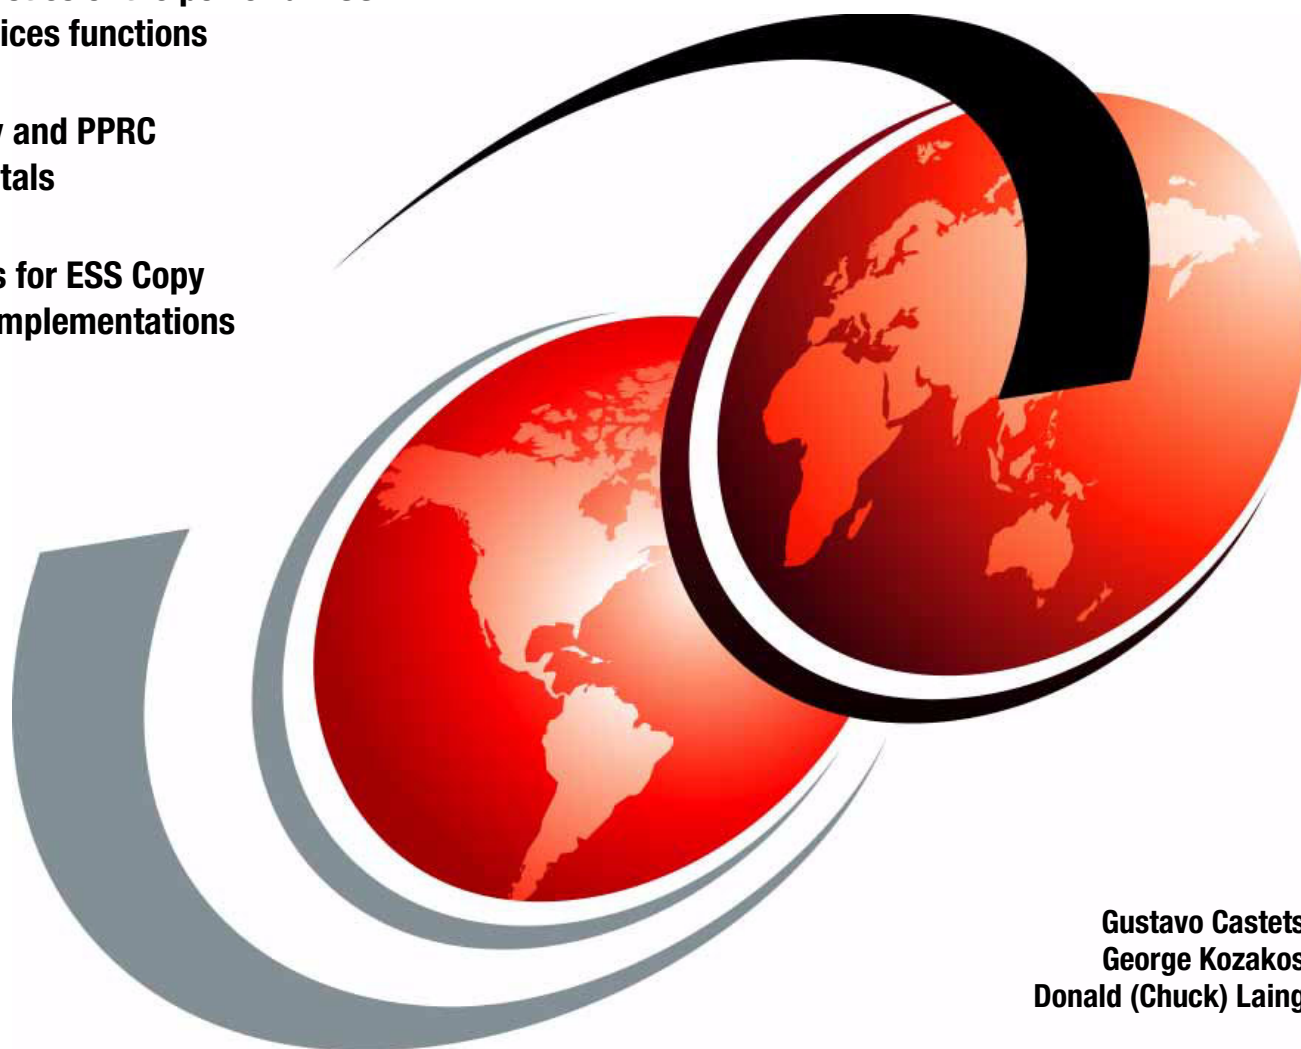# IBM

IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments

Characteristics of the powerful ESS Copy Services functions

FlashCopy and PPRC fundamentals

Guidelines for ESS Copy Services implementations

Gustavo Castets
George Kozakos
Donald (Chuck) Laing

Redbooks

IBM

International Technical Support Organization

IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments

February 2004

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

Fourth Edition (February 2004)

This edition applies to the IBM TotalStorage Enterprise Storage Server Model 800 and F.

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:

Redbooks(logo) ™	Enterprise Storage Server®	Redbooks™
@server™	ESCON®	S/390®
iSeries™	FlashCopy®	Seascape®
xSeries®	FICON™	Tivoli®
z/OS®	GDPS®	TotalStorage®
z/VM®	IBM®	VSE/ESA™
zSeries®	MVS™	Wave®
DB2®	OS/390®	

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), and Pentium 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.

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

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

Preface

This IBM® Redbook describes the copy functions available with the IBM TotalStorage Enterprise Storage Server (ESS). The powerful ESS Copy Services functions are explained in detail, and their respective characteristics are thoroughly covered. This redbook also gives information about how to manage the various ESS Copy Services functions, and finally discusses their implementations.

Because this redbook provides a broad understanding of the ESS Copy Services functions, as well as going into detail about the management interfaces and the implementation considerations, we recommend it for IT professionals who are planning the implementation of any of the ESS Copy Services functions in an open-systems environment.

This fourth edition of the redbook covers all of the functions and features of ESS Copy Services that were made available on the ESS with LIC level 2.2.0. and FlashCopy® Version 2 and PPRC Version 2. These functions include:

► FlashCopy Version 2 options: Incremental, Multiple Relationship, elimination of the Logical Subsystem (LSS) constraint, and Consistency Groups (Note: for S/390®, there is also Data Set FlashCopy, which will not be covered in this book).

► PPRC Version 2: Asynchronous Cascading PPRC.

This fourth edition of the redbook has also been updated with the latest ESS Copy Services functions available with LIC level 2.3.0:

► PPRC over Fibre Channel, enabling the use of Fibre Channel Protocol (FCP) as a communications link between the ESS PPRC primary and secondary machines.

► ESS Application Programming Interface (ESS API) support for ESS Copy Services configuration and use.

► Asynchronous Cascading PPRC enhancements, exploiting PPRC Failover and Failback modes.

► Copy Services domain support of up to eight ESSs per domain.

► Also, some Web user interface panels have changed with the introduction of LIC 2.3.0, and these changes are presented in this edition of the redbook.

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.

Gustavo Castets is a Certified IBM IT Specialist and Project Leader at the IBM International Technical Support Organization, San Jose Center. He has co-authored eight Redbooks™ and teaches IBM classes worldwide in the area of disk storage systems. Before joining the ITSO, Gustavo was based in Buenos Aires, where he worked in different technical sales support positions for more than 22 years for IBM Argentina.

George Kozakos is an OS/390® Technical Consultant in Australia. He has 17 years of experience in MVS™ systems programming. He holds a B.Sc. (Hons) degree in pure mathematics from Adelaide University. His areas of expertise include PPRC, XRC, and GDPS®.

Donald (Chuck) Laing is a Sr. Systems Management Integration Professional, specializing in open systems UNIX® disk administration in the IBM South Delivery Center (SDC). He has co-authored three previous IBM Redbooks on the IBM TotalStorage® Enterprise Storage Server®. He holds a degree in Computer Science. Chuck's responsibilities include planning and implementation of midrange storage products. His responsibilities also include department wide education and cross training on various storage products, such as the ESS and FAStT. He has worked at IBM for five and a half years. Before joining IBM, Chuck was a hardware CE on UNIX systems for 10 years and taught basic UNIX at Midland College for six and a half years in Midland, Texas.

Figure 1   The authors of the current version: Gustavo, George, and Chuck

The authors of previous editions of this book are:

Barry Mellish	Paresh Chudasma	Susanne Lukas	Miassimo Rosichini
Andrew Beyer	Rainer Wolafka	Vladmir Blazek	Cathy Warrick
Olivier Alluis	Torsten Rothenwaldt	Jiri Spacek	Nancy Roper
Toga Gursu	Carlin Smith		

Figure 2   Previous edition: (back row) Carlin, Cathy, Torsten (front row) Tolga, Jiri, Olivier

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

Werner Bauer	Steve Chase	James Goodwin
Craig Gordon	Mary Koegl	Denise Luzar
Ian Mac Quarrie	Gail Spear	John Wrobel

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, IBM 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

Summary of changes

This section describes the technical changes made in this edition of the book and in previous editions. This edition may also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-5757-03
for IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments
as created or updated on February 17, 2004.

February 2003, Fourth Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

► PPRC over Fibre Channel, enabling the use of Fibre Channel Protocol (FCP) as a communications link between the ESS PPRC primary and secondary machines

► ESS Application Programming Interface (ESS API) support for ESS Copy Services configuration and use

► Asynchronous Cascading PPRC enhancements, exploiting Failover and Failback modes

► Copy Services domain now supporting up to eight ESSs per domain

Changed information

► ESS Copy Services Web user interface, panels and options
► PPRC Failover and Failback descriptions and procedures
► ESS Model 800 characteristics

1

Introduction

In this chapter, we review the IBM Seascape® Architecture and the IBM TotalStorage Enterprise Server (ESS). We introduce the features of the ESS Copy Services functions and discuss, at a high level, some of the benefits, such as disaster recovery. We include the new features of FlashCopy Version 2 and PPRC Version 2:

- ► Data Set FlashCopy (only for zOS)
- ► Incremental FlashCopy
- ► Multiple Relationship FlashCopy
- ► FlashCopy across Logical Subsystems (LSSs)
- ► Asynchronous Cascading PPRC

1

1.1 Overview

As businesses become more and more dependent on information technology to conduct their operations and stay competitive, the availability of their processing facilities becomes crucial. Today, most businesses require a high level of availability, which extends to continuous availability, 24 hours a day and seven days a week (24x7) operation. A lengthy outage could lead to significant financial losses, loss of credibility with customers, and maybe even a total failure of business. Therefore, the ability to provide continuous availability for the major applications is more often than not a necessity for business survival. A key component of a highly available system is the storage subsystem. It is essential that data is available at all times and that downtime for data backup and software maintenance is minimized or more preferably eliminated.

These important demands on the storage system are fulfilled with the Enterprise Storage Server (ESS) and its Copy Services functions. The ESS Copy Services provide replication of mission critical data, point-in-time and incremental FlashCopy, dynamic synchronous mirroring to a remote site with Peer-to-Peer Remote Copy (PPRC), asynchronous copying to a remote site with Extended Remote Copying (XRC) for S/390 only, Peer-to-Peer Remote Copy Extended Distance (PPRC-XD), and Asynchronous Cascading PPRC.

1.2 Seascape Architecture

Seascape, the IBM storage enterprise architecture, is a blueprint for comprehensive storage solutions optimized for a connected world. The Seascape architecture outlines new concepts for storage by integrating leading technologies from IBM — including disk, tape, optical, powerful processors, and rich software function — to provide highly reliable, scalable, and versatile application-based storage solutions that span the range of servers from PCs to supercomputers.

Seascape solutions usher in a new era of data storage by helping your organization improve the ways it accesses, grows, manages, and secures data while protecting current storage investments. This storage architecture can help you implement a simplified, yet flexible storage infrastructure that will help you derive the most value from your information.

Here are the highlights of the Seascape architecture:

► It provides comprehensive storage solutions that integrate rich software, enterprise storage management, and a full range of disk, tape, and optical storage technology.

► It spans the range of servers from UNIX, Windows® NT, and midrange to S/390.

► It facilitates a wide array of connectivity and software functions to provide the foundation for universal data access and information sharing.

► It employs "snap-in" building blocks that enable rapid integration of new technologies while protecting existing storage investments.

► It provides solutions tuned to specific application needs.

► It leverages industry-leading technologies that enable the development of comprehensive, cost-effective storage solutions.

There are three main principles of the Seascape architecture:

1. Universal data access
2. Snap-in building blocks
3. Integrated storage servers

The implementation of the Seascape architecture in the ESS is illustrated in Figure 1-1.

Powerful Seascape Architecture

RAID 5 and RAID 10 Arrays Across Loops

Storage sharing among different platforms: IBM e-business servers and non-IBM servers

SSA loops with faster and greater capacity disk drives

zSeries parallel I/O

More powerful SMP processors with Turbo option

Incremental & Point-in-time copy solutions

IBM TotalStorage Enterprise Storage Server Model 800

Fault tolerant system

Disaster Recovery solutions: synchronous non-synchronous

High capacity scalable storage - up to 55.9 TB physical capacity

Figure 1-1   IBM Seascape Architecture

The Seascape architecture is the key to the development of IBM storage products. Seascape allows IBM to take the best of the technologies developed by the many IBM laboratories and integrate them, producing flexible and upgradeable storage solutions. This Seascape architecture design has allowed the IBM TotalStorage Enterprise Storage Server to evolve from the initial E models to the succeeding F models, and to the recently announced 800 models, each featuring new, more powerful hardware and functional enhancements, and always integrated under the same successful architecture with which the ESS was originally conceived.

1.3 IBM TotalStorage Enterprise Storage Server overview

The move to e-business on demand presents companies with both extraordinary opportunities and significant challenges. A whole new world of potential customers, automated and streamlined processes, and new revenue streams are being fueled by e-business. Consequently, companies also face an increase of critical requirements for more information that is universally available online, around the clock, every day of the year.

The IBM TotalStorage Enterprise Storage Server (ESS) provides un-matchable functions for all of the server family of e-business servers, and also for the non-IBM (that is, Intel®-based and UNIX-based) families of servers. Across all of these environments, the ESS features unique capabilities that allow it to meet the most demanding requirements of performance, capacity, and data availability that the computing business may require. The ESS Model 800 is the latest IBM storage product to be developed using IBM Seascape architecture.

As shown in Figure 1-2, the Enterprise Storage Server is the natural successor to the 3990 for IBM @server zSeries® and to the 7133 storage system in the open systems world. It provides all the functions that were available on the 3990, including peer-to-peer remote copy (PPRC), FlashCopy, extended remote copy (XRC), concurrent copy, and Peer-to-Peer Extended Copy (PPRC-XD) and brings them to open systems and iSeries™. Note that XRC is not available for open systems or iSeries.
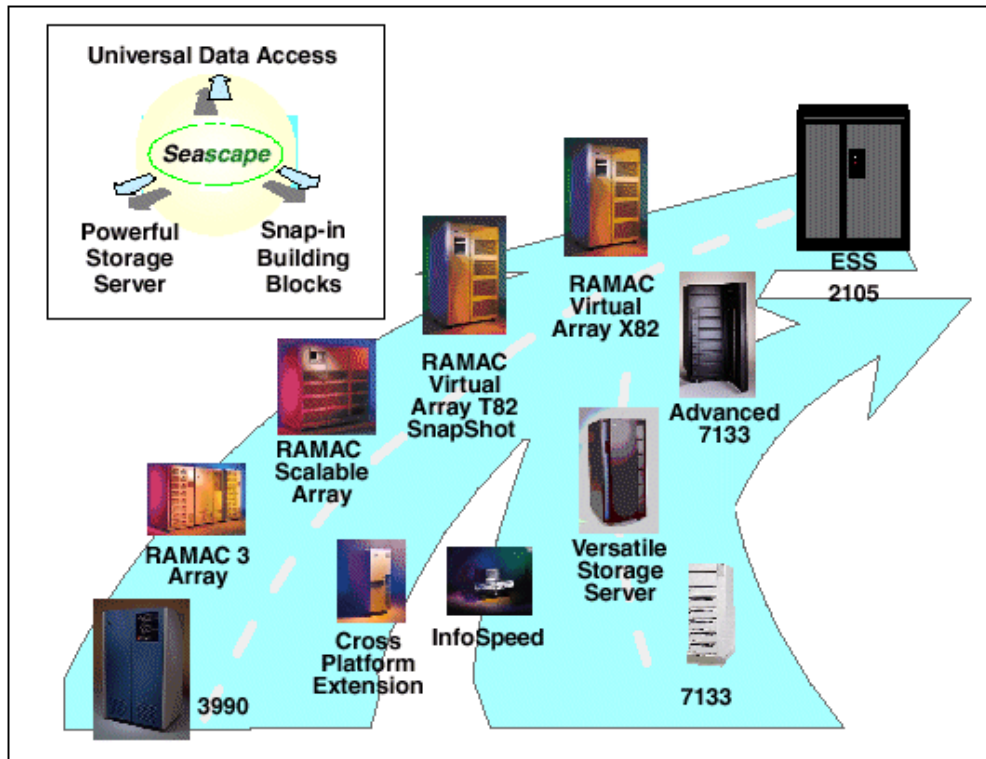
Figure 1-2   Storage evolution

The IBM Enterprise Storage Server (ESS) is a member of the Seascape family of storage products. The ESS was announced in June 1999 and was generally available in September 1999. Since its announcement, it has revolutionized the storage marketplace. It consists of a storage server and attached disk storage devices. The storage server provides integrated caching and RAID support for the attached disk devices. Since the original announcement of the ESS, it has continually evolved, and the latest model, the ESS 800, has several times the performance of the original E model (application dependent). The ESS can be configured in a variety of ways to provide scalability in capacity and performance.

Redundancy within the ESS provides continuous availability. It is packaged in one or more enclosures, each with dual line cords and redundant power. The redundant power system allows the ESS to continue normal operation when one of the line cords is deactivated.

The ESS provides the image of a set of logical disk devices to the attached servers. The logical devices are configured to emulate disk device types that are compatible with the attached servers. The logical devices access a logical volume that is implemented using multiple disk drives.

The following host I/O interface attachments are supported:

► **ESCON® host adapter:** The ESCON host adapter (HA) is the physical component of the ESS used to attach the host ESCON I/O interfaces and ESCON Director ports. The ESCON host adapter connects to an ESCON channel by means of an ESCON link and

accepts the channel command words (CCWs) from the host system. The ESS ESCON host adapters have two ports to connect to two ESCON links.

- ► **FICON™ host adapter:** The FICON host adapter (HA) is the physical component of the ESS used to attach the host FICON I/O interfaces and FICON Director ports. The FICON host adapter connects to a FICON channel by means of the FICON link and accepts the CCWs (channel command words) from the host system. The ESS FICON host adapter is in fact a Fibre Channel/FICON adapter card that can be configured either for FICON or for FCP use. Both 1 Gb and 2 Gb port cards are available.

- ► **SCSI host adapter:** The SCSI host adapter card is used to interface with hosts that are using the SCSI protocol. There are different versions of SCSI, some of which can be supported by the same adapter. The protocols that are used on the SCSI HA (the command set) can be either SCSI-2 or SCSI-3. Each SCSI adapter has two SCSI ports.

- ► **Fibre Channel host adapter:** The Fibre Channel host adapter (HA) is the physical component of the ESS used to attach to the servers' Fibre Channel I/O interfaces and SAN fabric ports. The ESS Fibre Channel host adapter connects to the server Fibre Channel I/O adapter by means of the Fibre Channel link and accepts the upper-layer commands (more than one protocol is supported by the Fibre Channel standard) from the host system. The ESS Fibre Channel host adapter is in fact a Fibre Channel/ FICON adapter card, can be configured either for FICON or for FCP use. It has one 2 Gb port for fiber connection.

1.4 Major components

The ESS can be broken down into several components. The storage server itself is composed of two clusters that provide the facilities with advanced functions to control and manage data transfer. Should one cluster fail, the remaining cluster can take over the functions of the failing cluster. The overall structure is shown in Figure 1-3 on page 6.
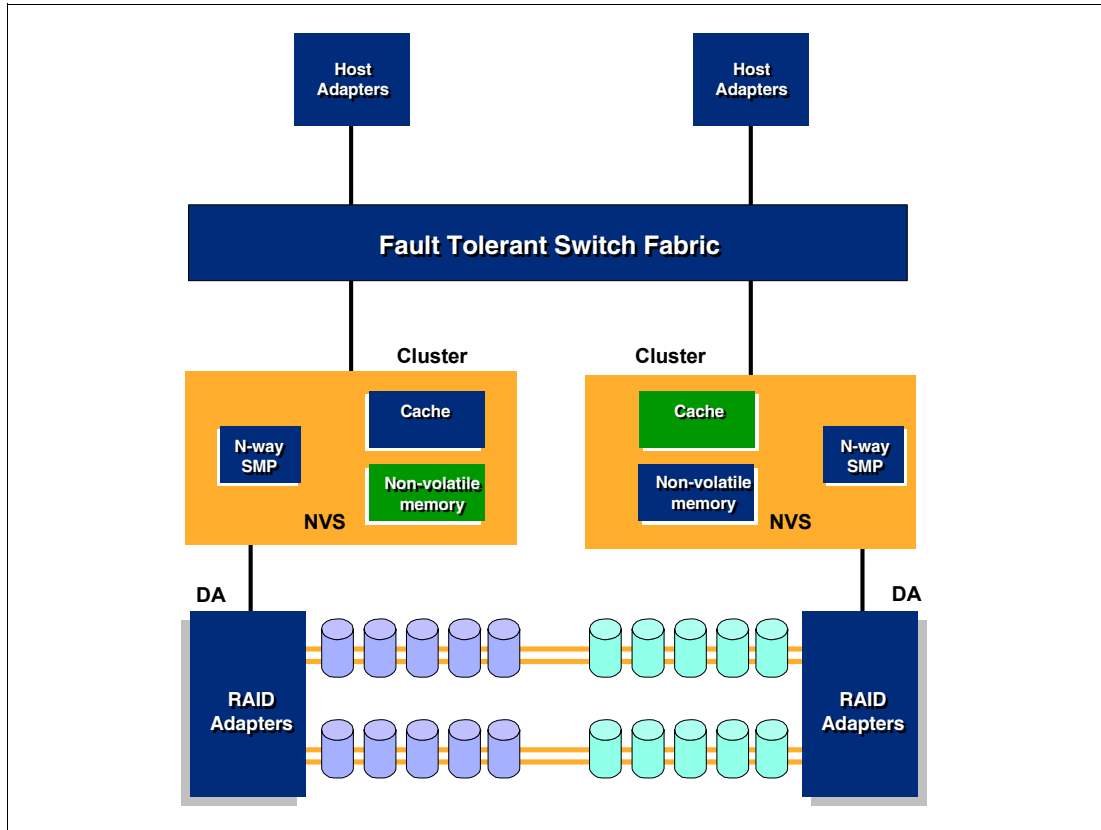
Figure 1-3   ESS high level design overview

A cluster is made up of the following subcomponents:

► **Host adapters**: Each cluster has one or more host adapters (HAs). Each host adapter provides one or more host I/O interfaces. A host adapter can communicate with either cluster complex.

► **Device adapters**: Each cluster has one or more device adapters (DAs). Each device adapter provides one or more storage device interfaces. Disk drives are attached to a pair of device adapters, one in each cluster, so that the drives are accessible from either cluster. At any given time, a disk drive is managed by only one device adapter.

