata.txt".

2. Type in the required information. There should be seven lines of information in the file. The first line should always be `Title: IBM FAStT Product`. The other line contains the company name, company address, contact name, contact phone number, alternate phone number, and machine location information. Do not split the information for a given item; for example, do not put the company address on multiple lines. Use only one line for each item.

   Note: When you type in the text for the userdata.txt file, the colon (:) is the only legal separator between the required label and the data. No extraneous data is allowed (blanks, commas, and so on) in the label unless specified. Labels are not case sensitive.

   The Title field of the userdata.txt file must always be "IBM FAStT Product". The rest of the fields should be completed for your specific FAStT Storage Server installation.

   Following is an example of a completed userdata.txt user profile.

Example 11-1 Sample userdata.txt

```
Title: IBM FAStT Product Company
name: IBM (73HA Department)
Address: 3039 Cornwallis Road, RTP, NC 27709 Contact
name: John Doe
Contact phone number: 919-254-0000
Alternate phone number: 919-254-0001
Machine location: Building 205 Lab, 1300
```

3. Save the userdata.txt file in ASCII format.

4. Store the userdata.txt file in the appropriate subdirectory of the FAStT Storage Server management station, depending on the operating system that is installed in the management station:
   - For Microsoft, Windows 2000, and Windows NT4, store the userdata.txt file in the %SystemRoot%\java\ directory if Event Monitor is installed or, if Event Monitor is not installed, in the Installed_Windows_driveletter:\Documents and Settings\Current_login_user_folder directory.
   - If your Windows 2000 or Windows NT4 installation uses the default installation settings, and the current login user ID is Administrator, the directories are c:\WINNT\java or c:\Documents and Settings\Administrator, respectively.
   - For AIX(R), store the userdata.txt file in the / directory.

– For RedHat Advance Server, store the userdata.txt file in the default login directory of the root user. In normal installation, this directory is /root.
– For SuSe 8, store the userdata.txt file in the default login directory of the root user. In normal installation, this directory is /root.
– For Novell NetWare, store the userdata.txt file in the sys:/ directory.

> **Note:** You must have a Storage Manager client session running to monitor failures of the FAStT Storage Server.

– For Solaris, store the userdata.txt file in the / directory.
– For HP-UX, store the userdata.txt file in the / directory.
– VMware ESX servers that are connected to a FAStT Storage Server will require a separate workstation for FAStT Storage Server management. Service Alert is only supported in a VMware ESX and FAStT environment by way of the remote management station.

Renaming the storage server and synchronizing the controller clock

When you register for Service Alert, you must change the existing node ID of each FAStT Storage Server. Service Alert uses this new name to identify which FAStT Storage Server has generated the problem e-mail. To rename the Storage Server refer to 5.3.2, “Starting the FAStT Storage Manager client” on page 137. Before you rename the storage subsystem, record the FAStT Storage Server machine type, model, and serial number.

Type the new name for the subsystem. You must use the following naming convention for the new name; any errors in the format of the new name can result in delays or denial of IBM service support. The new name cannot contain more than 30 characters. The format for the new name is:

ttttmmm/sssssss#cust_nodeid_reference

Where:

- *tttt* is the 4-digit IBM machine type of the product
- *mmm* is the 3-digit IBM model number for the product
- */* is the required separator
- *sssssss* is the 7-digit IBM serial number for the machine
- *#* is the required separator
- *cust_nodeid_reference* is the node ID as referenced by the customer

> **Important:** No extra characters are allowed before the "#" separator.

Following is a list of FAStT machine types and model numbers that you can use for reference.

Table 11-2 FAStT Storage Server machine and model numbers

Product	Machine type	Model number	Model number in your name
FAStT900	1742	90U, 90X	900
FAStT700	1742	1RU, 1RX	000
FAStT500	3552	1RU, 1RX	000

Product	Machine type	Model number	Model number in your name
FAStT200	3542	1RU, 1RX 2RU, 2RX	000

Following are some examples of a new storage subsystem name:

- 1742900/23A1234#IBM_Eng
- 1742000/23A1235#IBM_Acctg
- 3552000/23A1236#IBM_Mktg
- 3542000/23A1237#IBM_Mfg

Click **OK** to save the new name.

To synchronize the controller clock with the time in the FAStT Storage Server management station that monitors the alerts, refer to "Starting the FAStT Storage Manager client" on page 137

This step is optional. If performed, it facilitates the trouble-shooting session because the time that the alert e-mail is sent is about the same as the time that the errors occurred in the FAStT Storage Server.

All of the steps must be performed for each of the FAStT Storage Servers that support Service Alert.

Configuring the e-mail server

You must configure your e-mail server to enable it to send alerts. Refer to "Monitoring and alerting" on page 157 for instrustions on how to do this.

The email address you enter will be used to send every alert that is sent.

Configuring the alert destination

Refer to "Monitoring and alerting" on page 157 for instructions on how to do this.

In the E-mail address text box, type one of the following e-mail addresses, depending on your geographic location:

- For EMEA and A/P locations: callhome0@de.ibm.com
- For North America locations: callhome1@de.ibm.com
- For South and Central America, and Caribbean Island locations: callhome1@de.ibm.com

Validating the installation

Make sure that the IBM FAStT Event Monitor service is installed in the management station. If it is not installed, you must uninstall the FAStT Storage Manager client and reinstall it with the Event Monitor service enabled.

> **Note:** FAStT Event Monitor service is not supported on Novell Netware 6. You must use a management station with another operating systems installed, such as Windows 2000.

Testing the system

After all previous tasks have been completed, you are ready to test your system for Service Alert.