► **Cluster complex:** The cluster complex provides the management functions for the ESS. It consists of cluster processors, cluster memory, cache, nonvolatile storage (NVS), and related logic:

   – **Cluster processors:** The cluster complex contains four cluster processors (CP) configured as symmetrical multiprocessors (SMP). The cluster processors execute the licensed internal code (LIC) that controls the operation of the cluster. The ESS Model 800 offers the higher speed processors and a significant performance enchancement.

   – **Cluster memory / cache:** These are used to store instructions and data for the cluster processors. The cache memory is used to store cached data from the disk drives. The cache memory is accessible by the local cluster complex, by device adapters in the local cluster, and by host adapters in either cluster.

   – **Non-volatile storage (NVS):** This is used to store a nonvolatile copy of active written data. Each copy of the active write data is copied to the NVS area on the other cluster complex to obtain the redundancy. The NVS is accessible to either cluster-processor complex and to host adapters in either cluster. Data may also be transferred between the NVS and cache, if the algorithm detects the required data is on NVS.

The disk drives provide the primary nonvolatile storage medium for any host data stored within the ESS Storage devices. They are grouped into ranks and are managed by the clusters.

1.5 Third-generation hardware: ESS Model 800

The IBM TotalStorage Enterprise Storage Server Model 800 integrates a new generation of hardware from top to bottom, allowing it to deliver unprecedented levels of performance and throughput. Key features that characterize the performance enhancements of the ESS Model 800 are:

- ► The ESS Model 800 generally is capable of delivering twice the throughput of its predecessor Model F20.
- ► With the optional Turbo Processor (feature 3606), it is capable of providing 2.5 times the throughput of its predecessor Model F20, for increased scalability and response times.
- ► The Turbo II Processor option (feature number 3607), a follow-on replacement to the Turbo Processor (3606), offers up to a 30% improvement for high-hit ratio and high operations per second workloads as compared to the Standard Processor in an ESS Model 800.
- ► The 64 GB cache supports much larger system configurations and increases cache hit ratios, driving down response times.
- ► Double the internal bandwidth provides high sequential throughput for digital media, business intelligence, data warehousing, and life science applications.
- ► Larger NVS (2 GB Non-volatile storage) with twice the bandwidth allows greater scalability for write-intensive applications.
- ► Third-generation hardware provides response time improvements of up to 40% for important database applications.
- ► 2 Gb Fibre Channel/FICON host adapters provide doubled performance sustained and instantaneous throughput for both open systems and zSeries environments.
- ► RAID 10 can provide up to 75% greater throughput for selected database workloads compared to equal physical capacity configured as RAID 5. While most typical workloads will experience excellent response times with RAID 5, some cache-unfriendly applications and some applications with high random write content can benefit from the performance offered by RAID 10.
- ► Starting with LIC level 2.3.0 the ESS Model 800 can be configured with Arrays Across Loops (AAL), feature number 9903. AAL is a new configuration option whereby disk arrays are spread across two loops on the SSA device adapter pair. AAL allows you to take full advantage of the ESS subsystem bandwidth, because a RAID array is built across two SSA loops within a device adapter pair.
- ► 15,000 rpm drives provide up to 80% greater throughput per RAID rank and 40% improved response time as compared to 10,000 rpm drives. This allows driving the workloads to significantly higher access densities, while also experiencing improved response times.

1.5.1 Extensive capacity and scalability

The ESS model 800 offers increased performance and scalability with two new models with 8, 16, 24, 32, or 64 GB cache options. There are various configurations possible and the storage scalability goes up to 55.9 TB of physical capacity. There is concurrent support for all your major server platforms, including S/390, OS/400, Windows NT®, Windows 2000, NetWare, and most varieties of UNIX, including IBM AIX.

The extensive storage management capabilities include:

- ► FlashCopy for fast data duplication
- ► Peer-to-Peer Remote Copy (PPRC) and Peer-to-Peer Remote Copy Extended Distance (PPRC-XD) for your synchronous and asynchronous backup and disaster recovery needs
- ► Extensive StorWatch management capability via the Web

It offers superior performance with options and innovations to meet your changing requirements and high availability to support your e-business and other mission-critical applications.

The ESS continues to deliver on its SAN strategy, as was previewed in the July 27, 1999, announcement of the ESS. The ESS now provides up to sixteen 200 MB/sec. native Fibre up to 10 km distance with long wave and 300 m with short wave. Each single port adapter supports Fibre Channel Protocol (FCP) in a direct point-to-point configuration, point-to-point to a switch (fabric) configuration, or Fibre Channel-Arbitrated Loop (FC-AL) in a private loop configuration. For the complete and most up-to-date list of ESS fabric products which are supported, refer to:

http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

1.6 ESS Copy Services

ESS Copy Services is an optional feature of the IBM TotalStorage Enterprise Storage Server. It brings powerful data copying and mirroring technologies to open systems environments previously available only for mainframe storage.

This book deals with the three main features of the ESS Copy Services for the open systems environment (see Figure 1-4 on page 9):

- ► Peer-to-Peer Remote Copy (PPRC)
- ► Peer-to-Peer Remote Copy Extended Distance (PPRC-XD)
- ► FlashCopy

New ESS Copy Services Version 2 features that are added to existing ESS Copy Services functions are also covered:

- ► Incremental FlashCopy
- ► Multiple Relationship FlashCopy
- ► FlashCopy Consistency Groups
- ► Dataset Level FlashCopy (only for S/390 systems)
- ► Asynchronous Cascading PPRC and Inband FlashCopy

Also, this redbook discusses the following ESS facilities for copy services that became available starting with LIC level 2.3.0:

- ► PPRC over Fibre-Channel.
- ► ESS Application Programming Interface enhancements for ESS Copy Services support.
- ► Asynchronous Cascading PPRC has been enhanced to exploit failover and failback modes in a three-site solution.
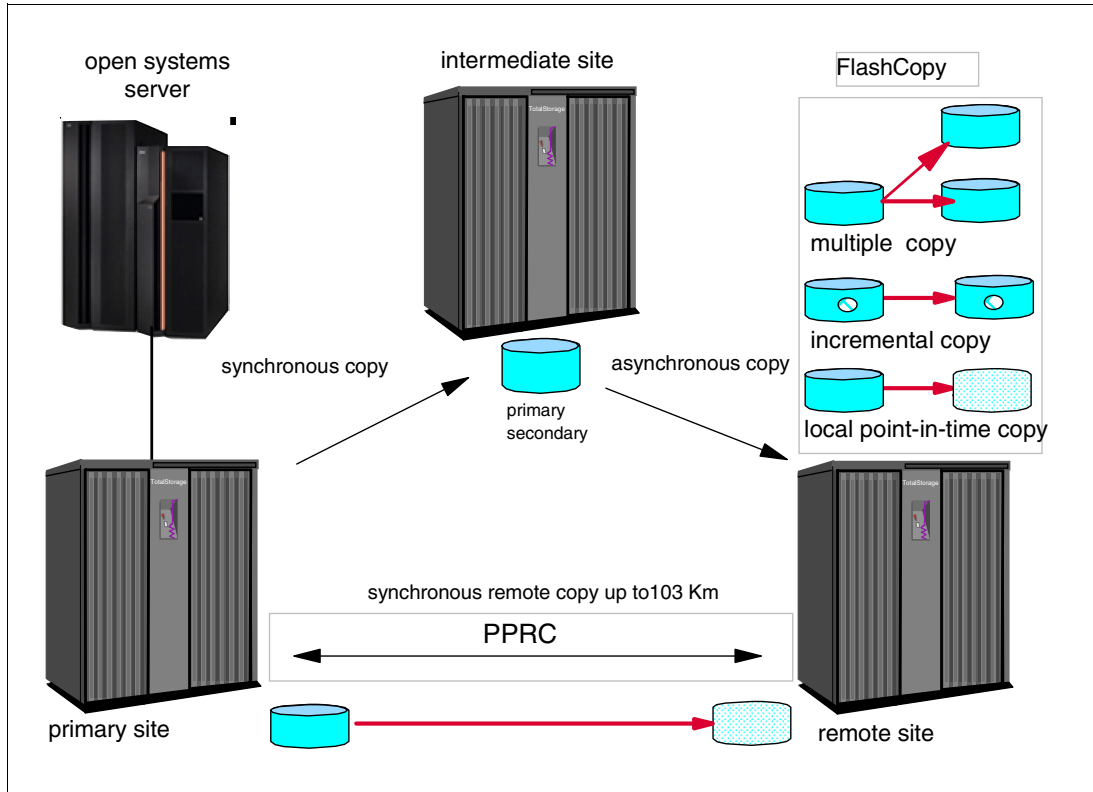
Figure 1-4   ESS Copy Services functions overview

1.6.1 Peer-to-Peer Remote Copy (PPRC)

PPRC is a function of a storage server that constantly updates a secondary copy of a volume to match changes made to a primary volume. The primary and the secondary volumes can be on the same storage server or on separate storage servers. In the case of two ESSs, the secondary ESS can be located at another site some distance away (see Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53).

PPRC is application independent. Because the copying function occurs at the disk subsystem level, the application has no knowledge of its existence.

The PPRC protocol guarantees that the secondary copy is up-to-date and consistent by ensuring that the primary copy will be written only if the primary receives acknowledgment that the secondary copy has been written.

1.6.2 Peer-to-Peer Remote Copy Extended Distance (PPRC-XD)

Peer-to-Peer Remote Copy Extended Distance adds flexibility to the IBM TotalStorage Enterprise Storage Server and PPRC, PPRC-XD is a non-synchronous long-distance copy option for both open systems and zSeries servers.

PPRC-XD can operate at very long distances, even continental distances, well beyond the 103 km (maximum supported distance for synchronous PPRC) with minimal impact on the applications. Distance is limited only by the network and channel extenders technology capabilities. This book presents an overview of the characteristics of PPRC Extended Distance for the open system environment. For additional information and how to use it, refer

to the redbook *IBM TotalStorage Enterprise Storage Server PPRC Extended Distance*, SG24-6568.

1.6.3 FlashCopy

FlashCopy makes a single point-in-time (PIT) copy of a LUN. This is also known as a time-zero (T0) copy. The target copy is available once the FlashCopy command has been processed (see Chapter 3, "FlashCopy" on page 33). FlashCopy provides an instant or point-in-time copy of an ESS logical volume. Point-in-time copy functions give you an instantaneous copy, or "view", of what the original data looked like at a specific point-in-time.

The point-in-time copy created by FlashCopy is typically used where you need a copy of production data to be produced with minimal application downtime. It can be used for online backup, testing of new applications, or for copying a database for data mining purposes. The copy looks exactly like the original source volume and is an instantly available, binary copy.

1.6.4 Incremental FlashCopy (V2)

Incremental FlashCopy is a new feature of FlashCopy that is available with CopyServices Version 2. This function is to be used in conjunction with the background copy option to track the changes on the source volume since the last FlashCopy relationship was invoked. When this option is selected, only the tracks that have been changed on the source are copied again to the target. The direction of the "refresh" can also be reversed, copying the changes made to the new source (originally the target volume) to the new target volume (originally the source volume).

1.6.5 Inband FlashCopy (V2)

The new inband management capability feature that comes with Copy Services Version 2 allows you to invoke FlashCopy on a remote site ESS. If you have two sites, one local and one remote, which are in a PPRC relationship, a FlashCopy task on the remote site ESS can be invoked from the primary site ESS via a PPRC inband connection.

1.6.6 Multiple Relationship FlashCopy (V2)

With Copy Services Version 2, one FlashCopy source may have up to 12 FlashCopy targets. This gives you more flexibility as you can initiate the multiple relationships using the same source volume without needing to wait for other relationships to end.

1.6.7 FlashCopy Consistency Groups

New options are available to facilitate the creation of FlashCopy Consistency Groups. With the FlashCopy Consistency Groups, the ESS will hold off I/O activity to a volume until the **Consistency Created** task with the **FlashCopy Consistency Group** option is issued.

1.6.8 Data Set Level FlashCopy (V2 for OS/390 only)

With this feature, data sets in the OS/390 environment can be copied independently from the volume level. This feature also provides the possibility for a volume to act as a source volume and a target volume at the same time.

> **Note:** To use this function, the OS/390 version should be V2R10 or higher.

1.6.9 Asynchronous Cascading PPRC (V2)

Another new feature that comes with Copy Services Version 2 is Asynchronous Cascading PPRC. The concept of the feature is that a PPRC secondary volume serves as a PPRC primary volume for a second PPRC relationship. The intermediate site target PPRC volume acts as the source volume for the remote site PPRC relationship. This feature gives you the flexibility to establish asynchronous mirroring between the intermediate and remote site as well as synchronous mirroring between the primary and intermediate site.

With LIC level 2.3.0 Asynchronous Cascading, PPRC has been enhanced to exploit the PPRC Failover and Failback modes. This can help reduce the time required to synchronize PPRC volumes after switching between sites during planned or unplanned outages.

1.6.10 PPRC over Fibre Channel

PPRC over Fibre Channel enables the use of Fibre Channel Protocol (FCP) as a communications link between the ESS PPRC primary and secondary machines.

1.6.11 ESS Application Programming Interface

ESS Application Programming Interface (ESS API) has been enhanced to support ESS Copy Services configuration and use.

1.6.12 Managing ESS Copy Services

ESS Copy Services provides a Command Line Interface (CLI) as well as a Web-based interface for setting up and managing its facilities (see Chapter 7, "ESS Command Line Interface" on page 301). The CLI interface allows administrators to execute Java™-based Copy Services commands from a command line. The Web-based interface, a part of the TotalStorage ESS Specialist, allows storage administrators to manage Copy Services from a browser-equipped computer (see Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103).

The ESS Copy Services functions can be used separately as stated above or they can be combined to produce solutions such as split mirror for remote tape vaulting backup solutions or static point-in-time volume replication. The allowable Copy Services combination can be seen in Table 1-1.

Table 1-1   Allowable combinations of copy operations on the same volume

If device is →	FlashCopy source	FlashCopy target	PPRC primary	PPRC secondary
May become				
PPRC primary	OK	OK	NO	OK (1)
PPRC secondary	OK	OK (2)	OK (1)	NO
FlashCopy source	OK (3)	NO	OK	OK
FlashCopy target	NO	NO	NO	NO
Notes: 1 - Asynchronous Cascading PPRC (V2) option: PPRC target volume on intermediate site can be used as PPRC source volume for the remote site PPRC target. 2 - The operation is allowed. Updates of the affected extents will result in implicit removal of the FlashCopy relationship. 3 - Multiple Relationship FlashCopy (V2).				

1.6.13 ESS Copy Services benefits

Copy Services will be of great use to customers with large IT systems, big data volumes, and a requirement for round-the-clock data availability.

Although each has its specific features, PPRC and FlashCopy are typically used as data backup tools for creation of test data and for data migration. They can also be used in disaster recovery scenarios (see Chapter 10, "Disaster recovery" on page 399).

Copy Services will provide the greatest benefit to the customer who:

► Needs to have disaster tolerant IT centers
► Is planning to migrate data between systems
► Is migrating workloads often
► Has to back up large amounts of data
► Needs to reduce the time the server has to be taken offline for backup
► Plans to test new applications
► Needs a copy of production data for data warehousing or data mining

Copy Services can be integrated with technologies such as Tivoli® Storage Manager (formerly ADSM), Logical Volume Manager (LVM) mirroring, or SAN Data Gateway mirroring to solve a wide variety of business issues. Other companies will most likely be selling partial solutions as a means to solve these problems. IBM, however, with its broad portfolio of products in this industry, has many experts available to discuss the right solution for your business and to help you design and implement a solution that will give you the maximum business benefit.

Advanced solutions with ESS Copy Services have been endorsed by many Independent Software Vendors (ISVs) worldwide.

2

Implementing ESS Copy Services

In this chapter, we introduce the terminology used when working with the ESS Copy Services Domains.

This chapter also describes the architecture of the Copy Services domains. It discusses some planning issues when preparing an installation of an ESS Copy services domain.

2.1 Copy Services terminology

In this section, we explain the terminology used when working within an ESS Copy Services environment.

2.1.1 Copy Services Version 1 (V1)

An ESS is running the Copy Services software Version 1 when the LIC level installed on the ESS is prior to 2.2.0.

2.1.2 Copy Services Version 2 (V2)

An ESS is running the Copy Services software Version 2 when the LIC level installed on the ESS is 2.2.0. or higher.

2.1.3 FlashCopy Version 1 (V1) and FlashCopy Version 2 (V2)

FlashCopy is a Copy Services feature. The FlashCopy feature must be installed and activated on the ESS before using the FlashCopy functions. The FlashCopy feature is available in two different versions (V1 and V2). Refer to Chapter 3, "FlashCopy" on page 33 for more details about the two versions.

- ► An ESS running Copy Services V1 (LIC level prior to 2.2.0) supports only FlashCopy V1.
- ► An ESS running Copy Services V2 (LIC level 2.2.0 or higher) supports FlashCopy V1 and FlashCopy V2.

2.1.4 PPRC Version 1 (V1) and PPRC Version 2 (V2)

PPRC is a Copy Services Feature. The PPRC feature must be installed and activated on the ESS before using the PPRC functions. The PPRC feature is available in two different versions (V1 and V2). Refer to Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53 for more details about the two versions.

- ► ESS running Copy Services V1 (LIC level prior to 2.2.0) supports only PPRC V1.
- ► ESS running Copy Services V2 (LIC level 2.2.0 or higher) supports PPRC V1 and PPRC V2.

2.1.5 Copy Services Domain

All ESS subsystem clusters participating in Copy Services relationships are grouped together in a *Copy Services Domain*. The Copy Services domain is managed by one or two ESS clusters designated as *active Copy Services servers*.

A Copy Services Domain is also called a Copy Services server group.

2.1.6 Active Copy Services Server

The Active Copy Services Server is the server that manages the Copy Services Domain. The active Copy Services Server is running the Copy Services server software and communicates with the Web browsers running the Copy Services WUI.

In Copy Services V1 (LIC prior to 2.2.0), this would normally be the Primary Copy Services server, unless it fails, in which case you would manually invoke the Backup Copy Services server as the active Copy Services server.

In Copy Services V2 (LIC 2.2.0 and higher), you can define two active Copy Services servers (ServerA and ServerB). Depending on your configuration, you can have the two Copy Services servers defined active at the same time (see 2.1.12, "Dual-active server configuration (Copy Services V2)" on page 16 and 2.1.13, "Mixed mode configuration (Copy Services V2 and Copy Services V1)" on page 16 for more information).

The active Copy Services server performs the following functions:

► It collects connectivity data and configuration data from all clients that are part of the Copy Services server group.

► It saves and manages Copy Services tasks.

► It sends tasks to the appropriate client.

► If you have a LIC level below 2.2.0., when the Primary server is the active server, it copies tasks to the Backup Copy Services server automatically.

► If you have LIC level 2.2.0 or higher, it copies tasks to the other Copy Services server automatically.

► It collects volumes status updates from clients.

► It communicates with the Web browser.

2.1.7 Copy Services client

A Copy Services client is software that runs on each cluster in your Copy Services server domain and performs the following functions:

► It communicates configuration, status, and connectivity information to the Copy Services server.

► It executes data-copy tasks on behalf of the Copy Services server.

2.1.8 Primary Copy Services Server (Copy Services V1)

In Copy Services V1 (CLI prior to 2.2.0), the Primary Copy Services server is a user-designated ESS cluster in your Copy Services server group that performs the role of the active Copy Services server until it fails. If the active Copy Services server running on the primary Copy Services server fails, an operator must manually restart the Backup Copy Services server to run as the active server.

2.1.9 Backup Copy Services Server (Copy Services V1)

The Backup Copy Services server is an ESS cluster in your Copy Services server group that can become active when the Primary Copy Services server fails. The Backup Copy Services server has the following characteristics:

► The server remains in a non-active mode, unless an operator manually restarts it in the event of a failure.

► When the server is non-active, it receives updates to the task repository from the other server.

► It runs on a different ESS cluster than the Primary Copy Services server.

► It saves tasks while running as the active Copy Services server. However, saved tasks are lost when the ESS cluster is reset to the Primary Copy Services server.

2.1.10 Copy Services server operating mode

A Copy Services server can operate in one of two modes: dual-active or single-active. Dual-active mode is the mode of operation if the IBM TotalStorage Enterprise Storage Server (ESS) license internal code (LIC) level is 2.2.0 or higher; otherwise, single-active is the mode of operation.

2.1.11 Single-active configuration (Copy Services V1)

The servers are running pre-2.2 LIC and operating in single-active mode. One server is the primary and the other is the backup. During normal operation in single-active mode, only one of the two Copy Services servers is active. The client-only ESS can be running pre-2.2 LIC or 2.2 LIC or higher. If a disaster occurs at your production site and the primary server is at your production site, user intervention is required. You must make the Backup server your active server at the recovery site for all ESSs that are running pre-2.2 LIC.

2.1.12 Dual-active server configuration (Copy Services V2)

You are in a dual-active configuration when all the ESSs in a Copy Services domain are using Copy Services V2 (2.2 LIC level or higher). Two ESS clusters can be designated to perform the role of the active Copy Services servers. In the dual-active mode, both Copy Services servers are active at the same time. There is no concept of a primary or Backup Copy Services server. ServerA is the designation for one of the Copy Services servers, and serverB is the designation for the other server. All the ESSs have clients registered with both servers. In the event of one active server failure, there is no manual action required, as the other server is still active.

2.1.13 Mixed mode configuration (Copy Services V2 and Copy Services V1)

You will be in this mixed-mode when the ESSs in your domains have different Copy Services versions. The following configurations are possible:

► One server is running Copy Services V2 (2.2 LIC or higher) and operating in dual-active mode.The other server is running Copy Services V1 (pre-2.2 LIC) and operating in single-active mode. Client-only ESSs can be running V1 or V2 (pre-2.2 LIC or 2.2 LIC and higher). With this configuration, it is better if the dual-active server is the active server (serverA) and the single-active server is the passive server (backup). If a disaster occurs at your production site and the server at the production site is the dual-active server, user intervention is required. You must make the server at the recovery site the active server for all the ESSs that are running Copy Services V1 (pre-2.2 LIC).

► Two servers (serverA and serverB) are running 2.2 LIC. or higher. At least one client-only ESS is running pre-2.2 LIC. If a disaster occurs at your production site and the server at the production site is the primary server for the ESS that is running pre-2.2 LIC, user intervention is required. You must make the server at the recovery site the active server for all client-only ESSs running pre-2.2 LIC.

2.1.14 The ESS Command Line Interface (CLI)

The ESS provides two command line interfaces: the ESS Management CLI and the ESS Copy Services CLI.

ESS Management CLI

The ESS Management CLI (Command Line User Interface) is software that runs on the host systems. The ESS Management CLI provides a set of commands that you can use to monitor

and to manage the ESS functions. Some commands provided by the ESS management CLI are very useful to view the configuration of the Copy Services environment. You can include the CLI commands into scripts to automate some of your procedures. For more details about the useful ESS Management CLI for Copy Services, refer to Chapter 7, "ESS Command Line Interface" on page 301.

Copy Services CLI

The CLI (Command Line User Interface) is software that runs on host systems. The Copy services CLI provides a set of commands you can use to invoke the Copy Services functions and monitor the status of the Copy Services tasks and volumes. You can include the CLI command into scripts to automate some of your procedures. For more details about the CLI, refer to Chapter 7, "ESS Command Line Interface" on page 301.

2.2 Network considerations for Copy Services

There are some Network considerations when you build the Copy Services domain.

▶ *The ESS clusters' communications*: This communication is necessary between the active servers and all the clients in the Copy Services domain. This communication is via the Local Area Network.

▶ *Copy Services management*: The normal way to manage the Copy Services is to use the Web graphical user interface. Also, you can use the CLI to invoke copy services functions. Both means (Web GUI and CLI) are using network connections to the Copy Services domain. Note that some configuration functions cannot be performed via the CLI, for example, creating a task is not possible via the CLI.

2.2.1 ESS cluster communication and network configuration

The active Copy Services servers and the Copy Services clients communicate together via a TCP/IP Ethernet connection. This communication is provided by the ESSNet.

The ESSNet is the ESS access facility that IBM installs with your ESSs. The ESSNet consists of a switch (hub) and other networking components and the dedicated IBM TotalStorage Enterprise Storage Server Master Console (*ESS Master Console*).

In a multiple site configuration, you need a physical network connection between the sites to connect the ESS clusters on the same LAN.

The Copy Services V2 (CLI 2.2.0 or higher) provides the option to use inband commands over the PPRC link to invoke FlashCopy tasks. You may use this specific option to invoke tasks at a remote site without maintaining the LAN connection. See Chapter 3, "FlashCopy" on page 33 for more information about how inband management works.

2.2.2 Management via the WUI and network considerations

To manage the ESS Copy Services, a Web server running in your ESS provides a Web interface that you can use to manage the ESS through a Web browser. The Web browser must be running on a workstation connected to the ESS through the IBM TotalStorage Enterprise Storage Server Network (ESSNet).

In addition to using a Web browser on the ESSNet Master Console to connect to your ESS, you can also use a Web browser running on your own workstation, either by connecting your workstation directly to the ESSNet hub, or by connecting your workstation to your intranet and

connecting your intranet to the ESSNet hub. If you use your own workstation, IBM recommends that it has at least 128 MB of memory.

2.2.3 Management using the CLI and network considerations

In addition to the WUI, you can use the Java based Command Line User interface (CLI) to manage your ESS and to invoke Copy Services functions. The CLI is software that is installed and runs on a host system (see Chapter 7, "ESS Command Line Interface" on page 301).

If you want to use the CLI commands to manage the ESSs in the domain, the requirement is that the host from which you want to invoke the commands can communicate via the Local Area Network to all the ESS clusters in your Domain.

If you want to use the CLI commands to invoke the Copy Services functions, the only requirement is that the host from which you want to invoke the commands can communicate via the Local Area Network to the active Copy Services ESS cluster in your Domain.

2.3 ESS Copy Services V1

In this section, we discuss the implementation of a domain when all the ESSs are running Copy Services V1 (LIC prior to 2.2.0).

2.3.1 Copy Services V1: Overview and requirements

Before you can use the PPRC or FlashCopy functions, you must have the appropriate feature codes installed on the ESS:

► The FlashCopy license feature ordered must be equal to or greater than the total capacity of the ESS. The FlashCopy License feature must be installed on all ESSs on which you want to establish FlashCopy pairs. The CopyServices V1 supports only the FlashCopy V1 Feature.

► The PPRC license feature must be equal to or greater than the total capacity of the ESS. The PPRC feature must also be purchased and installed on both the primary and secondary ESSs. The CopyServices V1 supports only the PPRC V1 Feature.

You should consider the following requirements when doing your planning:

► One ESS cluster in the domain must be defined as the Primary Server for Copy Services. Defining another cluster as a Backup server for the Copy Services Domain is optional.

► Starting with LIC level 2.3.0., the maximum number of ESSs that can be in the same Copy Services Domain is eight.

► Only one FlashCopy at a time can be active on a volume, however, you can perform a PPRC concurrently with FlashCopy on the same volume. The designated FlashCopy target volume cannot be a primary volume in a PPRC volume pair.

► The primary and the secondary volume of a PPRC pair can only be in one relationship at a time.

► The primary and secondary volumes for a PPRC pair must be an ESS type. You cannot establish a PPRC pair from an ESS to a different external device type.

► The source and target volume for a FlashCopy must reside in the same LSS of an ESS.

► You need to manage PPRC using the ESS Copy Services Web interface, therefore, Ethernet and TCP/IP connectivity is needed between all of the participating ESS subsystems, and the Web browser initiating and managing the PPRC activities.

- ► You can use the Command Line User Interface to invoke the Copy Services Functions from the command line. Ethernet and TCP/IP connectivity is needed between the server hosting the CLI software and the Copy Services servers in the Copy Services Domain.

- ► The source and target logical volumes must be the same size, or the target must be larger in size to establish a FlashCopy relationship.

- ► The primary and secondary logical volumes must be the same size, or the secondary must be larger in size to establish a PPRC relationship.

- ► There are some limitations when establishing a path from one LSS to another LSS. See Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53 for path limitations.

- ► A Copy Services Domain can handle a maximum of 2048 copy services pairs. This number includes all the primary and secondary PPRC pairs plus all the source and target FlashCopy pairs.

2.3.2 ESS Copy Services V1: Defining the Copy Services Domain

A Copy Services Domain is a group of ESS clusters participating in Copy Services relationships.

Starting with LIC level 2.3.0, up to eight ESSs can reside in the same Copy Services Domain. You will choose, among the clusters in the domain, a pair of clusters and define them as servers for the Copy Services Domain:

- ► One Cluster to be defined as the Primary Copy Services server (mandatory).

- ► One Cluster to be defined as the Backup Copy Services server (optional but recommended).

All the other clusters will be clients for the Copy Services.

Domain definitions

It is mandatory to define the Primary Copy Services server in all the ESS clusters participating in the Copy Services Domain. This definition is performed using the ESS master console.

Optionally, a Backup server for the Copy Services can be defined on all the ESS clusters in the Copy Services domain. This definition is performed using the ESS master console.

The definition of the Backup Copy Services server is not mandatory, but we strongly suggest that a Backup server be defined in your Copy Services Domain. The Primary and Backup Copy Services servers can reside on the same ESS; however, we suggest defining the Primary and Backup server in two different servers.

During the configuration, you must also provide a list of all the ESS Clients in the Copy Services Domain to the Primary Copy Services server. This list is known as the "Copy Services clients" list. You must provide the ESS serial numbers, the ESS clusters' host names, and IP addresses of all the ESSs in your Copy Services Domain (including the Primary and Backup server). These definitions are performed using the ESS master console.

2.3.3 Planning for a disaster with ESS Copy Services V1

A key decision you must make in planning for a disaster is deciding where to place your Primary and Backup Copy Services servers in a two-site environment. There are two options:

- ► The typical setup is that you place your Primary Copy Services server at your production site (the place where most of your applications are running).

► Another option is to place your Primary Copy Services server at your recovery site.

There are trade-offs associated with each option:

► Benefits of placing your Primary Copy Services server at your production site:

If a disaster or emergency situation at your production site causes you to lose connections to the Primary Copy Services server, you can switch control to your recovery site. The ESS Copy Services code copies tasks synchronously between the Primary and Backup servers. This configuration might be the desirable setup in open systems environments. The reason for this is that the ESS Copy Services CLI uses TCP/IP links to communicate between the host system and the Copy Services server, which may introduce network delays. However, with this configuration, if your primary site fails, there can be delays in bringing up applications at the recovery site.

► Benefits of placing your Primary Copy Services server at your recovery site:

Choosing this configuration can gain you some efficiency in terms of recovering from a disaster. Because the Primary server is at the recovery site, you do not have to perform manual recovery steps to switch to a different Copy Services server. If you have planned in advance and have created tasks for disaster recovery, then you can run those tasks and bring your production systems back up.

See the *IBM TotalStorage Enterprise Storage Server: Introduction and Planning Guide,* GC26-7294 for more information about factors you should consider in making the decision about where to place your ESS Copy Services servers.

2.4 ESS Copy Services V2

In this section, we discuss the implementation of a domain when the ESSs are running Copy Services V2 (LIC 2.2.0 or higher). We also describe the implementation of a domain with a mixed environment of Copy Services (V1 and V2).

2.4.1 Copy Services V2: Overview and requirements

Before you can use the PPRC or FlashCopy functions, you must have the appropriate feature codes installed on the ESS:

► The FlashCopy license feature ordered must be equal to or greater than the total capacity of the ESS. The Copy Services V2 supports FlashCopy V1 and FlashCopy V2 feature codes.

► The PPRC license feature must be equal to or greater than the total capacity of the ESS. The PPRC feature must also be purchased and installed on both the primary and secondary ESSs. The Copy Services V2 supports PPRC V1 and PPRC V2 feature codes.

► To be able to benefit from Copy Services Version 2 functions, the LIC level of the ESS must be 2.2.0 or higher.

You should consider the following requirements when planning:

► You can establish multiple FlashCopy relationships at one time using the same source volume. The maximum number of FlashCopy relationships using the same source volume is 12. Among all the FlashCopy relationships using the same source volume, only one can be an incremental FlashCopy.

► A secondary (target) volume of a PPRC relationship can become the primary volume of another PPRC relationship. This possibility is provided by the Asynchronous Cascading PPRC functionality.

- You can perform a PPRC concurrently with FlashCopy on the same volume. The designated FlashCopy target volume cannot be a primary volume in a PPRC volume pair.

- The primary and secondary volumes for PPRC must be an ESS type. You cannot establish a PPRC from an ESS to a different external device type.

- The source and target volume for FlashCopy must reside in the same ESS. The source and target volumes can be in different LSSs.

- You need to manage PPRC using the ESS Copy Services Web interface, therefore, Ethernet and TCP/IP connectivity is needed among all of the participating ESS subsystems and the Web browser initiating and managing the PPRC activities.

- You can use the Command Line User Interface to invoke the Copy Services functions from a command line. Ethernet and TCP/IP connectivity is needed between the server hosting the CLI software and the Copy Services servers in the Copy Services Domain.

- The source and target logical volumes must be the same size, or the target must be larger in size to establish a FlashCopy relationship.

- The primary and secondary logical volumes must be the same size, or the secondary must be larger in size to establish a PPRC relationship.

- There are some limitations when establishing a path from one LSS to another LSS. See Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53 for path limitations.

- When using the inband commands to invoke FlashCopy tasks, the source volume you select for the FlashCopy at the remote site must be the secondary PPRC volume of the PPRC pair. (The PPRC secondary volume becomes the source volume in a FlashCopy pair.)

- A Copy Services Domain can handle a maximum of 2048 copy services pairs. This number includes all the primary and secondary PPRC pairs plus all the source and target FlashCopy pairs.

2.4.2 ESS Copy Services V2: Defining the Copy Services Domain

A Copy Services Domain is a group of ESS clusters participating in Copy Services relationships.

Starting with LIC level 2.3.0 up to eight ESSs can reside in the same Copy Services Domain. You will choose, among the clusters in the domain, a pair of clusters and define them as servers for the Copy Services Domain:

- One cluster to run as active ServerA (mandatory).

- One cluster to run as active ServerB (optional, but recommended).

All of the other clusters will be clients for the Copy Services.

Domain definitions
It is mandatory to define the ServerA in all of the ESS clusters running Copy Services V2 and participating in the Copy Services Domain. These definitions can be performed using the WUI or using the ESS master console.

Optionally, in all of the ESS clusters running Copy Services V2 and participating in the Copy Services Domain, you can also define a ServerB as the other active Copy Services server. This definition is performed using the WUI or using the ESS master console.

During the configuration, you must also provide to the ServerA and the ServerB a list of all the ESS clients in the Copy Services Domains. This list is known as the Copy Services clients list. You must provide the ESS serial numbers, the ESS clusters host names, and IP

addresses of all the ESSs in your Copy Services Domain (including the ServerA and ServerB clusters). The definition is performed by using the WUI or using the ESS master console.

The procedure is detailed in the Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169.

2.5 Copy Services Domain configurations

In this section we discuss the following operating modes:

- ► The dual-active configuration: All the ESSs in the domain are running Copy Services V2.
- ► The mixed configuration: The ESSs in your domain have different versions of Copy Services. Two configurations are possible in the mixed mode.

2.5.1 Dual-active configuration

In this configuration, all the ESSs in the Copy Services Domain are running V2 and two ESS clusters are defined as active servers (ServerA and ServerB). Both servers will be active at the same time. In the event of one active server failure, there is no manual action required, as the other server is still active.

Planning for disaster recovery in dual-active mode

A key decision you must make in planning for a disaster is deciding where to place your Server A and Server B in a two-site environment.

For maximum protection against a disaster, ensure that one dual-active server is the IP address of an ESS cluster at your production site and the other dual-active server is the IP address of an ESS cluster at your recovery site.

It is preferable that ServerB resides at the production site and ServerA at the recovery site.

Attention: When the two servers resume communication after a disaster recovery, the task repository in ServerB is overwritten by the one in ServerA, when the ServerA becomes active.

Configuration example

We give here an example of a two-site configuration with 3 ESSs. All the ESSs are running Copy Services V2.
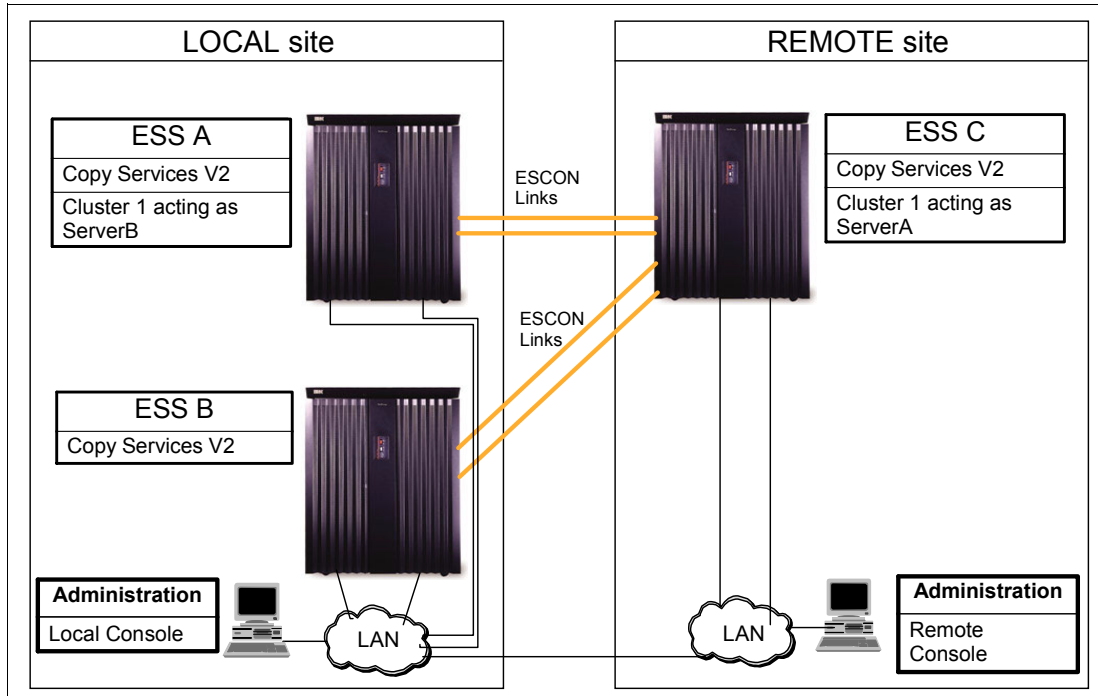
Figure 2-1   Example of a domain

Figure 2-3 on page 28 and the Table 2-3 on page 29 provide details about this dual-active configuration. Here are some specific points about this example:

► All the ESSs are running Copy Services V2 and two servers are defined as active servers (ServerA and ServerB). ServerA is at the remote site, and serverB is at the local site.

► The ServerA and ServerB IP address are defined on all the clusters with Copy Services V2. The clients list is defined on both ServerA and ServerB. A Domain Wide Reset procedure was performed from ServerA to initiate the communication between the V2 servers and the V2 clients.

► Task definitions can be performed either by starting the Copy Services WUI with ServerA or with ServerB.

► In case of a failure of ServerA or ServerB, no manual action is required, since the other server can take over the Copy Services management.

Table 2-1   Example of configuration: dual-active servers and V2 clients

	ESS A	ESS B	ESS C
Site	LOCAL	LOCAL	REMOTE
Copy Services	V2 (LIC 2.2 or higher)	V2 (LIC 2.2 or higher)	V2 (LIC 2.2 or higher)
Cluster 1 IP Address	10.10.1.21	10.10.1.31	10.10.1.41
Cluster 1 is...	Active server (V2) (ServerB)	Client (V2)	Active server (V2) (ServerA)
Cluster 1 Copy Services definitions	ServerA=10.10.1.41 ServerB=10.10.1.21 and Clients list	ServerA=10.10.1.41 ServerB=10.10.1.21	ServerA=10.10.1.41 ServerB=10.10.1.21 and Clients list and "Domain Wide Reset" procedure

	ESS A	ESS B	ESS C
Cluster 2 IP Address	10.10.1.22	10.10.1.32	10.10.1.42
Cluster 2 is...	Client (V2)	Client (V2)	Client (V2)
Cluster 2 Copy Services definitions	ServerA=10.10.1.41 ServerB=10.10.1.21	ServerA=10.10.1.41 ServerB=10.10.1.21	ServerA=10.10.1.41 ServerB=10.10.1.21

2.5.2 Mixed configuration

In this configuration, you have a mix of ESSs running V1 and V2 Copy Services software. There are two possible configurations:

► **Configuration A**: One server is running Copy Services V2 (2.2 LIC or higher) and operating in dual-active mode.The other server is running Copy Services V1 (pre-2.2 LIC) and operating in single-active mode. Client-only ESSs can be running V1 or V2 (pre-2.2 LIC or 2.2 LIC) and higher. With this configuration, it is better if the dual-active server is the active server (serverA) and the single-active server is the passive server (backup). If a disaster occurs at your production site and the server at the production site is the dual-active server, user intervention is required. You must make the server at the recovery site the active server for all the ESSs that are running Copy Services V1 (pre-2.2 LIC).

► **Configuration B**: Two servers (ServerA and ServerB) are running 2.2 LIC. or higher. At least one client-only ESS is running pre-2.2 LIC. If a disaster occurs at your production site and the server at the production site is the primary server for the ESS that is running pre-2.2 LIC, user intervention is required. You must make the server at the recovery site the active server for all client-only ESSs running pre-2.2 LIC.

> **Attention:** In a mixed configuration, any ESS cluster that is running pre-2.2.0 LIC must have its primary Copy Services server defined as the IP address for ServerA.

Important considerations when defining the domains in configuration A

With your mixed configuration A, you cannot have, in the same domain and at the same moment, one active Copy Services server running V2 software and one active Copy Services server running V1 software. The single-active server must be in non-active (passive) mode. The single-active server will be manually switched to the active mode only if the dual-active server is not active or not available (this can happen, for example, in a disaster recovery scenario).

Perform the following steps to put the Backup server in a passive (non-active) state:

1. Establish a browser connection to the ESS Launch panel at the single-active server.

2. Click the **Tools** button.

3. Perform a reset to primary procedure by clicking on the **Reset to Primary** button.

4. You also need to establish a browser connection to all the other ESS clusters running Copy services V1 (and acting as clients) and perform steps 2 and 3 to initiate the communication with the active server.

This procedure is detailed in Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103.

Planning for disaster recovery in mixed configuration A

For maximum protection against a disaster, one server must be the IP address of an ESS cluster at your production site and the other server should be the IP address of an ESS cluster at your recovery site.

Two configurations are supported:

► The dual-active server is at your production site and the single-active server is at your recovery site.

► The dual-active server is at your recovery site and the single-active server is at your production site.

Configuring a mixed configuration A

► On all ESSs running Copy Services V2 (LIC 2.2.0 or above):

– The ESS cluster defined as ServerA is a cluster running Copy Services V2.

– The ESS cluster defined as ServerB is a cluster running Copy Services V1.

► On the ESS running Copy Services V1 (LIC prior 2.2.0)

– The ESS cluster defined as the Primary server is a cluster running Copy Services V2 (this cluster must be the ServerA cluster defined for ESSs running Copy Services V2).

– The ESS cluster defined as Backup server is a cluster running Copy Services V1.

In this configuration, one Copy Services server (the ServerA) operates in Dual-active mode, while the other Copy Services server (the Backup server) operates in Single-active mode.

Be aware of the following:

► The Copy services clients with Copy Services V1 still must be configured by defining a "Primary server" and a "Backup server".

► You should define the Primary server with the IP address of the ServerA, and the Backup server with the IP address of the Backup server cluster.

► All the ESS clusters in the domain with LIC prior to 2.2.0 will run the Copy Services V1 Client software. Those clients can only communicate with one active server. You will have to initiate manually the communication to the active server using the **Reset to primary** procedure. Those actions are described in Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103.

Example of mixed configuration A

We give here an example of a two site configuration with 3 ESSs. Two of the ESSs are running Copy Services V2 and one is running Copy Services V1.

LOCAL site   REMOTE site

ESS A

Copy Services V2

Cluster 1 acting as
ServerA

ESS B

Copy Services V2

ESCON
Links

ESCON
Links

ESS C

Copy Services V1

Cluster 1 defined
Backup server.
And in non-active
mode.

Administration

Local Console

LAN

LAN

Administration

Remote
Console
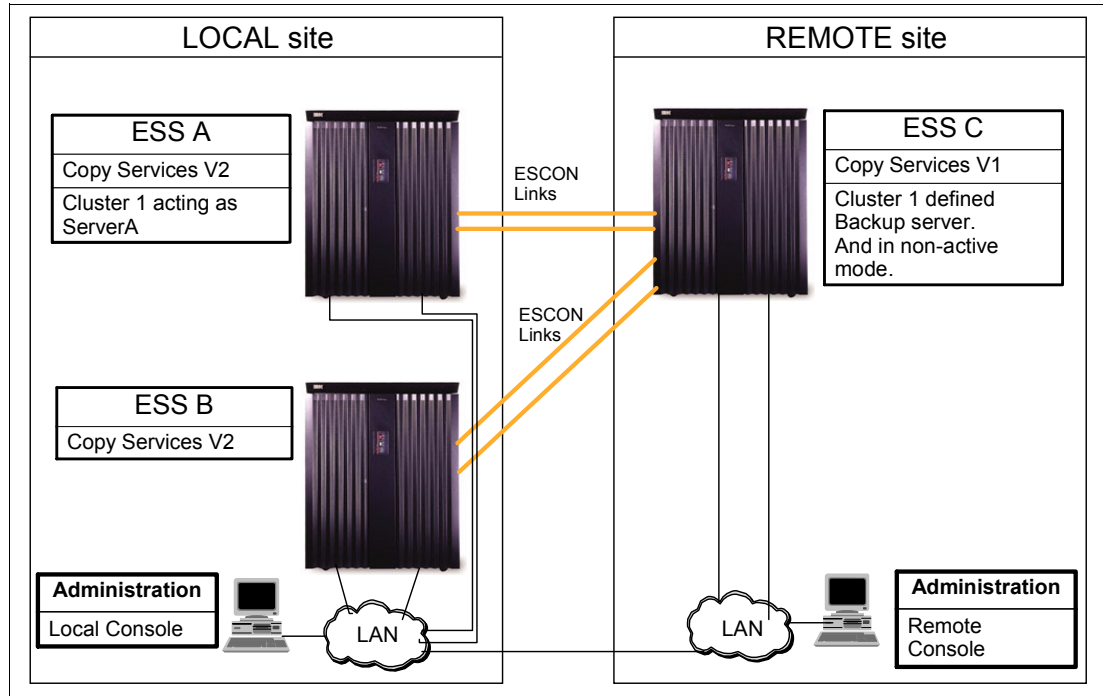
Figure 2-2   Example of a domain

Figure 2-2 and Table 2-2 on page 27 provide details about this mixed configuration A. Here are some specific points about this example:

► One cluster at the local site running Copy Services V2 is defined as ServerA. The ServerA is in dual-active mode. One cluster running Copy Services V1 at the remote site is defined as the Backup server. The Backup server is in single-active mode and is not active.