Call your IBM Support Center. Tell the representative that you are ready to test the Service Alert process. The IBM representative will work with you to test your system setup and ensure that FAStT Service Alert is working properly.

A test that you will perform, with the help of the Support Center, is to manually fail a non-configured drive in the FAStT Storage Server using the FAStT Storage Manager client. If all of the drives are configured, you can turn off a redundant power supply in the FAStT Storage Server or FAStT expansion enclosure. When the drive fails or the power supply is turned off, a Service Alert is sent to the IBM e-mail address that you specified in "Configuring the e-mail server" on page 395.

Note: Do not turn off the power supply if this is the only one that is powered on in your storage server or expansion enclosure. Turning off the power supply is the preferred test because it allows the testing of the FAStT Storage Server Event Monitor service. This service monitors the FAStT Storage Server for alerts without needing to have the FAStT Storage Manager client running in the root user session.

.

C

Recovery Guru events and critical events

This appendix provides a listing of the Recovery Guru events and information about critical events.

Recovery Guru events

Battery Nearing Expiration
Battery Canister Nearing Expiration
By-passed Drive
Channel Miswire
Controller Memory Parity Error
Data Rate Negotiation Failed
Degraded Logical Drive
Drive Channel Minihub Data Rate Mismatch
Drive - Loss of Path Redundancy
Drive Enclosure - Loss of Path Redundancy
Dual Primary Logical Drive Conflict
Dual Secondary Logical Drive Conflict
ESM - Loss of Communication
ESM Canister Firmware Version Mismatch
ESM Canister Miswire
Failed Battery
Failed Battery Canister
Failed Drive - Unassigned or Standby Hot Spare
Failed Drive SCSI Channel
Failed ESM Canister
Failed or Removed Fan Canister
Failed GBIC/SFP
Failed Minihub Canister
Failed FlashCopy Logical Drive
Failed Logical Drive - During Modification
Failed or Removed Power Supply Canister
Failed Logical Drive - Drive Failure
Failed Logical Drive - Awaiting Initialization
Fibre Channel Link Errors - Threshold Exceeded
Incompatible Minihub Canister
Impending Drive Failure (High Data Availability Risk)
Impending Drive Failure (Medium Data Availability Risk)
Impending Drive Failure (Unassigned or Standby Hot Spare)
Lost AC Power
Maximum Temperature Exceeded
Mirror Communication Error - Unable to Contact Fabric
Mirror Communication Error - Unable to Contact Storage Subsystem
Mirror Communication Error - Unable to Contact Logical Drive
Mirror Data Unsynchronized
Missing Mirror Repository Logical Drive
Missing Logical Drive
Nominal Temperature Exceeded
Non-Compliant Partitions
Non-Compliant FlashCopy
Offline Controller
Offline Array
Remote Mirroring Feature - Out of Compliance
Removed Battery Canister
Removed ESM Canister
FlashCopy Repository Logical Drive Capacity Full
FlashCopy Repository Logical Drive Capacity Threshold Exceeded
Enclosure ID Conflict
Enclosure ID Mismatch

Unable to Update Remote Mirror
Uncertified Drive
Unknown Failure Type
Unrecovered Interrupted Write
Logical Drive - Hot Spare In Use
VolumeCopy Operation Failed
VolumeCopy Premium Feature - Out of Compliance

Storage Manager Critical Events

This section 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.

Table 11-3 Critical events

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 anymore. 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 that 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:** Communication 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.

..

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 410. Note that some of the documents referenced here may be available in softcopy only.

- *IBM SAN Survival Guide,* SG24-6143
- *Fibre Array Storage Technology A FAStT Introduction*, SG24-6246
- *IBM TotalStorage Solutions for xSeries,* SG24-6974
- *IBM TotalStorage: Implementing Linux with IBM Disk Storage*, SG24-6261

Other publications

These publications are also relevant as further information sources:

- *IBM TotalStorage FAStT Best Practices Guide*, REDP-3690
- *IBM TotalStorage FAStT Storage Manager 8.4 Library Guide and Common Index,* GC26-7600
- *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide,* GC26-7561
- *IBM TotalStorage FAStT Storage Manager Concepts Guide,* GC26-7560
- *IBM TotalStorage FAStT Storage Manager 8.4 Installation and Support Guide for Intel-based Operating System Environments,* GC26-7589

Online resources

These Web sites and URLs are also relevant as further information sources:

- FAStT product documentation

  http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt/
- IBM TotalStorage

  http://www.ibm.com/storage/
- IBM FAStT Storage products

  http://www.storage.ibm.com/disk/fastt/index.html/
- IBM FAStT Storage interoperability matrix

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

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

S

T

U

V

W

IBM

Redbooks

IBM TotalStorage: FAStT600/900 and Storage Manager 8.4

IBM TotalStorage: FAStT600/900 and Storage Manager 8.4

New FAStT models

Storage Manager V8.4

VolumeCopy

This IBM Redbook reviews, in detail, the hardware of the IBM TotalStorage FAStT600 and FAStT900 Storage Servers and presents the new features introduced with the FAStT Storage Manager Version 8.4 (SM 8.4).

After reviewing the concepts and functions used in planning and managing the FAStT Storage Server, this publication offers a step-by-step guide to using the Storage Manager to create arrays, logical drives, and other basic (as well as advanced) management tasks.

This publication also contains practical information on diagnostics and troubleshooting, critical event descriptions, and usage of scripts, and the Command Line Interface.

Finally, this publication includes chapters dedicated to the Copy Services premium features, and covers the VolumeCopy feature provided in the latest release of the Storage Manager software.

This publication 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 V8.4. 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

SG24-7010-00 ISBN 0738499064