► The ServerA and ServerB IP address are defined on all the Clusters with Copy Services V2. The clients list is defined on ServerA. A Reset Copy Services procedure was performed on all the V2 clusters or a Domain Wide Reset procedure was performed from serverA to initiate communication between the ServerA and the V2 clients.

► In each of the ESS clusters running Copy services V1 at the remote site, the primary server is defined with the IP address of the ServerA. The Backup server is defined with the IP address of one cluster running Copy Services V1. A Reset to Primary procedure was performed to initiate the communication with the active server and to put the Backup server in non-active mode.

► Task definitions must be performed by starting the Copy Services WUI with ServerA because the server defined as ServerB is not active. The tasks that are not supported by the Backup server will exist only on the server A and their task name will be appended with ~## characters (where ## is a sequence starting from 00 to 99).

► In case of a failure of ServerA, we have to perform a Reset to Backup procedure on all ESS clusters running Copy Services V1. After the procedure, the single-active server (the Backup server) will be the active server for Copy Services at the remote site. You can run existing tasks, and create and run new tasks from the active Backup server. The new tasks will be lost after the recovery of the ServerA.

Table 2-2   Example of mixed configuration

	ESS A	ESS B	ESS C
Site	LOCAL	LOCAL	REMOTE
Copy Services level	V2	V2	V1
Cluster 1 IP Address	10.10.1.101	10.10.1.111	10.10.1.121
Cluster 1 is...	Active server (V2) (Server A)	Client (V2)	Non-active server (V1) (Backup server)
Cluster 1 Copy Services definitions	ServerA=10.10.1.101 ServerB=10.10.1.121 and Clients list and Reset Copy Services procedure	ServerA=10.10.1.101 ServerB=10.10.1.121 and Reset Copy Services procedure	Primary=10.10.1.101 Backup=10.10.1.121 and Reset to Primary procedure
Cluster 2 IP Address	10.10.1.102	10.10.1.112	10.10.1.122
Cluster 2 is...	Client (V2)	Client (V2)	Client (V1)
Cluster 2 Copy Services definitions	ServerA=10.10.1.101 ServerB=10.10.1.121 and Reset Copy Services or Domain Wide Reset procedure	ServerA=10.10.1.101 ServerB=10.10.1.121 and Reset Copy Services procedure	Primary=10.10.1.101 Backup=10.10.1.121 and Reset to Primary procedure

Configuring a mixed configuration B

► On all ESSs running Copy Services V2 (LIC 2.2.0 or above):

– The ESS cluster defined as ServerA is a cluster running Copy Services V2.

– The ESS cluster defined as ServerB is a cluster running Copy Services V2.

► On the ESS running Copy Services V1 (LIC prior 2.2.0)

– The ESS cluster defined as the Primary server is a cluster running Copy Services V2 (this cluster must be the ServerA cluster defined for ESSs running Copy Services V2).

– The ESS cluster defined as the Backup server is running Copy Services V2 (this cluster must be the same as the ServerB for ESSs running Copy Services V2).

Be aware of the following:

► The Copy services clients with Copy Services V1 still must be configured by defining a "Primary server" and a "Backup server". You should define the Primary server with the IP address of the ServerA, and the Backup server with the IP address of the ServerB cluster.

► All the ESS clusters in the domain with LIC prior to 2.2.0 will run the Copy Services V1 Client software. Those clients can only communicate with one active server. You will have to initiate manually the communication to the Primary server using the **Reset to primary** procedure. Those actions are described in 5.3.3, "Restarting ESS Copy Services" on page 108.

► The active ServerB that is defined as the Backup server for Copy services V1 clients will not be able to get the V1 client configuration, unless you perform a **Reset to Backup** procedure on V1 clients. But then, there is no communication anymore from the V1 clients to the ServerA (defined as the Primary server for V1 clients).

Example of mixed configuration B

We give here an example of a two-site configuration with 3 ESSs. Two of the ESSs are running Copy Services V2 and one is running Copy Services V1.
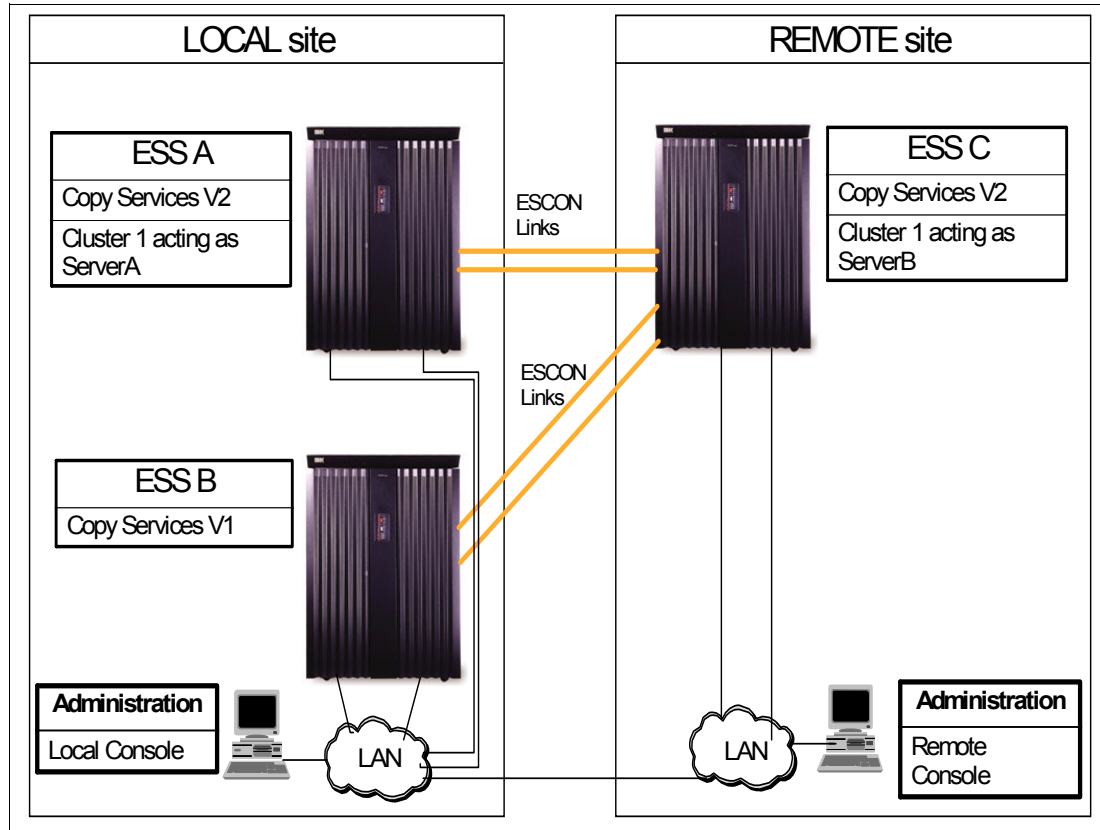
Figure 2-3   Example of a domain

Figure 2-3 and Table 2-3 on page 29 provide details about this mixed configuration B. Here are some specific points about this example:

► ServerA and ServerB are active on two clusters running Copy Services V2. ServerA is at the local site, and ServerB is at the remote site.

► The ServerA and ServerB IP address are defined on all the Clusters with Copy Services V2. The clients list is defined on both ServerA and ServerB. A **Domain Wide Reset** procedure was performed from the ServerA to initiate the communication between the V2 servers and the V2 clients.

► In the ESS running Copy services V1, the two clusters are running the client software. The cluster acting as ServerA is defined as the Primary server in each cluster. The ServerB is defined as the Backup server in each cluster. And a **Reset to Primary** procedure was performed in each cluster to initiate the communication between the V1 client and ServerA.

► We decided to configure the ServerA at the local site because the ESS B with copy Services V1 is also at the local site. This will avoid any manual action if the communication is lost with the remote site.

► Task definitions must be performed by starting the Copy Services WUI with ServerA because ServerB cannot get the configuration of the ESS running the Copy Services client software V1.

► In case of a failure of ServerA, we have to perform a **Reset to Backup** procedure on the ESS B clusters running V1. After the procedure, we will start the Copy Services WUI with ServerB and we will be able to run existing tasks, and create tasks and run new tasks from ServerB. The new tasks will be lost after the recovery of the ServerA.

Table 2-3   Example of configuration: dual-active servers and V1 Clients

	ESS A	ESS B	ESS C
Site	LOCAL	LOCAL	REMOTE
Copy Services	V2	V1	V2
Cluster 1 IP Address	10.10.1.21	10.10.1.31	10.10.1.41
Cluster 1 is...	Active server (V2) (Server A)	Client (V1)	Active server (V2) (ServerB)
Cluster 1Copy Services definitions	ServerA=10.10.1.21 ServerB=10.10.1.41 and Clients list and Domain Wide reset procedure	Primary=10.10.1.21 Backup=10.10.1.41 and Reset to Primary procedure	ServerA=10.10.1.21 ServerB=10.10.1.41 and Clients list
Cluster 2 IP Address	10.10.1.22	10.10.1.32	10.10.1.42
Cluster 2 is...	Client (V2)	Client (V1)	Client (V2)
Cluster 2 Copy Services definitions	ServerA=10.10.1.21 ServerB=10.10.1.41	Primary=10.10.1.21 Backup=10.10.1.41 and Reset to Primary procedure	ServerA=10.10.1.21 ServerB=10.10.1.41

2.5.3 Compatibility matrix

When you define the copy services domain, you must provide only a pair of servers. Table 2-4 describes the supported Copy server pairs configuration in a Copy Services Domain.

To read this table, select, in the first row and in the first column, the two kinds of Copy Services' servers you want to define in your domain. Then, using the intersection between the column and the row, you can check if this server pair is supported.

Table 2-4   Supported pairs of server in the Copy Services domain

	ServerA	ServerB	Primary	Backup	None
ServerA	NO	YES	NO	YES	YES
ServerB	YES	NO	YES[a]	NO	NO

a. This solution is not recommended but is allowed for a short period of time (migration for example). Please note that the primary server task repository will be sent to ServerB correctly, however, ServerB tasks will not be sent to the primary server.

Example of allowed server pairs in a copy services domain

► You cannot have two different ServerAs in the same domain.

► It is allowed to have one ServerA and one Backup server.

► It is not allowed to have one ServerA and one Primary server.

► It is allowed to have only one ServerB defined in the domain (the None column refers to "no other cluster defined as a Copy Services Server").

2.5.4 Task management

The Copy Services servers running Version 2 software will act in the following ways regarding the Task repository:

► At Initialization (when the servers register with each other):

  – Task repository in ServerA will overwrite the task repository in ServerB for the tasks that are supported by both servers.

  – ServerA and ServerB will maintain the tasks that are not supported by the other server. In dual-active mode, tasks that can only be supported by ServerA, but not Server B, exist only on ServerA. Similarly, tasks that can only be supported by ServerB, but not ServerA, only exist on ServerB. These tasks are saved in the repository by appending the characters ~## to the end of the task name, where ## is a sequence number from 00 to 99.

► If you modify a task when two servers are running and communicating with each other (normal configuration and operation):

  – Tasks that are supported by both servers will be stored and maintained in both servers.

  – Tasks that are supported by only the initiating server will be saved with the names appended ~## in the initiating server only.

  – If one or more tasks in a group task are not supported by the other server or their name has ~xx appended, all tasks in that group that are supported by both servers will be cloned. All the tasks within that group will have names with ~## appended. The group task name also has ~## appended.

► If you perform a LIC Upgrade on one Server:

  – Tasks with ~## appended are not automatically updated in the upgraded server.

  – Task with ~## appended must be manually renamed to remove the ~## to cause it to be replicated to the other server.

2.5.5 Inband management considerations

The Copy Services V2 allows you to invoke some FlashCopy commands via inband commands. That means that the active Copy Services server invokes the FlashCopy task to a client using a PPRC link. When using inband commands, the network connection to the client ESS is not used. This option will be useful, for example:

► If you lose the network connection between the sites.

► If the distance between the two sites makes it difficult or impossible to maintain a network connection between the sites.

You should consider the following points when planning to use inband commands for FlashCopy management:

► The source volume you select for the FlashCopy at the remote site must be the secondary PPRC volume of a PPRC pair. (The PPRC secondary volume becomes the source volume in a FlashCopy pair.)

► The network connection is mandatory when you define the tasks.

► There is no support for Inband queries. You cannot get the FlashCopy volumes' information and status. For example, you would have no clue if the background copy for an Inband FlashCopy completed on a pair unless the network connection is set up between the Copy Services server and the ESS hosting the FlashCopy pair.

2.5.6 Migration from V1 to V2

When migrating one or more ESSs from Copy Services V1 to V2 in a domain, ServerA should be defined as the old Primary and ServerB as the old Backup. The V1 ESSs should remain unchanged. This will insure that the task definitions are kept during the migration.

You should terminate all FlashCopy relationships before upgrading an ESS with this new LIC. The FlashCopy relationships are not maintained during the upgrade from V1, pre-LIC 2.2.0, moving to V2, LIC 2.2.0. You do not have to terminate the PPRC relationships, since active PPRC relationships are maintained. You should not define new PPRC relationships until the migration process is completed.

Here are the different points you have to consider when planning for a migration of the ESS LIC levels in a Domain from V1 to V2:

► Check in 2.5.3, "Compatibility matrix" on page 29 for the supported Copy services server pair configuration.

► Check in 2.5.4, "Task management" on page 30 for the behavior of each server regarding the task management. Remember that when a ServerA becomes active, it overwrites the task repository on a ServerB.

► You can define the "old" Backup server as the new ServerA.

► If you want to define a ServerA that is not an "old" Primary or Backup server, you can temporarily change the server definition for the Backup server in the domain. The temporary definition would be that the temporary Backup server will be the ServerA after the migration. Restart the WEB Copy Services after the temporary definition. In this way, the tasks in the Primary will be transferred to the "new" ServerA. After that, one can define the final configuration. This actually can be used as the migration steps to keep the tasks for mixed V1 and V2 ESSs.

3

FlashCopy

Today, more than ever, organizations require their applications to be available 24 hours per day, seven days per week. They require high availability, minimal application downtime for maintenance, and the ability to perform data backups with the shortest possible application outage.

The prime reason for data backup is to provide protection in case of source data loss due to disaster, hardware failure, software failure, or user errors.

Data copies can also be taken for the purposes of program testing or data mining by database query applications. However, normal copy operations take a long time, requiring the prime application to be offline. With the need for 24x7 data processing, there is a need to have an instant copy of the data.

FlashCopy allows you to move effectively towards such solutions.

In this chapter, we discuss the ESS FlashCopy function.

3.1 Overview

FlashCopy provides an instant or point-in-time (PIT) copy of an ESS logical volume. Point-in-time copy functions give you an instantaneous copy, or "view", of what the original data (source) looked like at a specific point in time. This is the so-called *time-zero copy*.

When a FlashCopy task is invoked, the process of establishing the FlashCopy pair and creating the necessary control bitmaps takes only a few seconds to complete. Thereafter, you have access to a time-zero copy of the source volume. As soon as the pair has been established, you can read and write to both the source and the target volumes.

The point-in-time copy created by FlashCopy is typically used where you need a copy of the production data to be produced with minimal application downtime. It can be used for online backup, testing of new applications, or for creating a database for data-mining purposes. The copy looks exactly like the original source volume and is an instantly available, binary copy. See Figure 3-1 for an illustration of FlashCopy concepts.
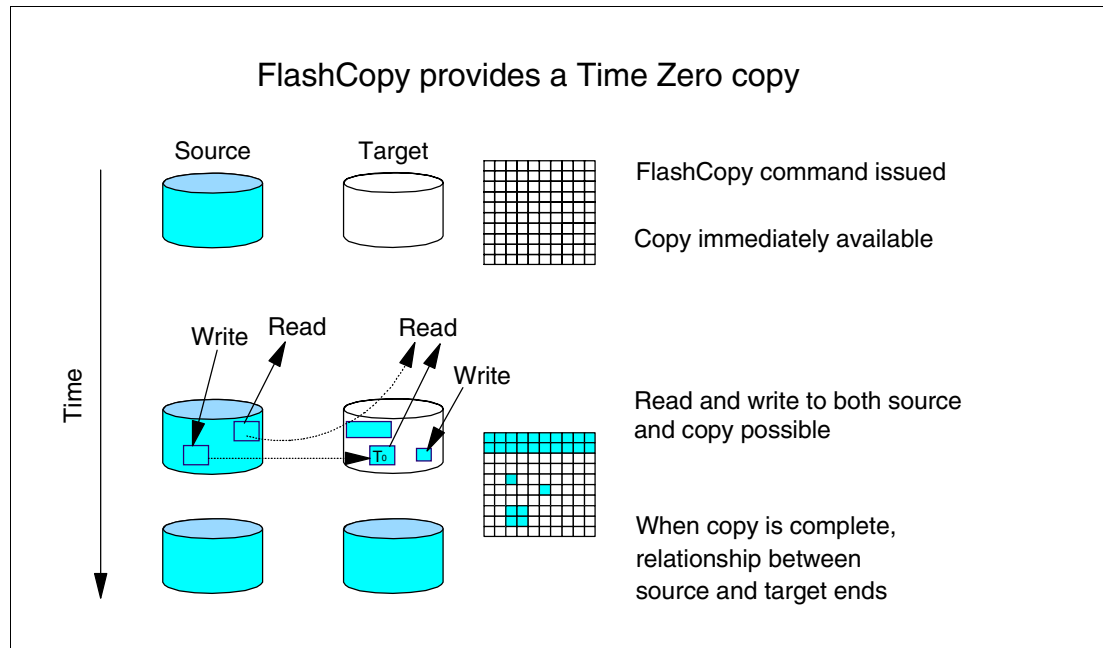
Figure 3-1   FlashCopy concepts

In FlashCopy V1, the target volume is restricted to be in the same logical subsystem (LSS) as the source volume. The source and target volumes can be on the same or on different arrays (also on same or different SSA loops within one DA pair), but only if they are part of the same LSS. In FlashCopy V2, which we discuss later in this chapter, the LSS constraint has been removed, so it is possible to have the source and the target volumes in different logical subsystems (LSS) within a single ESS.

3.2 FlashCopy essentials

When you set up the FlashCopy, a relationship is established between the source and the target volume and a bitmap of the source volume is created. Once this relationship is established and the bitmap created, the target volume can be accessed as though all the data had been physically copied. While a relationship between the source and target volume exists, a background process copies the tracks from the source to the target.

The relationship ends when the physical background copy process has completed or when the relationship is ended by the issuing of the **Withdraw FlashCopy pair** task.

At the time when the FlashCopy is started, the target volume is, in a sense, empty. The background copy process copies data from the source to the target. The FlashCopy bitmap keeps track of which data has been copied from the source to target. If an application wants to read some data from the target that has not yet been copied to the target, the data is read from the source; otherwise, the read is satisfied from the target volume. When the bitmap is updated for a particular piece of data, it signifies that the source data has been copied to the target volume. Further updates to the same area are ignored by FlashCopy. This is the essence of the time-zero point-in-time copy mechanism.

Before data is destaged to a track on the source that has not yet been copied, the original track is copied to the target volume. Reads that are subsequently directed to this track on the target volume are now satisfied from the target volume instead of the source volume. After some time, all tracks will have been copied to the target volume, and the FlashCopy relationship will end, unless the **Persistent FlashCopy** option was used.

You cannot create a FlashCopy on one type of operating system and make it available to a different operating system, unless they support the same structure of storing data. You can make the target available to another host running the same type of operating system. Also, remember that the FlashCopy creates an exact binary copy of the source volume. So if you want to access the target from the same system as the source, special conditions must be met for different platforms (see Chapter 8, "Open systems specifics" on page 327 for more details).

3.3 FlashCopy V2 enhancements

FlashCopy V2 includes support for all previous FlashCopy functions. There are also new enhancements available for the open-systems environment. They include:

► Elimination of the logical subsystem constraint
► Multiple Relationship FlashCopy
► Incremental FlashCopy with Reverse Restore option
► FlashCopy Consistency Groups
► Inband commands over the PPRC link
► Establish time improvement

3.3.1 FlashCopy across different logical subsystems (LSS)

This is a new feature in FlashCopy V2. From now on, the source and the target volumes can span logical subsystems. We can FlashCopy to any volume within a single ESS.

This gives us more flexibility and simplified capacity management and administration. For more information about logical subsystems, see Appendix C, "Logical subsystems" on page 487.

3.3.2 Multiple Relationship FlashCopy

This feature provides the capability for the source volume to have FlashCopy relationships with multiple target volumes simultaneously. This allows you to initiate up to 12 FlashCopy relationships on a given source volume, without needing to first wait for or cause the previous relationships to end.

The limitations using Multiple Relationship FlashCopy are:

► The target volume can only have one source volume.
► The target volume cannot be used as the source volume at the same time (for another FlashCopy relationship).

The source and the target volumes can be spread across any LSS within a single ESS.

3.3.3 Incremental FlashCopy and the Reverse Restore option

This feature provides the capability to "refresh" a volume involved in a FlashCopy relationship. With Incremental FlashCopy, the initial relationship between the source and the target volume is maintained and read/write operations are allowed on both the source and the target (see 6.11.6, "Incremental FlashCopy" on page 216). When a subsequent FlashCopy establish is initiated, only the data that has changed since the last time-zero or incremental copy is copied. This helps reduce background copy completion time when only a subset of the data has changed.

The direction of the "refresh" can also be reversed, in which case the volume previously defined as the target becomes the source for the volume previously defined as the source (and is now the target). Again, only the data that has changed is copied to the volume previously defined as the source. See Figure 3-2 for an illustration of Incremental FlashCopy concepts.
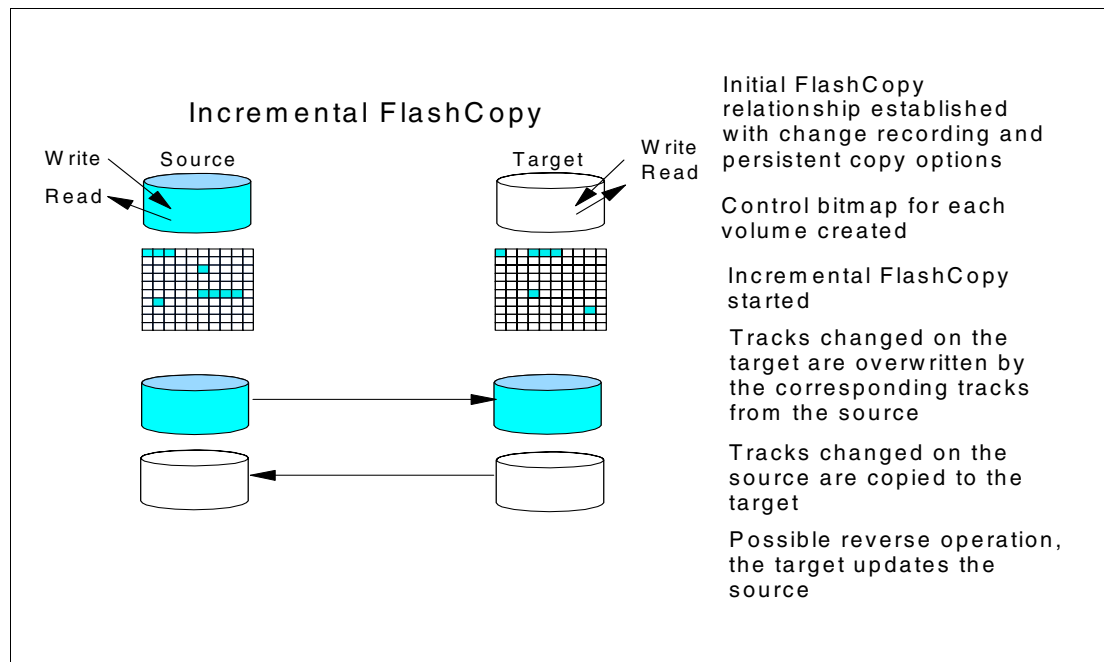
Figure 3-2   Incremental FlashCopy concepts

In Figure 3-2, we have two volumes: a source volume and a target volume. After establishing the FlashCopy relationship between these two volumes with the **Start Change Recording** option enabled (the **Persistent FlashCopy** option is automatically selected), the ESS creates the control bitmap for each volume in this relationship. All tracks that change on either of the volumes are marked in the corresponding bitmap. The Incremental FlashCopy is used to create a point-in-time copy of the source volume, but not by copying the entire volume again. It updates the target volume only with data that has changed on the source volume since the last point-in-time copy. After the **Establish FlashCopy pair** task with **Increment FlashCopy**

option is run, the tracks that have changed on the target are overwritten by the corresponding tracks from the source. Thereafter, the tracks that have changed on the source are copied to the target (tracks that have already been copied will not be copied again). This ensures that the target volume becomes an updated point-in-time copy of the source volume without the need to copy the entire volume again. Figure 3-3 shows an example of how Incremental FlashCopy works.

Incremental FlashCopy example

										Initial FlashCopy established:
Volume A	A1	A2	A3	A4	A5	A6	A7	A8	A9	A1=B1.....A9=B9
Volume B	B1	B2	B3	B4	B5	B6	B7	B8	B9	

Read/Write operations to A and B

Volumes A and B changed:

Volume A	A1c	A2	A3c	A4c	A5	A6	A7	A8c	A9	A1c not equal to B1c; A2 not equal to B2c; A3c not equal to B3; A6=B6
Volume B	B1c	B2c	B3	B4	B5c	B6	B7	B8	B9	

Incremental FlashCopy started

Volume A	A1c	A2	A3c	A4c	A5	A6	A7	A8c	A9	Volume B exact copy of Volume A, only changed tracks copied
Volume B	A1c	A2	A3c	A4c	A5	B6	B7	A8c	B9	

Figure 3-3   Incremental FlashCopy example

In Figure 3-3, we have two volumes: A is the source volume, and B is the target volume. We establish the initial FlashCopy pair with the **Start Change Recording** option enabled. Two bitmaps are created, one to record which tracks have been updated on volume A, the second to record tracks updated on volume B. Each of the volumes have nine tracks: A1 to A9 on volume A, B1 to B9 on volume B. After establishment of the FlashCopy with the **Start Change Recording** option enabled, tracks A1, A3, A4, and A8 have changed on volume A (marked as A1c, A3c, A4c, and A8c in Figure 3-3) and tracks B1, B2, and B5 have changed on volume B (marked as B1c, B2c, and B5c in Figure 3-3). After you run the **Establish FlashCopy pair** task with the **Increment FlashCopy** option from A to B, the target will contain the following tracks: A1c, A2, A3c, A4c, A5, B6, B7, A8c, and B9. Tracks that have not changed (marked with the "equal to" sign in Figure 3-3) are not copied.

> **Note:** The background copy process must be finished before you use the **Reverse Restore** option.

3.3.4 FlashCopy Consistency Groups

New options are available to facilitate the creation of FlashCopy Consistency Groups. With the FlashCopy Consistency Groups, the ESS will hold off I/O activity to a volume until the **Consistency Created** task with the **FlashCopy Consistency Group** option is issued (see Figure 3-4 on page 38).

The **Freeze FlashCopy Consistency Group** option for the **Establish FlashCopy pair** task (see 6.11.9, "Freezing a FlashCopy consistency group" on page 224) provides the capability to ensure data consistency across multiple FlashCopy volumes in an LSS, across LSSs, and even across multiple ESSs (remember that FlashCopy cannot be used to copy data between two or more ESSs, so each source volume must have its target within the same ESS). It

causes volumes to remain in a *Queue Full (QF)* condition, until the **Consistency Created** task with the **FlashCopy Consistency Group** option is run.

The QF status signals to the host that the tagged SCSI command queue is full and that the I/O request has not been placed in the queue. Then the host system will attempt to requeue the requests in the logical unit's device queue.

In other words, FlashCopy Consistency Groups provide the capability to temporarily queue (at the host's level) subsequent write operations to the FlashCopy source volumes that are part of the Consistency Group, until the QF condition is reset by the **Consistency Created** task or the time-out value expires (the default is two minutes; see "Consistency Group timeout" on page 191). To learn more about how to reset the QF condition, see 6.11.10, "Thawing a FlashCopy consistency group" on page 228.

Consistency Groups with FlashCopy

FlashCopy

Ensures data consistency accross multiple volumes, multiple LSS, and even across multiple ESSs

Freeze comand holds off I/O activity to the volumes, which are part of the consistency group

Volumes in "long-busy" condition until reset or timeout

Reset done by "Run Consistency Group after a Freeze command" task (one per LSS)
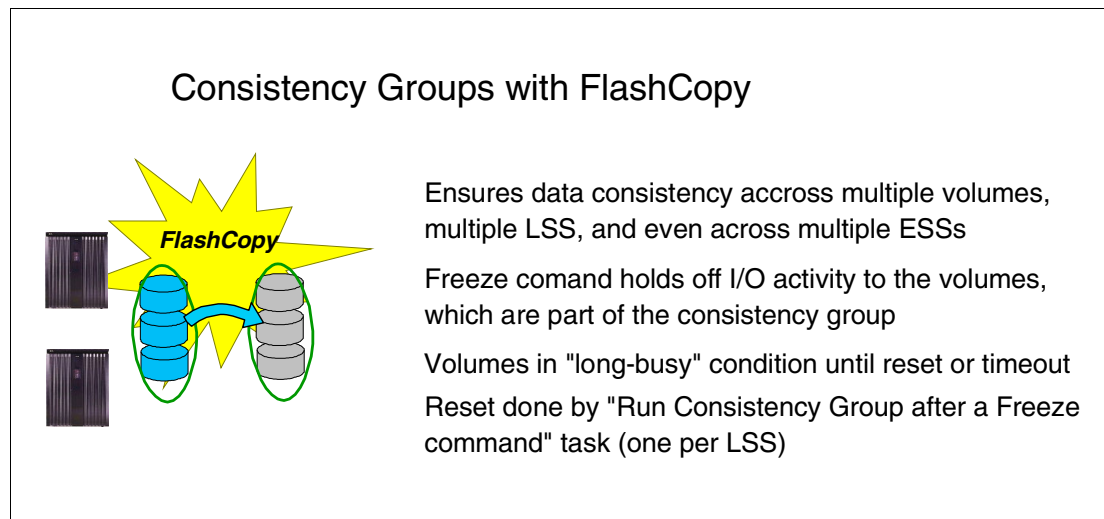
Figure 3-4   Consistency Groups with FlashCopy

The default time-out value can be changed using the ESS Copy Services Web user interface (WUI) (see 6.6.4, "LSS properties" on page 190). Consistency groups can be used to help create a consistent point-in-time copy across multiple volumes, and even across multiple ESSs.

> **Note:** If the Consistency Group source volumes are used with a journaled file system (like AIX JFS) and the source LUNs are not unmounted before running FlashCopy, it is likely that **fsck** will have to be run on the target volume. This is because the file system meta-data is written directly to the source volume without using any system buffers, unlike the data, which is usually written to the source using a periodic *sync* operation, which flushes the buffers to the source volume.

3.3.5 Inband commands over the PPRC link

In a PPRC environment, commands to manage FlashCopy at the remote site can now be issued from the local site and transmitted over PPRC links (see 6.11.11, "Inband FlashCopy" on page 229).

FlashCopy commands are issued to the primary volume of a PPRC pair at the local site. The PPRC pair acts as a conduit to the remote site for the execution of the command at the remote site. The source volume that you select for the FlashCopy operation at the remote site *must be* a secondary PPRC volume of the PPRC pair.

This new function eliminates the need for a network connection to the remote site solely for the management of FlashCopy. Remember that you still need the network connection to the remote site for the purpose of creating the inband FlashCopy task, because you must be able to select the volumes that you intend to use for the FlashCopy relationship.

> **Note:** You must first establish PPRC paths and the PPRC pair from the local ESS logical subsystem to the remote ESS logical subsystem containing the volumes that are used for the Inband FlashCopy task.

3.3.6 Establish time improvement

Performance improvements in FlashCopy V2 are designed to provide up to a 10 times reduction in the time required to complete the **Establish FlashCopy pair** task. With this significant reduction in establish time, operational interruption is further minimized and the benefits of FlashCopy can be extended into new application environments.

3.4 FlashCopy tasks and options

In this section, we discuss various options for both FlashCopy V1 and FlashCopy V2. There are two basic types of FlashCopy tasks that can be used with various options to create the ESS Copy Services tasks:

► **Establish FlashCopy pair**: This type of task establishes a FlashCopy pair between source and target volume, so they are in the FlashCopy relationship. To learn more, see 6.11.1, "Establishing a FlashCopy pair" on page 204.

► **Withdraw FlashCopy pair**: This type of task is used to withdraw the FlashCopy pair, so the relationship between source and target volume ends. To learn more, see 6.11.3, "Withdrawing a FlashCopy pair" on page 209.

> **Note:** In this chapter, we will use references to other chapters, mainly to the ESS Copy Services Web user interface (WUI), which is now divided into two separate chapters, as there are two versions of the ESS Copy Services and WUI. The reader should refer to the sections according to the section names specified, but in the chapter corresponding to the version of the WUI used.

3.4.1 FlashCopy V1 establishment options

With FlashCopy V1, you can use the following options for the **Establish FlashCopy pair** task type to create the ESS Copy Services tasks:

No background copy

You can suppress the background copy process using the **No background copy** (NOCOPY) option. This may be useful if you need the copy only for a short time, such as making a backup to tape. If you perform FlashCopy with the NOCOPY option, data will only be copied from the source volume to the target volume if a track on the source volume is modified.

Generally the FlashCopy relationship between source and target volumes remains indefinitely and has to be broken manually. The best way to break the relationship is to use the **Withdraw FlashCopy pair** task.

If you need to create a permanent physical copy of the data for backup or disaster recovery purposes, you can perform a **Withdraw FlashCopy pair** operation with the **FlashCopy Start**

Background Copy option (see "FlashCopy Start Background Copy" on page 41). All data will be copied from the source volume to the target volume. After the copy is complete, the FlashCopy relationship is terminated, unless the FlashCopy relationship is persistent, in which case, the relationship remains. Refer to "Persistent FlashCopy" on page 40 for further information.

> **Note:** After a FlashCopy pair is established with the **No background copy** option, an automatic withdrawal of the FlashCopy pair is caused when all sectors or cylinders on the source volume are modified. This rarely occurs, as it is unlikely that the whole contents of a disk volume is modified.

Accelerated destage mode

Select this option to cause a FlashCopy source volume track, which is being modified and is in a FlashCopy relationship, to be destaged from cache sooner than it would be if normal cache algorithms were applied. This feature minimizes (for the source volume) the number of modified tracks that are resident in cache. With FlashCopy Version 1, this option might (under specific circumstances) speed up the completion of the background copy process.

This feature has no effect with FlashCopy Version 2, as the cache algorithms have changed.

Permit establish if target is online

Select this option if you want to establish a FlashCopy relationship, even if the target is online, to the S/390 or zSeries host.

> **Note:** This option does not apply to the open-system volumes.

Persistent FlashCopy

If you want to retain the FlashCopy relationship after the background copy completes, you can use the **Persistent FlashCopy** option. The FlashCopy relationship between the source and the target volume remains indefinitely and must be broken by a **Withdraw FlashCopy pair** task (see 5.10.3, "Withdrawing a FlashCopy pair" on page 135).

This has no impact on host performance once the background copy is finished. The **Persistent FlashCopy** option is usually used to protect the target volume from being overwritten by another ESS copy services task. With FlashCopy version 2, it is a prerequisite for using Incremental FlashCopy.

3.4.2 FlashCopy V1 withdraw options

With FlashCopy V1, you can use the following options for the **Withdraw FlashCopy pair** task type to create the ESS Copy Services tasks.

FlashCopy withdraw to target

Use this option if you want to remove the I-am-a-copy icon from the target volume (5.10.3, "Withdrawing a FlashCopy pair" on page 135). Typically, you remove the indicator because you decide it is no longer necessary to mark the volume as a copy of another volume. This function might be useful also in other situations (like if there is a source volume missing). See "Withdraw FlashCopy to the target" on page 137 for further details.

> **Note:** To use this option, you must select the target volume of the established FlashCopy pair as a source and also as a target in the Volumes panel of the ESS Copy Services Web user interface (WUI).

FlashCopy Start Background Copy

This option allows a customer to change a FlashCopy from a NOCOPY to a COPY, which will speed up the completion of creating the target volume. Only the remaining tracks are copied to the target volume.

This is an option for the **Withdraw FlashCopy pair** task type. See 5.10.4, "FlashCopy Start Background Copy" on page 139 for further details.

> **Note:** To use this option, you must select the source volume of the established FlashCopy pair as a source and also as a target in the Volumes panel of the ESS Copy Services Web user interface (WUI).

3.4.3 FlashCopy V2 establishment options

In addition to FlashCopy V1 options, with FlashCopy V2, you can use the following options for the **Establish FlashCopy pair** task type to create the ESS Copy Services tasks.

Inband command

Select this option to establish a FlashCopy relationship between two volumes at a remote ESS using inband commands; the FlashCopy source volume must be the PPRC secondary volume (see 6.11.11, "Inband FlashCopy" on page 229). Inband commands are issued to a primary volume of a PPRC pair at the local ESS and sent across PPRC paths (acting as a conduit) to a remote ESS to establish FlashCopy at the remote site.

This eliminates the need for the network connection to the remote site solely for FlashCopy management. Remember that you still need the network connection to the remote site for the purpose of creating the inband FlashCopy task, because you must be able to select the volumes which you intend to use for the FlashCopy relationship.

> **Note:** You must first establish PPRC paths from the local ESS logical subsystem to the remote ESS logical subsystem containing the volumes that are used for the Inband FlashCopy task.

Freeze FlashCopy Consistency Group

This option causes the source volume(s) to go into the Queue Full condition. All writes to the volume(s) are queued by the host and are written after the condition is reset, either using the **Consistency Created** task with the **FlashCopy Consistency Group** option or by the time-out value, which has a default value of two minutes. See 6.6.4, "LSS properties" on page 190 to see how to modify this value. Consistency groups can be used to help create a consistent point-in-time copy across multiple volumes, and even across multiple ESSs.

Inhibit writes to target

If you select this option, any write operation to the target volume of the FlashCopy pair will fail as long as the FlashCopy relationship exists. This is used if you want to protect the target volume from any updates from the host side.

Be careful while accessing the target volume from the same host system as the source while write operations to the target volume are inhibited. This might cause serious problems (depending on the operating system or volume manager software used), as information which has to be changed on the target volume to differentiate between the two volumes (because the target volume is an exact binary copy of the source volume) is not written.

Note: Remember that while accessing the target volume (even in read only mode), the operating system might need to update some piece of data on the target volume. So if writes are inhibited to the target volume while accessing it, it might produce unpredictable results depending on the volume manager or operating system used.

Start Change Recording

This option is used to create the bitmaps that record the changed tracks on both volumes within a FlashCopy pair. It is a prerequisite when you want to use the Incremental FlashCopy. This option has to be selected (**Persistent FlashCopy** option is autoselected) while you are establishing the initial FlashCopy pair that you intend to use with the Incremental FlashCopy (see 6.11.6, "Incremental FlashCopy" on page 216).

If you want to run subsequent Incremental FlashCopy tasks on the same volume pair, this function has to be selected each time you increment the FlashCopy.

Incremental FlashCopy

This option is used for incrementing the existing FlashCopy pair. The option ensures that only the changed tracks are copied from the source to the target volume (for more details see 3.3.3, "Incremental FlashCopy and the Reverse Restore option" on page 36). Before using this option, there must be an established FlashCopy pair with the **Start Change Recording** option available.

Reverse Restore

Use this option to reverse an established FlashCopy pair. After establishing the FlashCopy with the **Reverse Restore** option enabled, the volume previously defined as a source becomes a target for the volume previously defined as a target (and now is the source). To be able to use this option, there must be an established FlashCopy pair with the **Start Change Recording** option available. For more detailed information about how this function works, refer to 3.3.3, "Incremental FlashCopy and the Reverse Restore option" on page 36.

Note: The background copy process must be finished prior to using this option.

3.4.4 FlashCopy V2 withdraw options

In addition to FlashCopy V1 options, with FlashCopy V2, you can use the following options for the **Withdraw FlashCopy pair** task type to create the ESS Copy Services tasks.

Inband command

Select this option to withdraw the FlashCopy relationship between two volumes at a remote ESS using inband commands (see 6.11.11, "Inband FlashCopy" on page 229). Inband commands are issued to a primary volume of a PPRC pair at the local ESS and sent across PPRC paths (acting as a conduit) to a remote ESS to withdraw the FlashCopy at the remote site.

This eliminates the need for the network connection to the remote site solely for FlashCopy management. Remember that you still need the network connection to the remote site for the purpose of creating the inband FlashCopy task.

Note: The Inband command can also be used with the **FlashCopy Start Background Copy** option.

3.4.5 Consistency Created task with FlashCopy V2

In addition to the **Establish FlashCopy pair** and **Withdraw FlashCopy pair** tasks, there is a new type of task called **Consistency Created**, which is issued on the LSS level (from the Logical Subsystems panel in the ESS Copy Services Web user interface (WUI)).

For FlashCopy, it is created with the following option.

FlashCopy Consistency Group (One LSS Selected)

This option is used to reset the Queue Full condition after the **Establish FlashCopy pair** task with the **Freeze FlashCopy Consistency Group** option was run. All write operations to the source volume(s) that were queued by the host are written as soon as this task completes. See 6.11.10, "Thawing a FlashCopy consistency group" on page 228 to learn how to create the task. For more information about FlashCopy Consistency Groups, see 3.3.4, "FlashCopy Consistency Groups" on page 37.

> **Note:** Because this task is created on the LSS level, it is run on all the FlashCopy pairs that are in the Queue Full condition within the selected LSS.

3.5 Planning for FlashCopy

Because FlashCopy invariably will be used on production systems, you should carefully plan the setup of your environment and test it thoroughly. This is an important step to minimize the possibility of error and potential rework.

3.5.1 Hardware and software requirements

If you want to use FlashCopy, you need to comply with the following prerequisites:

► Have a FlashCopy feature purchased and enabled on your Enterprise Storage Server (ESS) by the Customer Engineer (CE). The feature code is dependent on the total disk capacity of your ESS, rather than on the capacity of the volumes that will use FlashCopy. It is also dependent on whether you want to use FlashCopy V1 or FlashCopy V2. Refer to Appendix B, "Copy Services feature codes" on page 479 for the feature code information.

► On the server that will have the FlashCopy target volumes attached, you need to have enough SCSI target IDs and/or SCSI/Fibre Channel LUNs available (not occupied by volumes). The ESS can have up to 15 SCSI target IDs each with up to 64 LUNs on one SCSI channel and up to 4095 LUNs on a Fibre Channel port.

► You need TCP/IP connectivity between the ESS and the host system that will initiate FlashCopy (usually this is the system that will access the FlashCopy target) in order to use the Copy Services Command Line Interface (CLI). You can achieve that by connecting the ESS to the company intranet. You have to install the CLI on the host that will be using it.

► If you have Independent Software Vendor (ISV) software installed, you need to contact the ISV regarding their support for ESS Copy Services.

► Review your volume manager software considerations for FlashCopy (refer to Chapter 8, "Open systems specifics" on page 327):
  – AIX LVM
  – VERITAS VxVM
  – HP SAM
  – SUN Solaris Solstice DiskSuite

3.5.2 Configuration planning

The most important consideration that has to be taken into account while using FlashCopy V1 is to have an available volume (LUN) in the logical subsystem (LSS) where the source volume resides. With FlashCopy V2, you can use any LUN within a single ESS as a target. The target LUN has to be of the same size as the source or bigger. The space for target data has to be available even if only the **No background copy** option will be used.

For an understanding of LSS concepts, see Appendix C, "Logical subsystems" on page 487.

> **Note:** If using FlashCopy V1, you may need to review your configuration to have a target volume available in the source LSS. That will usually mean copying some of your data to another LSS, using either the host system or the PPRC links.

3.5.3 Resource planning

When planning your ESS volume layout, it is important to consider the capacity you may need for FlashCopy targets. Bear in mind that the disk space you need is real disk space. For FlashCopy V1, you must also consider that the target volume is restricted to the same LSS as its source volume. So, when you allocate additional storage for FlashCopy targets, consider how much space in each LSS you need to leave unallocated.

With Version 1, you cannot initiate a FlashCopy session on a source and target that are already in a FlashCopy session. You need to wait for the FlashCopy task to complete or you can withdraw the pair manually. If you have used the **No background copy** option, you always need to withdraw the pair (unless all tracks on the source volume have changed).

3.5.4 Data consistency considerations

It is very important to verify that the copy of the data you will be using is fully consistent using a proper file system check procedure provided by your operating system. If you are going to automate your FlashCopy procedures, consider including this check each time when you make the FlashCopy target available to the host. In all cases before starting the FlashCopy procedure, the target volume must be unmounted; this is to ensure that there is no data in any system buffers that can be flushed to the target and could potentially corrupt it.

3.5.5 Test plan and disaster recovery plan

If you plan to use FlashCopy, you need to test your setup. Do not forget that you are dealing with a binary copy of the data that was done out of the control of your operating system. Prepare a test plan and, if you are using FlashCopy for backup/restore, a recovery plan.

3.6 Operational considerations

The following sections are intended to help you manage your FlashCopy pairs and, in particular, to manage the target you create.

3.6.1 Monitoring and managing FlashCopy pairs and volumes

FlashCopy pairs and tasks can be managed by both the ESS Copy Services Web User Interface (WUI) and the Command Line Interface (CLI) on the host.

The ESS Copy Services Web User Interface (WUI) will allow you to manage FlashCopy volumes and tasks (for more information, see 5.10, "Performing FlashCopy with the Web User Interface" on page 130 and 6.11, "Performing FlashCopy operations with the WUI" on page 203). You can establish and withdraw a FlashCopy by clicking on the graphical representations of the volumes in the WUI. If you wish to perform a FlashCopy from the CLI (for more information, see Chapter 7, "ESS Command Line Interface" on page 301), you must create a FlashCopy task within the WUI and save it. You can either execute your tasks from the WUI or call them with the `rsExecuteTask.sh` command in the CLI.

Using the CLI with predefined tasks minimizes the danger of a human error when handling physical volumes by their volume numbers or names from the ESS Copy Services Web User Interface, and it enables automation.

3.6.2 Using a FlashCopy target volume

Remember that if you have established a FlashCopy with the **No background copy** option, you need to withdraw the FlashCopy pair after you have finished using the FlashCopy target volume. If you choose to perform a full copy, the relationship will be withdrawn automatically when the background copy task ends. The performance issues of using FlashCopy with full copy and **No background copy** options are discussed in 3.7, "Performance considerations" on page 46.

If you will be using FlashCopy for data backup purposes, change your recovery procedure so that you will be able to recover even when the data has been backed up by a different backup client than the original owner of the LUN, or it has been backed up from a different location in the file system (the target mount point).

You can, of course, perform the FlashCopy from the target volume to the original LUN using the **full copy** option (but only *after* the background copy has completed).

> **Note:** Do not attempt to defragment or optimize a FlashCopy source volume while the FlashCopy background copy task is running. This can significantly degrade the performance.

3.6.3 Automation

Different operating systems allow different levels of automation. The automation can be done using batch or script files. This is useful, for example, to create an automatic daily backup procedure.

In the script file that you use for creating the target volume(s), you should include:

1. Quiescing of an application (switching on the backup mode). The proper quiesce procedure is provided in your application's documentation.

2. Flushing data to the source volumes. This can be accomplished by unmounting the source volumes, but sometimes it may be necessary to shut down the source server.

3. Unmounting the FlashCopy target volumes from the target system (if they were already used).

4. Establishment of FlashCopy pair(s) using the Copy Services Command Line Interface with the `rsExecuteTask.sh` command.

5. Resuming an application (terminating the backup mode) using the procedure described in your application's documentation.

6. Hardware scan for new disks on a target system (in case their definitions are not already present).

Optional steps:

7. Verifying the consistency of FlashCopy target volumes.

8. Mounting the target volumes on a target system.

After you finish using the target volume (for example, backup to tape has completed), you should run a script including the following steps:

► Unmounting the target volumes from a target system (if you did not do this before running the FlashCopy).

► Withdrawal of the FlashCopy pair (in case the **No background copy** option was used) with `rsExecuteTask.sh`.

For detailed information about automation on different operating systems, see Chapter 8, "Open systems specifics" on page 327 and Chapter 9, "IBM AS/400 and iSeries" on page 373.

3.7 Performance considerations

The next section is intended to give you the considerations involved when setting up the Copy Services of the Enterprise Storage Server (ESS) in order to achieve better performance.

This should help you understand the performance impact of ESS Copy Services. As there are many different parameters that have an influence on performance, such as applications, kind of workload, and the configuration of the Enterprise Storage Server, this information should serve as a guideline when planning for ESS Copy Services.

Keep in mind that the general ESS performance considerations, such as the volume placement or amount of storage per host adapter, still apply when planning for FlashCopy.

3.7.1 Placement of source and target volume

As we discussed earlier in this book, with FlashCopy V1, the source and target volume of a FlashCopy pair must be in the same logical subsystem (LSS). In certain configurations of the ESS, a single LSS includes more than one RAID 5 or RAID 10 Array (rank). In such configurations, we recommend that you use a target and source volume from different ranks to make your FlashCopy pair. Furthermore, we recommend that you use ranks for your source and target volumes that reside on different SSA loops of the Device Adapter (DA) pair, if this is possible in your ESS configuration (in open-systems environment, this is only possible if the ESS is set up to use 8 Fixed Block logical subsystems). This will distribute the I/O load over more disks, as data is copied from the source to the target volume.

With FlashCopy V2, we recommend that you use different logical subsystems for the source and the target volumes, so the I/O operations are distributed over different SSA Device Adapters.

If you are making a FlashCopy with the **Full copy** option when all data from the source is physically copied to the target (default), and you do not have to work with the target volume directly after the FlashCopy was issued, we recommend waiting until the background copy is finished. This will give you better performance for both the source and the target volume, as host I/O requests will not interfere with I/O of the FlashCopy task. The progress of the

FlashCopy background copy process can be determined with the Copy Services Web User Interface and the completion of the process by using the CLI.

3.7.2 No background copy option

Consider whether you want to select the **No background copy** option for FlashCopy, or not. If you have mainly read access to the source and the target of your FlashCopy pair, we recommend using the **No background copy** option to minimize I/O traffic to the RAID arrays. Keep in mind that when selecting the **No background copy** option, the relationship between the FlashCopy source and target stays until the pair is withdrawn manually (unless all tracks on the source volume have been changed).

3.7.3 Number of simultaneous FlashCopy pairs

Consider the number of FlashCopy pairs you have active at the same time. The time for a single FlashCopy pair establish and the copy time to finish will increase with the amount of FlashCopy pairs you have established at the same time as data is copied in between all pairs. Try to logically group FlashCopy pairs together for ease of use.

3.8 Practical examples using FlashCopy

In this section, we discuss various practical examples of using FlashCopy.

3.8.1 Moving and migrating data

Any time you need to move data from one server to another, FlashCopy can be useful. Do not forget to quiesce disk access before making a FlashCopy and verifying the consistency of the data before attaching it to the target server.

3.8.2 Moving workload

The same way you move the data from one server to another, you can move workload between servers. To move workload, you can usually only reassign the volume in the ESS Specialist. For details, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448, or for more information on the ESS Storage Management CLI, see Chapter 7, "ESS Command Line Interface" on page 301.

3.8.3 Backup

FlashCopy does not usually speed up your physical backup but it allows you to run your application while you back up; therefore, the physical backup completion time becomes less important for you. You may need fewer tape drives and a less powerful Backup server. You need to shut down your application only for the time the FlashCopy task is started and can restart it almost immediately when the task is completed.

While using FlashCopy V1, the performance of the LSS in which the FlashCopy pair resides may be reduced by a very small percentage during the background copy process.

You can keep FlashCopy targets online after you back them up, so you will be able to copy the files that need to be restored from the FlashCopy target rather than having to restore from tape.

> **Attention:** Be careful if you have the two volumes online to the same host with the same vol IDs.

FlashCopy also enables you to do backups whenever you want, not only during the off-shift period, because you do not need to wait for a reduced server load in order to do the backup.

You can use all your existing backup software to do backups. However, in case you intend to do full and incremental (differential) backups, you need to check how your software records the files that have been backed up. In some cases, it marks the files as "archived" on the target disk. This change, of course, is not reflected in the FlashCopy source volume, so anytime you withdraw the FlashCopy target and you attempt incremental backup on the next FlashCopy of that source, you back up all the files again.

There are three methods you can use to restore the data when the target is still available:

1. Mapping the target to the source host and the original mountpoint/drive letter using the ESS Specialist or the ESS Storage Management CLI.

2. Copying the data back to source either using the standard operating system means or doing a reverse FlashCopy, that is, from the original target back to the source. Note that you can only do this if the full background copy has completed.

3. Creating a new FlashCopy of the target and assigning it to the source host.

Database offline backup

The scenario shown in Figure 3-5 provides very fast backups and restores of a database. Logically, the data is being copied from copy (A) to copy (B).
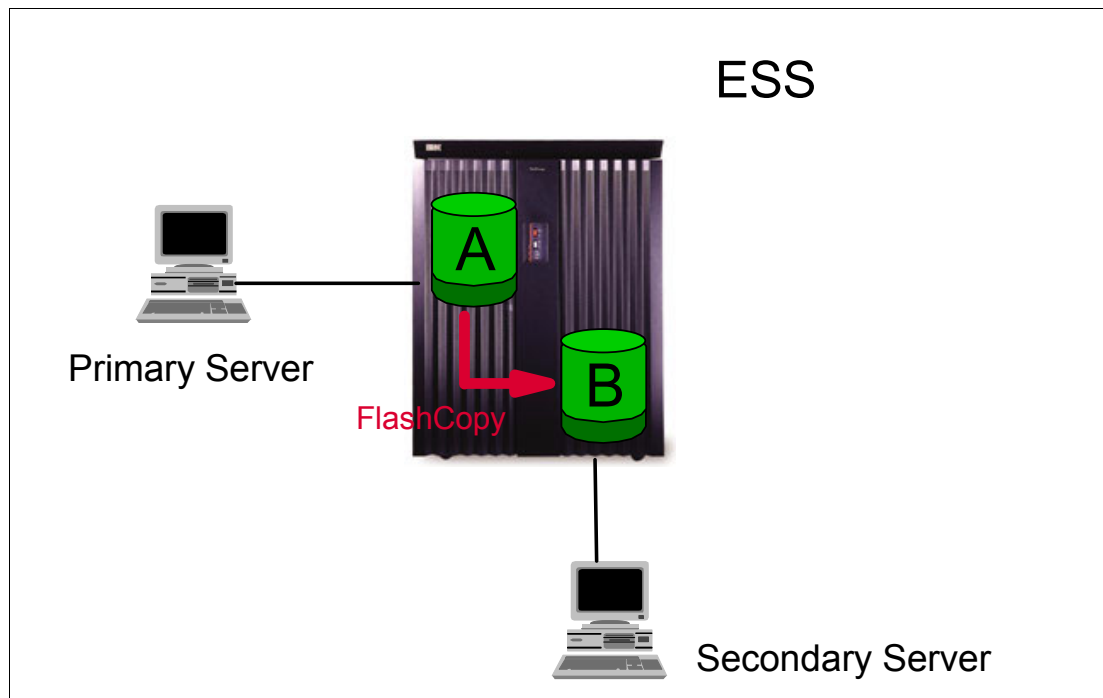
Figure 3-5   Database Backup scenario

The steps to perform the backup are as follows:

If application consistency is required, include these first two steps. For a discussion of consistency, please go to 10.2, "Data consistency" on page 402.

1. Stop production services and the database server.

2. Flush the file system buffer on the production server and secondary server.

Then continue with these steps:

3. Start your FlashCopy task with the WUI or the CLI.

4. As soon as the relationship has been established, restart production services and the database server.

5. Refresh the secondary file system (make the secondary server aware of the new disks if needed).

6. Attach the database on the secondary site.

7. Perform database integrity checks at the secondary site.

8. Start production services on the secondary site (optional).

Database online backups

If you require 24x7 availability, you can use the scenario shown in Figure 3-5 on page 48 to perform an online backup instead of an offline backup of your database. While the backup is not necessarily in a known state, it is still consistent and can be used for recovery. In general, the online backup is done by quiescing the data base on the production server, performing the FlashCopy and then restarting the data base. You can accomplish this by using the supported backup features of your database. For further information, see your database documentation or refer to the ESS whitepapers on the following Web site:

http://www.storage.ibm.com/disk/ess/whitepaper.htm

3.8.4 Application testing

You can test new applications and new operating system releases against a FlashCopy of your production data. Therefore, you can eliminate the risk of data corruption, and your application does not need to be taken offline for an extended period of time in order to perform the copy of the data.

3.8.5 Other examples

Data mining is a good example of where FlashCopy can help you. Data mining can now extract data without affecting your application.

In the following sections, we give examples for the usage of FlashCopy and PPRC. The examples are intended to show solutions that are possible with ESS Copy Services. They are not related to a specific operating system and could be used on all supported Operating Systems.

Moving from a single-host environment to a clustered environment

In a clustered environment, two or even more servers are accessing the same resources, such as disk drives. There are different clustering models available that provide the following benefits. A cluster increases the availability of the data, as the control of the resources from a server that is not available anymore is transferred to the remaining server(s). The failover, failback, and access to the resources is controlled by special cluster software running on all servers within a cluster. Depending on the clustering model, the workload may also be shared between the nodes within a cluster.

If you are currently running your applications in a non-clustered environment and you are planning to move to a clustered-environment, ESS FlashCopy functionality can be used to

create a test environment for the cluster very quickly and easily with minimal impact on production.

In order to test the cluster, you need to prepare the new host systems of the cluster in advance. This is needed to create an identical environment from the host point of view.

Issue the FlashCopy command on all the volumes you want to use in the cluster environment later on. The FlashCopy target volumes are shared with the new hosts of the cluster. Disk sharing with the Enterprise Storage Server is very easy, as the FlashCopy target LUNs only need to be assigned to multiple host adapters. This is done with the ESS Specialist Web User interface. It is also possible to move to Fibre Channel connectivity in the clustered environment if the shared target volumes are connected to a Fibre Channel host port.

Once the FlashCopy pair is established, the shared target volumes with the time-zero copy of the production data could be tested in the clustered environment. However, we recommend waiting until the data copy from the source to the target is finished. This will ensure that there is minimal impact on production I/O when testing on the FlashCopy targets.

This example is illustrated in Figure 3-6.
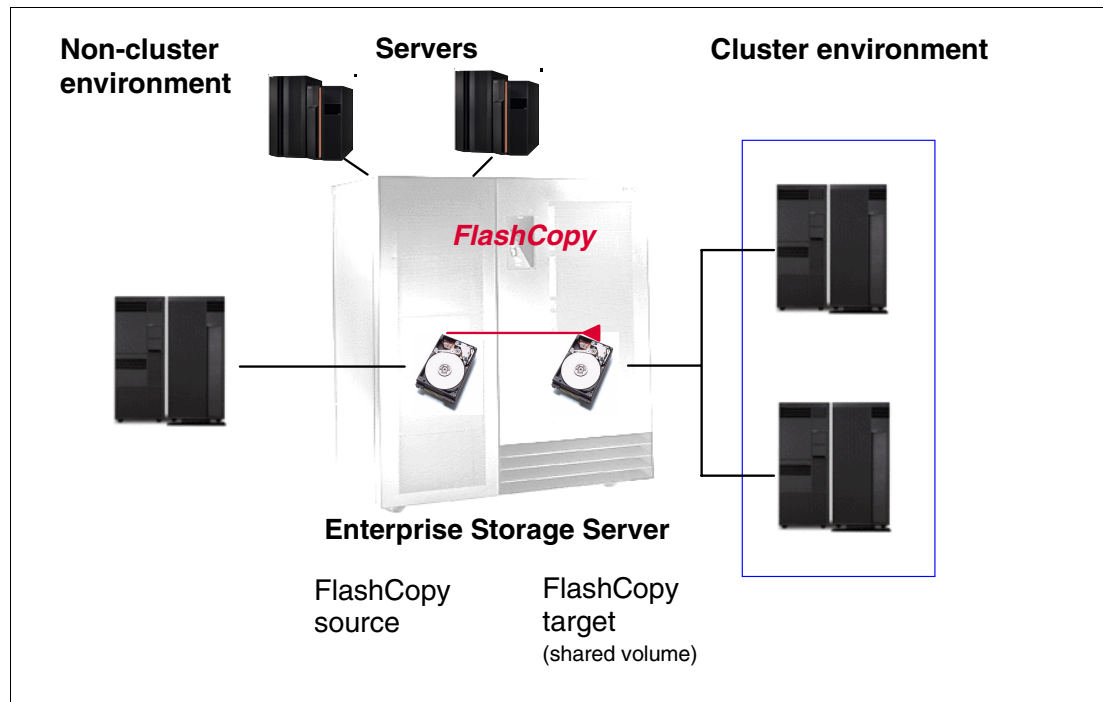
Figure 3-6   Moving cluster environment using FlashCopy

Testing a cluster environment

If there is already a clustered environment, you can use FlashCopy to duplicate the shared data on the target volume(s). In this case, both the FlashCopy source and target volumes are LUNs that are shared on the ESS. This may be applicable if you want to test your cluster without having any impact on the production environment. Refer to Chapter 8, "Open systems specifics" on page 327 for information on the issue of having duplicate vol IDs online to the same host.

Data backup with minimal impact on production

ESS FlashCopy functionality can be used to do online backups of production data. As the data on the FlashCopy target is immediately available after the FlashCopy relationship is

established, you need to ensure a consistent state of the data on the source volume for a few seconds only. After the short time of FlashCopy establishment, the data on the target could be written to back up media immediately. This decreases the time of backup windows of production data significantly.

4

Peer-to-Peer Remote Copy (PPRC)

In this chapter, we describe the ESS Peer-to-Peer Remote Copy (PPRC) function in detail, covering its two variations: the synchronous PPRC (PPRC-SYNC) and the non-synchronous PPRC Extended Distance (PPRC-XD).

This chapter also discusses the combined implementation of PPRC-SYNC and PPRC-XD when configuring an Asynchronous Cascading PPRC environment. The chapter is organized as follows:

► PPRC concepts
► Synchronous PPRC
► PPRC Extended Distance
► Asynchronous Cascading PPRC
► Planning for PPRC
► PPRC Connectivity
► Practical uses of PPRC

This chapter helps you to understand how PPRC works and to plan for its usage. Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103, Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169, and Chapter 7, "ESS Command Line Interface" on page 301 show you how to set up PPRC step by step. PPRC is also covered in Chapter 10, "Disaster recovery" on page 399.

4.1 PPRC overview

PPRC is a function of a storage server that constantly updates a secondary copy of a volume to match changes made to a primary volume. The primary and the secondary volumes can be on the same storage server or on separate storage servers.

PPRC differs from FlashCopy in two essential ways. First, as the name implies, the primary and secondary volumes can be located at some distance from each other. Second, and more significantly, PPRC is not aimed at capturing the state of the source at some point in time, but rather aims at reflecting all changes made to the source data at the target.

The ESS implementation of PPRC provides two modes of operation:

Synchronous PPRC Maintains a consistent copy. All modifications that any attached host performs on the primary volume are also performed on the secondary logical volume before operation completion is indicated.

PPRC Extended Distance Maintains a fuzzy copy. All modifications that any attached host performs on the primary volume are also performed on the secondary volume at a later point in time. The original order of update is not strictly maintained.

PPRC is application independent. Because the copying function occurs at the disk subsystem level, the host's operating system or application has no knowledge of its existence. In contrast to that, host-based mirroring is controlled by software at the operating system or file system level: The storage subsystem does not know about that. Table 4-1 summarizes characteristics of both approaches.

Table 4-1   Comparison of PPRC and host-based mirroring

Peer-to-Peer Remote Copy	Host-based mirroring
Operation is performed by storage subsystem, transparent for host operating system. The functionality is the same for all operating systems and applications.	Operation is performed by host software or host bus adapter, transparent for storage subsystem. The functionality depends on capabilities of the operating system or host bus adapter.
Read and write operations are sent to the primary volume only.	Write operations are sent to both volumes. Read operations are sent to any volume, depending on read policy.
There is an unidirectional relationship from the primary to the secondary volume. Failure recovery is different for the primary and secondary volume.	The relationship between the volumes is symmetric. Failure recovery is identical for both volumes.

4.2 Synchronous PPRC

PPRC can be viewed as having two phases: The first is an initial copy phase involving a bulk transfer of all of the data at the primary site to the secondary site. Large amounts of data are copied in large units. During a second phase, modifications are transferred that occur at the primary site; this may include updates that occurred during the first phase. This continuous phase involves transferring smaller units of data, according to the I/O size of the applications.

4.2.1 Synchronous PPRC initial copy

When a PPRC copy pair relationship has been initially established or re-established after being suspended, the volumes are in a *duplex pending* state (see 4.2.3, "Synchronous PPRC volume states" on page 56). While in this state, PPRC maintains a bitmap to keep a record of updated tracks. On an initial copy pair establish, PPRC will parse the volume copying of all the tracks (as if all the bits are set). On the second parse of the volume, PPRC will only copy tracks that were changed since the last parse.

While in this state, the PPRC algorithm will signal a write complete to the application server as soon as the write I/O has been committed to the primary ESS cache/NVS structure. Once the primary and secondary volumes have been synchronized, the volume state changes to *duplex* and the PPRC algorithm now signals the completion of a write cycle, once it receives acknowledgement from the secondary ESS that the data has been committed to its cache and NVS. The sequence of events is illustrated in Figure 4-1. (The icons are from the ESS Copy Services version 1 panels; for version 2 icons, see Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169.)
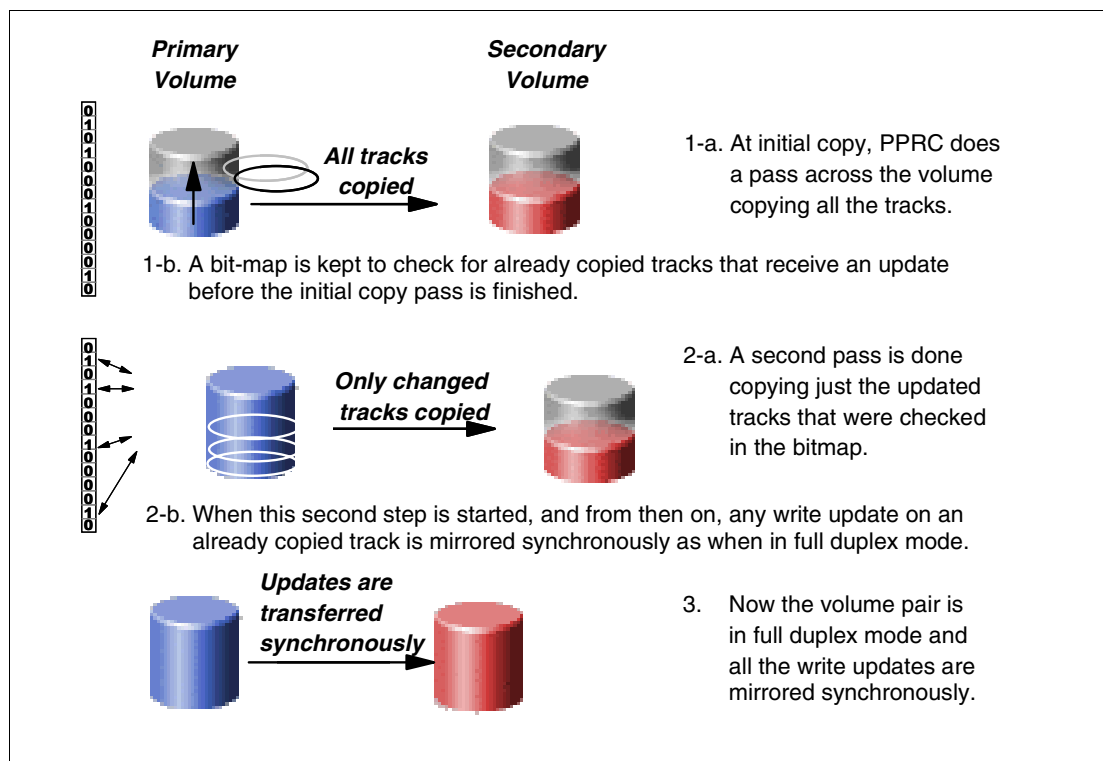
Figure 4-1   PPRC Initial volume synchronization

4.2.2 Synchronous PPRC continuous operation

The synchronous PPRC protocol guarantees data consistency by ensuring that the *write complete* is only received by the host application once the remote copy has been committed to the secondary ESS and acknowledged by both ESSs. This is illustrated in Figure 4-2 on page 56.
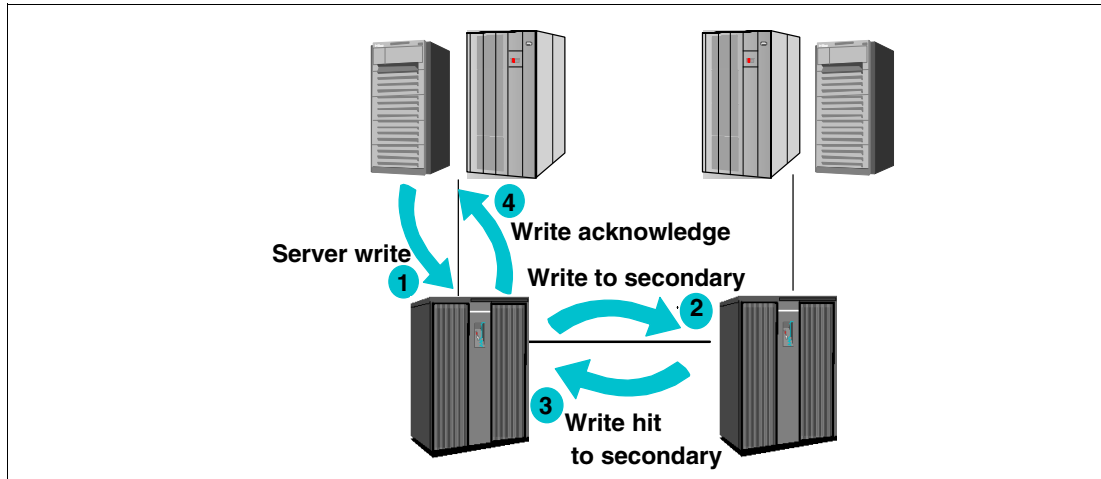
Figure 4-2   PPRC write I/O cycle

As shown in Figure 4-2:

1. The host server requests a write I/O to the primary ESS. The write is staged through cache and non-volatile storage (NVS).

2. PPRC dispatches the write over an ESCON channel to the secondary ESS. The write hits the secondary ESS's cache and NVS, which is sufficient to initiate an acknowledgement of the remote write to the primary ESS.

3. The primary ESS waits for acknowledgment of the remote write. If the secondary write fails, the acknowledgement does not return to the host server, causing an I/O timeout to the host server, which in turn causes a retry from the host server.

4. The write returns to the host server's application.

Once acknowledgement of the write has been received by the primary, both the primary and secondary write I/Os are eligible for destaging to disk. Destaging from the cache to the disk drives on both the primary and the secondary ESS is performed asynchronously.

If acknowledgement of the remote write is not received within a fixed period of time, the write is considered to have failed, and is rendered ineligible for destaging to disk. At this point, the application receives an I/O error, and in due course, the failed write I/O is *aged-out* of each NVS.

4.2.3 Synchronous PPRC volume states

Volumes within the Enterprise Storage Server used for synchronous PPRC can be found in one of the states illustrated in Figure 4-3 on page 57. (The icons are from the ESS Copy Services version 1 panels; for version 2 icons, see Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169.)
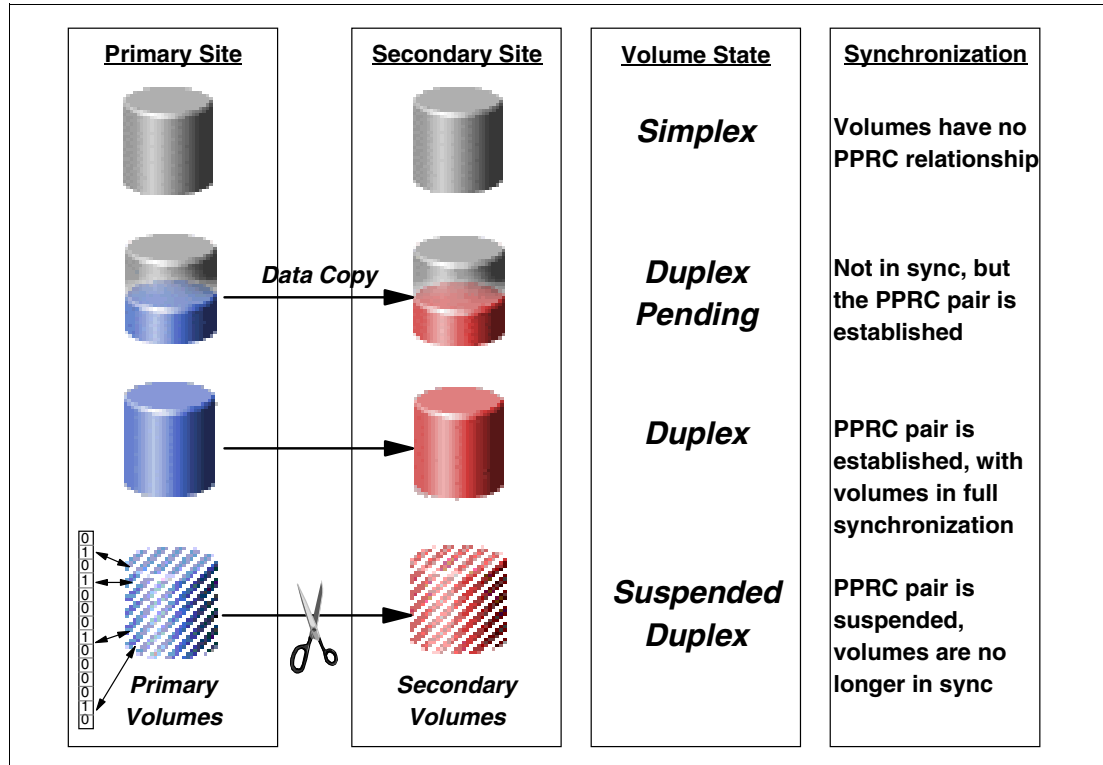
Figure 4-3   PPRC volume states

► **Simplex:** The simplex state is the initial state of the volumes before they are in a PPRC volume pair relationship, or after the PPRC relationship has been terminated. Both volumes are accessible only when in this state.

► **Duplex pending:** Volumes are in duplex pending state after the PPRC copy relationship is established, but the source and target volume are out of sync. In that case, data still needs to be copied from the source to the target volume of a PPRC pair. That may be the case either after the PPRC relationship was just established (or re-established after being suspended), or in the case where the PPRC volume pair re-establishes after a storage subsystem failure. The PPRC secondary volume is not accessible when the pair is in duplex pending state.

► **Duplex:** The duplex state is the state of a volume copy pair whose members are in sync; that is, both source and target volumes contain exactly the same data. This state is also referred to as *full copy mode*. The PPRC secondary volume is not accessible when the pair is in duplex state (if the **Permit read from secondary option** was not activated).

► **Suspended:** Volumes are in the suspended state when the source and target storage subsystems cannot communicate anymore, or when the PPRC pair is suspended manually. In this state, writes to the primary volume are not mirrored onto the secondary volume. The secondary volume becomes out of sync. During this time, PPRC keeps a bitmap record of the changed tracks in the primary volume. Later, when the volumes are re-synchronized, only the tracks that were updated will be copied.

4.2.4 Synchronous PPRC tasks and options

There are three types of ESS Copy Services tasks to establish, suspend, and terminate synchronous PPRC. Each of these tasks can be used with several options.

Establish synchronous PPRC copy pair: This operation starts the transition of two simplex volumes into a pair in the duplex pending or duplex state. You have to select one out of the four options below to specify the amount of data to be transferred during initialization:

- ▶ **Copy entire volume** copies all tracks from the primary volume to the secondary, overwriting the whole secondary volume. The volume pair changes from state simplex to duplex pending and finally to duplex after the copy is completed.

- ▶ **Do not copy volume** establishes the PPRC pair without any data transfer. The volume pair state goes to duplex immediately. This option presumes that the volumes are already synchronized. It allows you to restart a PPRC operation in the shortest time possible when you are sure that there were no data modifications.

- ▶ **Copy out-of-sync cylinders only** requires that the volumes are in the suspended state with a bitmap record of the changed tracks in the primary volume. The volume pair state changes from suspended to duplex pending and finally to duplex after the copy is completed. You choose this option to restart a PPRC operation when there are data modifications on the primary volume.

- ▶ **PPRC Failover** and **PPRC Failback** reverse the primary and secondary roles during a failover or failback process. The amount of data to be transferred depends on the volumes' state. We will discuss this in Chapter 10, "Disaster recovery" on page 399.

The following additional options can be selected:

- ▶ **Permit read from secondary** allows host servers to read from the PPRC secondary volume. The PPRC pair must be in a duplex state in order for the host server to read the volume. This option is helpful for operating systems with a Logical Volume Manager that reads data structures from the header of a secondary volume to update internal databases.

- ▶ **Suspend PPRC after establish complete** suspends the PPRC pair immediately after the establish-pairs command completes. This option is useful in combination with PPRC-XD (see 4.3, "PPRC extended distance (PPRC-XD)" on page 59).

- ▶ **Asynchronous Cascading PPRC** will be discussed in 4.4, "Asynchronous Cascading PPRC" on page 65.

Please note that the options **Critical volume mode** and **Permit establish if target is online** are not available for open-systems volumes.

Suspend PPRC copy pair: This operation stops transferring data to the target volume. Because the primary ESS keeps track of all changed cylinders on the source volume, you can resume PPRC operations at a later time. You have to select one of the two options:

- ▶ **Schedule task with source logical subsystem** is the default option. Both volumes will go to state suspended.

- ▶ **Schedule task with target logical subsystem** may be necessary when the primary LSS is not available. The primary volume might appear later still in state duplex, not creating a bit map record of changed tracks.

Terminate PPRC copy pair: This operation ends the relationship between the volumes. Similar to **Suspend PPRC copy pair**, you have to select one out of two options:

- ▶ **Schedule task with source logical subsystem** is the default option. Both volumes will go to state simplex.

► **Schedule task with target logical subsystem** may be necessary when the primary LSS is not available. The secondary volume will go to the simplex state, the primary volume will be in a suspended state.

Operations and options will be discussed again with respect to task creation in Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169 and Chapter 10, "Disaster recovery" on page 399.

4.3 PPRC extended distance (PPRC-XD)

In this mode of operation, PPRC copies data from the primary volume to the secondary volume in a *non-synchronous* manner. This means that I/O write completions are returned to the application once they have been committed to the primary ESS. Updates on the secondary volume are performed at a later point in time. The original order of updates is not strictly maintained.

The non-synchronous characteristics of PPRC-XD, combined with the throughput of the efficient track mirroring technique, and with enhanced microcode support for channel extenders, makes PPRC-XD well suited for remote copy solutions at distances beyond 103 km, without having to incur distance latency penalties that synchronous write I/O cycles exhibit. This behavior means that distance does not directly impact the performance of the application.

4.3.1 PPRC-XD operation and volume states

When operating in PPRC-XD, the volume pairs are in the *duplex pending XD* state. This state is in addition to those already described in 4.2.3, "Synchronous PPRC volume states" on page 56. (The icons in Figure 4-4 are from the ESS Copy Services version 1 panels; for version 2 icons, see Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169.)
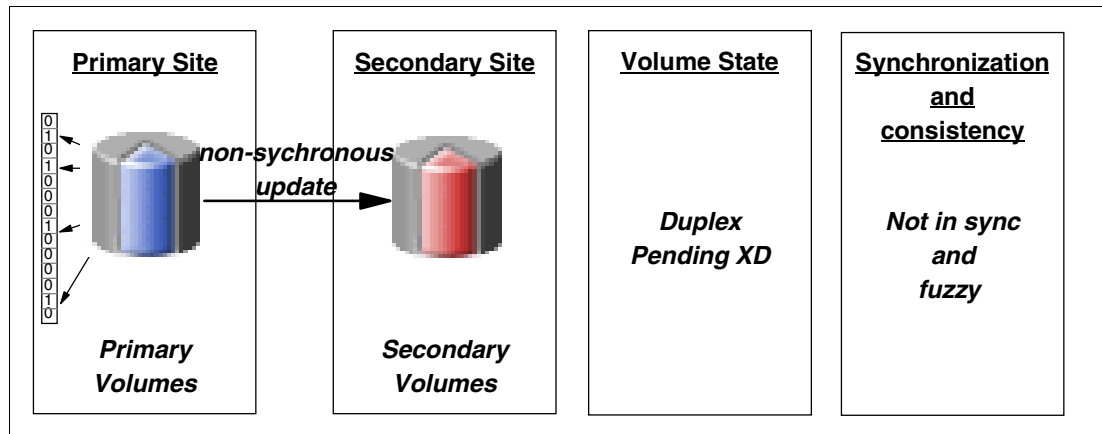
Figure 4-4   PPRC Extended Distance duplex pending XD volume state

While the volume pairs are in the duplex pending XD state, PPRC is doing a non-synchronous mirroring of the primary volumes' updates on to the secondary volumes. PPRC-XD maintains a bitmap of the changed tracks for all the primary volumes in the duplex pending XD state. It periodically cycles through each volume's bitmap for updated tracks and then sends the updates in a batch of 30 tracks for transmission to their secondary counterparts. Once the batch has been committed to the cache/NVS of the secondary ESS and the original track bit has not been set again, the bit is cleared to avoid unnecessary

retransmissions. This is a very efficient throughput-oriented method of non-synchronous mirroring.

The efficient extended distance technique of PPRC Extended Distance is achieved using sophisticated algorithms. For example, if changed data is in the cache, then PPRC sends only the changed sectors. There are also sophisticated queueing algorithms to schedule each volume for processing of their updated tracks. These PPRC-XD operational details are beyond the scope of this document.

The volume pair will remain in the duplex pending XD condition, not reaching the duplex state until instructed to do so. The operation to exit the duplex pending XD state will direct the volume pair either to go to the duplex state (a *go-to-sync* operation), or to the suspended state (a suspend pair operation), or to the simplex state (a terminate pair operation). The permissible volume state transitions are illustrated in Figure 4-5. (The icons are from the ESS Copy Services version 1 panels; for version 2 icons, see 6.5.1, "Volume panel icons" on page 181.)
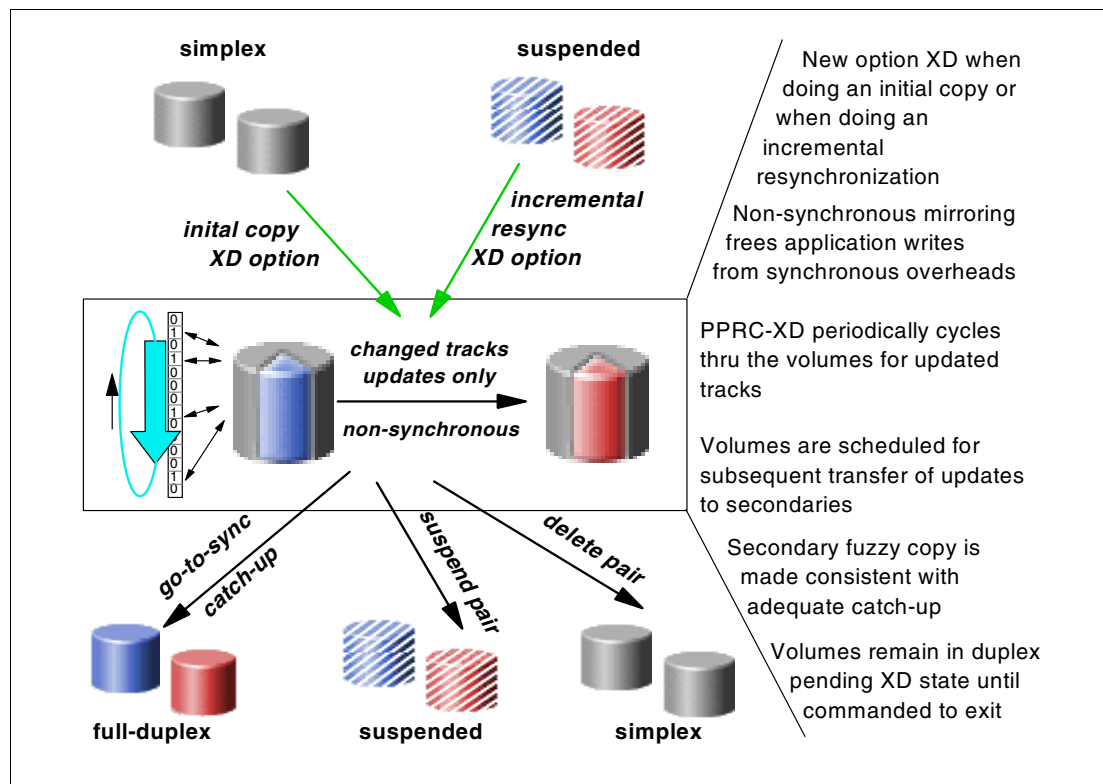
Figure 4-5   PPRC-XD volume state transitions

4.3.2 PPRC-XD tasks and options

Similar to synchronous PPRC, there are three types of ESS Copy Services tasks to establish, suspend, and terminate PPRC-XD with several options for each tasks.

Establish PPRC Extended Distance copy pair: This operation starts the transition of two simplex volumes into a pair in a duplex pending or duplex state. You have to select one of the four options below to specify the amount of data to be transferred during initialization:

▶ **Copy entire volume** copies all tracks from the primary volume to the secondary, overwriting the whole secondary volume.

- **Do not copy volume** establishes the PPRC pair without any data transfer. This option presumes that the volumes are already synchronized. It allows you to restart a PPRC operation in the shortest time when you are sure that there were not any data modifications. (This option is not available for PPRC-XD with ESS Copy Services V1.)

- **Copy out-of-sync cylinders only** requires that the volumes are in suspended state with a bitmap record of the changed tracks in the primary volume. You choose this option to restart a PPRC operation when there were data modifications on the primary volume.

- **PPRC Failback** reverse the primary and secondary roles during a failback process. The amount of data to be transferred depends on the volumes' state. We will discuss this subject in Chapter 10, "Disaster recovery" on page 399.

The option **PPRC Failover** is available only for synchronous PPRC. The following additional options can be selected:

- **Permit read from secondary** allows host servers to read from the PPRC-XD secondary volume. Unlike a synchronous PPRC pair, the pair is not required to be in a duplex state for the host server to read the secondary volume.

- **Asynchronous Cascading PPRC** will be discussed in 4.4, "Asynchronous Cascading PPRC" on page 65.

Please note that the options **Critical volume mode** and **Permit establish if target is online** are not available for open-systems volumes. The option **Suspend PPRC after establish complete** is available only with synchronous PPRC.

Suspend PPRC pair: This operation stops transferring data to the target volume. Because the primary ESS keeps track of all changed cylinders on the source volume, you can resume PPRC operations at a later time. You have to select one of the two options:

- **Schedule task with source logical subsystem** is the default option. Both volumes will go to suspended state.

- **Schedule task with target logical subsystem** may be necessary when the primary LSS is not available. The primary volume might appear later still in duplex state, not creating a bit map record of changed tracks.

Terminate PPRC pair: This operation ends the relationship between the volumes. Similar to **Suspend PPRC pair**, you have to select one of the two options:

- **Schedule task with source logical subsystem** is the default option. Both volumes will go to state simplex.

- **Schedule task with target logical subsystem** may be necessary when the primary LSS is not available. The secondary volume will go to the simplex state, and the primary volume will be in the suspended state.

Operations and options will be discussed again with respect to task creation in Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169 and Chapter 10, "Disaster recovery" on page 399.

4.3.3 Creating a consistent point-in-time copy

While the copy pair volumes are in the duplex pending XD state, the secondary volumes maintain a *fuzzy copy* of the data:

- Because of the asynchronous characteristics, at any time there will be a certain amount of updated data that is not reflected at the secondary volume. This data corresponds to the tracks that were updated since the last volume bitmap scan was done. These are the out-of-sync tracks.

► Because of the bitmap scan method, writes are not ensured to be applied on to the secondary volume in the same sequence as they are written to the primary volume.

When terminating the PPRC relationship to establish host access to secondary volumes, the first issue may cause loss of transactions. Since a file system's or database's consistency depends on the correct ordering of write sequences, the second issue may cause inconsistent volumes. Therefore, to use secondary volumes by the host systems, you need to make them *point-in-time (PIT) consistent*:

► The secondary volumes have to *catch up* to their primary counterparts. PPRC-XD catch-up is the name of the transition that occurs to a PPRC-XD pair when it goes from its normal out-of-sync condition until it reaches a full sync condition. At the end of this transition, the primary and secondary volumes become fully synchronized.

► Then the application needs to be quiesced and the PPRC volume pairs temporarily suspended. This is necessary to ensure consistency not only at the volume level, but also at the application level.

► A FlashCopy of the secondary volumes onto tertiary volumes should now be performed, followed by resuming of the PPRC-XD pairs.

These tertiary volumes are then a consistent point-in-time copy of the primary volumes. Figure 4-6 on page 63 summarizes this procedure that provides a consistent point-in-time (PIT) copy of the data:

1. Quiesce the application updates.

2. Synchronize the volume pairs by one of these methods:

   – Perform the catch-up by issuing a *go-to-sync operation*. This is done by the PPRC task **Establish synchronous PPRC copy pair** with the option **Copy out-of-sync cylinders only** to minimize the time needed for data transfer.

   Now the volume pair leaves the duplex pending XD state and will reach the duplex state (refer back to Figure 4-5 on page 60). From this moment, if the pairs were not immediately suspended, primary write updates would be synchronously transmitted to the recovery site.

   You may not want any synchronous copy operations to occur once the volumes are in duplex state, especially if the volumes you are mirroring are separated by long distances (beyond 103 km). For this reason, you may select the additional option **Suspend PPRC after establish complete** when creating the **Establish synchronous PPRC copy pair** task.

   – Perform the catch-up by waiting until all application updates are transferred to the secondary site.

   You can monitor the number of out-of-sync tracks in the volume information panel or by issuing `rsQuery` commands (see Chapter 7, "ESS Command Line Interface" on page 301). If you triggered the catch-up by issuing go-to-sync, then you can also watch for completion of the synchronous PPRC task using the `rsQueryComplete` command.

3. Suspend the PPRC copy pairs after they reach the full duplex state. (With a go-to-sync operation, you can select the additional task option **Suspend PPRC after establish complete** to avoid this step.) If you use PPRC consistency groups (will be discussed in 4.5, "PPRC data consistency" on page 72), issue a **Freeze** operation.

   At this point, we have a set of consistent secondary volumes.

4. You can resume the application (updates will not be transferred to the secondary volumes because the pairs are suspended).

5. Perform a FlashCopy on the secondary volumes.

6. Resume PPRC Extended Distance mode for the copy pair.
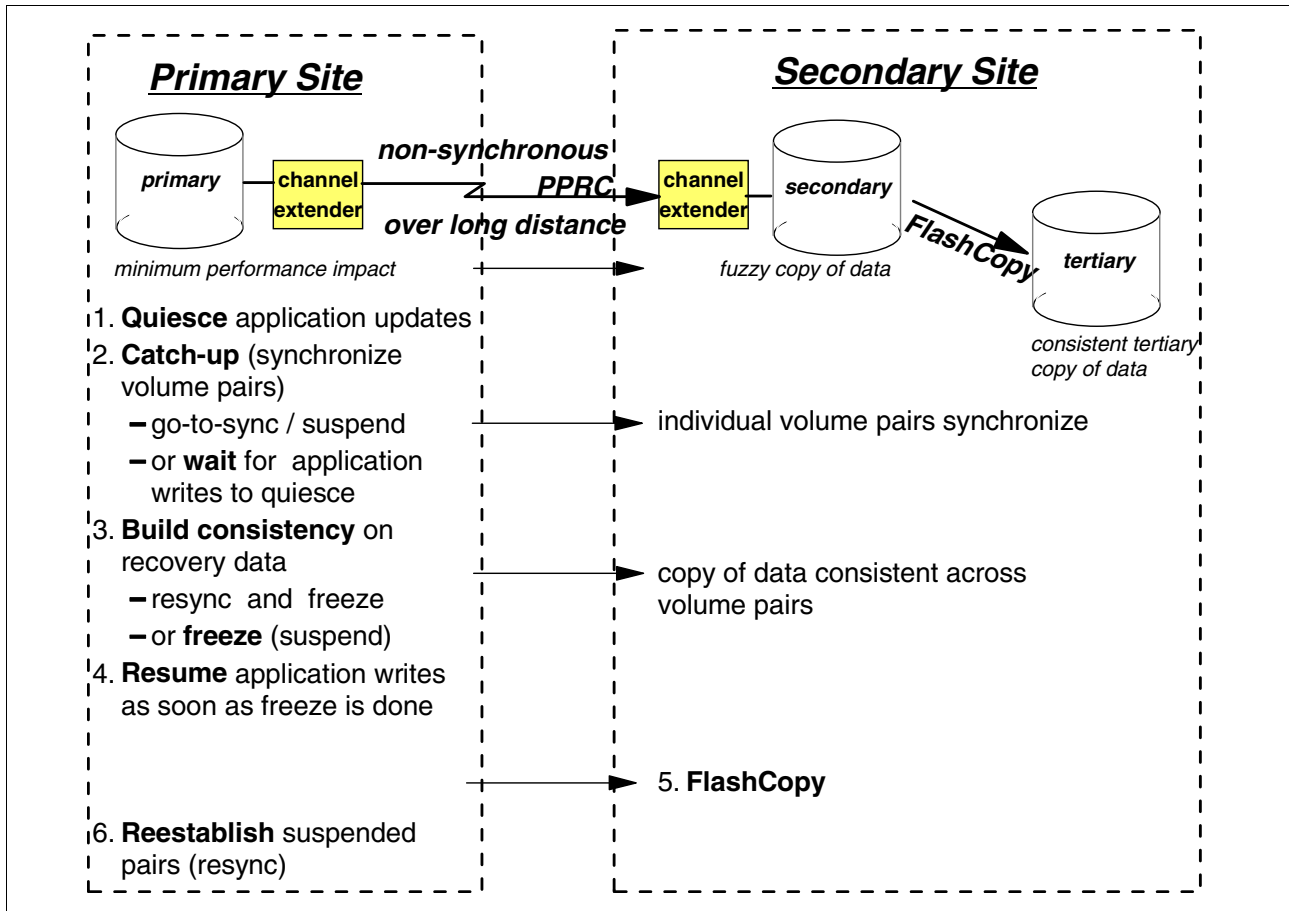
Figure 4-6   Create a PPRC-XD consistent copy

For application recovery based on point-in-time copies, you have to plan for appropriate checkpoints to briefly quiesce the application and synchronize the volumes pairs. When the recovery of the application is done, you must remember that, while in an active PPRC-XD relationship, the secondary volumes always have a current fuzzy copy of the primary volumes. So you have to keep the tertiary volumes where you "FlashCopied" the last globally consistent catch-up. As you can realize, this tertiary copy will not be reflective of the current update, but rather, any updates up until the last global catch-up operation.

4.3.4 Synchronous PPRC and PPRC-XD comparison

Synchronous PPRC and PPRC-XD differ by several characteristics. Some of the more important ones are data consistency and currency of data at the secondary site, impact of latency for applications, and bandwidth requirements. Table 4-2 on page 64 summarizes these differences.

Table 4-2   Comparison of synchronous PPRC and PPRC-XD characteristics

Synchronous PPRC	PPRC Extended Distance
If write W_1 completes at primary site before W_2 is initiated, then secondary site never has modifications from W_2 without also having W_1.	No consistency guarantee; point-in-time consistency with quiesce/resume procedure.
Data at the secondary site is up-to-date when write has been acknowledged to host.	Data transfer delayed (indefinite amount of time) after write has been acknowledged to host.
Latency impact for applications increases proportionally with distance between sites.	Latency impact is independent of the distance between sites.
Bandwidth must be sufficient to carry the instantaneous peak write load.	Bandwidth must be greater than the average write bandwidth.

Synchronous PPRC provides a consistent copy, which means that the secondary volume is as good as possible a complete match of the primary volume. Dependent writes are applied on the secondary volumes in the same sequence as they are applied on the primary volumes. This makes synchronous PPRC ideal for rolling disaster recovery scenarios. PPRC-XD is an asynchronous data mover that provides a fuzzy copy of the data on the secondary volume. As a consequence, periodic go-to-sync operations to build consistent PIT copies at the recovery site are necessary. Due to this, recovery is possible to the last PIT copy.

Because PPRC-XD does not delay the acknowledgement of a write complete until updates have been made to the secondary volume, PPRC-XD preserves the application write performance better than synchronous PPRC over longer distances. With PPRC Extended Distance, the ESSs hosting the primary and secondary volumes can be separated by distances well beyond the 103 km supported with PPRC synchronous transmissions (continental distances) and are only limited by the channel extender and network technology capabilities. As synchronous mirroring cannot occur over long distances without a serious application performance impact, the best trade-off is to accept the periodical brief application quiesces, in order to build consistent PIT copies of the data. Another way may be to combine synchronous PPRC and PPRC-XD in a cascading relationship (see 4.4, "Asynchronous Cascading PPRC" on page 65).

Application write performance under synchronous PPRC is dependent on the available bandwidth, so only critical data should be copied. PPRC-XD allows you to better exploit your available bandwidth capacity, therefore allowing you to include more of your data to be protected.

As summarized in Figure 4-7 on page 65, PPRC Extended Distance is a recommended solution for data copy, data migration, off-site backup, and transmission of inactive database logs with excellent performance, which is particularly relevant when implemented over continental distances. PPRC Extended Distance can also be used for application recovery solutions based on periodic point-in-time backup copies of the data if the application tolerates periodic brief interruptions (application quiesce).

PPRC Extended Distance Positioning

> **PPRC Extended Distance is a recommended solution for data copy, data migration and offsite backup — over continental distances — with excellent application performance.**
>
> **It can be used for disaster recovery implementations if application can be quiesced and non-zero data loss Recovery Point Objective (RPO) is admissible.**

- It can be used over continental distances with excellent application performance.
  - The distances are only limited by the network and channel extenders capabilities.
- Applications' write operations do not have distance-dependent latency impact.
- Fuzzy copy of data at the recovery site (sequence of dependent writes may not be respected at the recovery site).
- Recovery data can become a consistent point in time mirror of the primary data, if appropiate application checkpoints are set to do global catch-ups.
  - Pairs are synchronized with application group consistency.
- Synchronizations can be done more frequently, because of short catch-ups
  - RPO still non-zero, but improves substantially.

Figure 4-7   PPRC Extended Distance positioning

For an in-depth discussion of PPRC characteristics, see "Advanced functions for storage subsystems: Supporting continuous availability" by Azagury, et al, in *IBM Systems Journal*, Vol. 42 (2003) No. 2 (June), found at:

http://www.research.ibm.com/journal/sj42-2.html

4.4 Asynchronous Cascading PPRC

A feature introduced with ESS Copy Services V2 called Asynchronous Cascading PPRC allows you to combine synchronous PPRC and PPRC-XD, providing another method for disaster recovery.

4.4.1 Asynchronous Cascading PPRC operation and volume states

With Asynchronous Cascading PPRC, a volume is allowed to be simultaneously both a PPRC primary and a PPRC secondary, combining three volumes in two PPRC pairs. Figure 4-8 on page 66 shows the principles of this configuration.
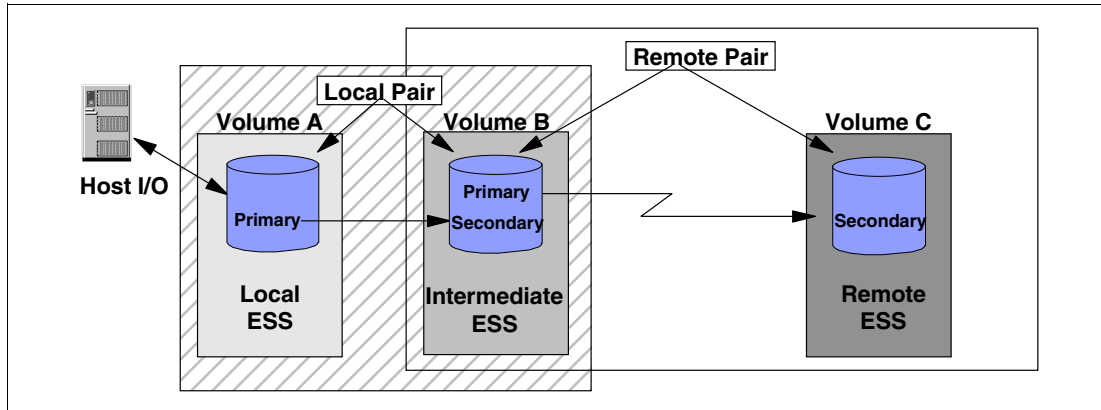
Figure 4-8   Volume pairs in an Asynchronous Cascading PPRC relationship

In Figure 4-8, we have two PPRC volume pairs: A–B and B–C. Volume B acts in two roles simultaneously: It is the secondary volume of the PPRC relationship A–B and the primary volume of the PPRC relationship B–C. As with any PPRC relationship, volumes A and B, as well as volumes B and C, must be in a different LSS (in the same or different ESS). There is a minimum of two different ESSs and a maximum of three different ESSs in an Asynchronous Cascading PPRC relationship. Figure 4-9 shows the volume panel of an ESS with LSS 16 as local and LSS 17 as the intermediate LSS. Volume 700-18767 has two status icons: the red as secondary of the local synchronous PPRC relationship, the blue as primary of the remote PPRC-XD relationship.
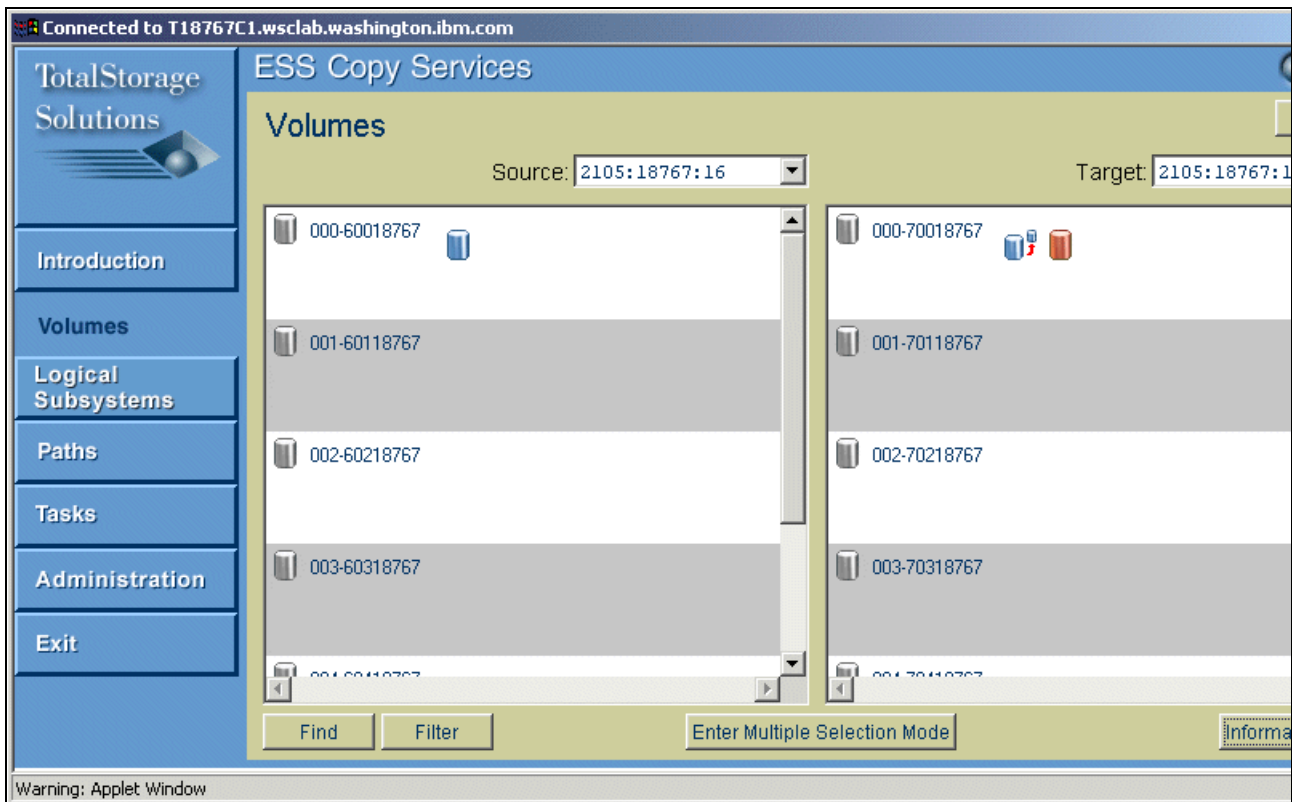
Figure 4-9   Asynchronous Cascading PPRC intermediate volume states

In practice, the three volumes will be often in different sites, called the *local*, the *intermediate*, and the *remote site*. Accordingly, we denote the pair A–B as the *local pair* and B–C as the

remote pair. When considering the whole chain of data transfer, volume A can be considered as the primary, volume B as the secondary, and volume C as a *tertiary* volume.

The state of both PPRC pairs is maintained independently. You can suspend or resume A–B without changing the state of B–C and vice versa. It is important to notice that volume B has two PPRC states, depending on which relationship is considered. Figure 4-10 shows the volume information panel of the intermediate volume from the example above, indicating that this volume acts as a PPRC primary as well as a secondary.
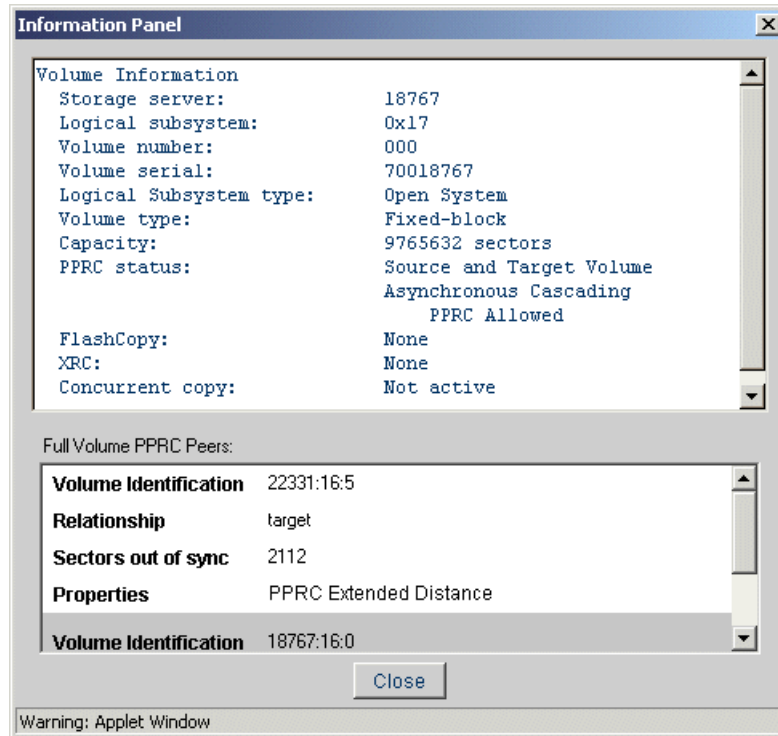
Figure 4-10   Asynchronous Cascading PPRC intermediate volume information

> **Note:** There is no additional volume involved in the intermediate LSS. With respect to the total costs of ownership, this saves disk space and license fees compared with solutions based on an additional FlashCopy volume in the intermediate site.

4.4.2 Asynchronous Cascading PPRC tasks and options

To maintain Asynchronous Cascading PPRC relationships, you apply the same types of ESS Copy Services tasks to establish, suspend, and terminate PPRC for each of the two pairs separately. The available options depend on the mode of operation you use for the specific volume pair (see 4.2.4, "Synchronous PPRC tasks and options" on page 58 and 4.3.2, "PPRC-XD tasks and options" on page 60). The only difference is that you must select the option **Asynchronous Cascading PPRC** when establishing the B–C pair:

► **Asynchronous Cascading PPRC** allows the PPRC primary volume of that relationship to be also the secondary volume in another PPRC relationship.

The cascading pairs may be established in any order. We recommended the following procedure:

1. Establish the remote pair with the option **Do not copy volume**.
2. Establish the local pair with the option **Copy entire volume**.

This avoids the unnecessary transfer of invalid data from the secondary to the tertiary volume. If you alternatively establish the local pair first with **Copy entire volume** option, then you have to specify **Copy entire volume** also for the remote pair. A cascading relationship may be added to existing pairs too.

4.4.3 Asynchronous Cascading PPRC mode combinations

Table 4-3 shows the supported combinations of PPRC operation modes in a cascading relationship.

Table 4-3   Asynchronous Cascading PPRC mode combinations

Local volume pair	Remote volume pair
Synchronous PPRC	PPRC Extended Distance
	Synchronous PPRC
PPRC Extended Distance	PPRC Extended Distance

The combination PPRC Extended Distance to synchronous PPRC is not supported because the intermediate volume as a fuzzy copy cannot be the source for a consistent copy at the tertiary volume.

Synchronous PPRC to PPRC-XD is the typical usage intended for disaster recovery. The combination PPRC-XD to PPRC-XD allows you to maintain multiple copies of data without significant impact to application I/O performance. In contrast to that, a combination of two synchronous PPRC relationships will have a severe impact to application I/O because each write operation has to wait for two PPRC transfers. This combination may be useful to guarantee a consistent copy at the remote site if there is little write I/O and the whole distance is beyond the maximum distance of 103 km supported for a single synchronous PPRC.

Although our terminology for local, intermediate, and remote ESSs may suggest that Asynchronous Cascading PPRC requires three ESSs, there are other variants possible. The only one requirement is that the two volumes forming a basic relationship A–B or B–C must be in different Logical Subsystems. These LSSs may be in the same or in different ESSs. So there are variants of Asynchronous Cascading PPRC with three LSSs in the same, in one or two ESS(s). From a practical point of view, only two variants make sense:

► Three ESSs with one volume per ESS, as shown in Figure 4-8 on page 66.

   When the local pair runs in synchronous mode, you should keep the distance between the local and intermediate ESSs as close as possible to minimize the performance impact. The intermediate ESS should be in a secure environment separated from the local ESS and with separate power to reduce the possibility of an outage affecting both locations.

   The remote ESS is typically in a different city, any distance away from the intermediate ESS. If the remote PPRC pair operates asynchronously from the local PPRC pair, then the effect of long distance on host response time is minimized.

► Two ESSs with volumes A and B in the same local ESS and volume C in the remote ESS, as shown in Figure 4-11 on page 69.

   If the remote pair operates in PPRC-XD mode, this configuration is very similar to a single PPRC-XD relationship. Having the intermediate volume leads to procedures for creating a consistent remote point-in-time copy with less application impact.
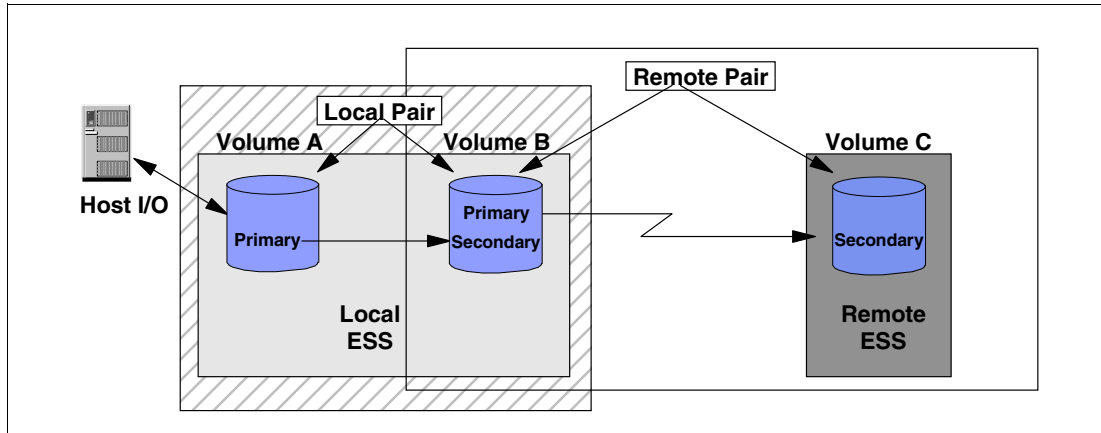
Figure 4-11   Asynchronous Cascading PPRC with two sites

4.4.4 Creating a consistent copy with PPRC-XD as a remote pair

The combination of synchronous PPRC and PPRC-XD brings up again the question of how to create a consistent copy at the remote site. As discussed in 4.3.3, "Creating a consistent point-in-time copy" on page 61, the secondary volume of a PPRC-XD relationship is not consistent and generally not useful for recovery unless all of the data has been copied up to some point in time. PPRC-XD secondary volumes must periodically be made consistent and an additional copy of the data kept.

We describe how to create a consistent backup copy of volume A at the remote site under the following assumptions:

► The data on volumes A and B is identical (volume state is full-duplex).
► Volumes A and B were established in PPRC synchronous mode.
► Volumes B and C were established in PPRC Extended Distance mode.

Under these assumptions, the following procedure (summarized in Figure 4-12 on page 71) creates a consistent point-in-time copy at the remote site:

1. At the local ESS:

   Issue a **Freeze** operation (see "Freeze operation" on page 78) to the A–B volume pair. This ensures that no further updates occur to volume B and updates are allowed to complete on volume C. Volumes A and B (the synchronous PPRC pair) are now suspended in a consistent state.

2. At the intermediate ESS (in a two-ESS configuration: the local ESS):

   a. Convert volumes B and C to synchronous mode to allow pending tracks and updates to be transmitted to the remote site.

   b. Monitor the copy process of volumes B and C for the transition to full duplex state.

   When volumes B and C are in full duplex state, you can create a copy at the remote ESS that is consistent to the time of the freeze.

> **Note:** When using Fibre Channel links between site A and site B, if the synchronous PPRC is active with application I/O still running at site A, and if a go-to-SYNC is triggered for the XD pair, then the XD pair will automatically suspend.
>
> We recommended that when go-to-SYNC is issued for the XD pair, the synchronous PPRC pairs should be suspended (this will normally be done via a freeze operation, as described in our example).

3. At the remote ESS:

   While volumes A and B are suspended and after the synchronization process of volumes B and C is complete, use FlashCopy to create a PIT-consistent copy using volumes C and D.

4. At the intermediate ESS (in a two-site configuration: the local ESS):

   a. Suspend volumes B and C.

   b. With volumes B and C in a suspended state, re-establish volumes B and C as a PPRC-XD pair. Select the **Copy out-of-sync cylinders only** option.

5. At the local ESS:

   a. Resume operations after the freeze. When you issue the freeze to volume A, the established paths between the LSSs were disabled. You have to re-establish the paths.

   b. Resynchronize the suspended A–B volume pair so that all the changed data is copied from volume A to B. Select the **Copy out-of-sync cylinders only** option.
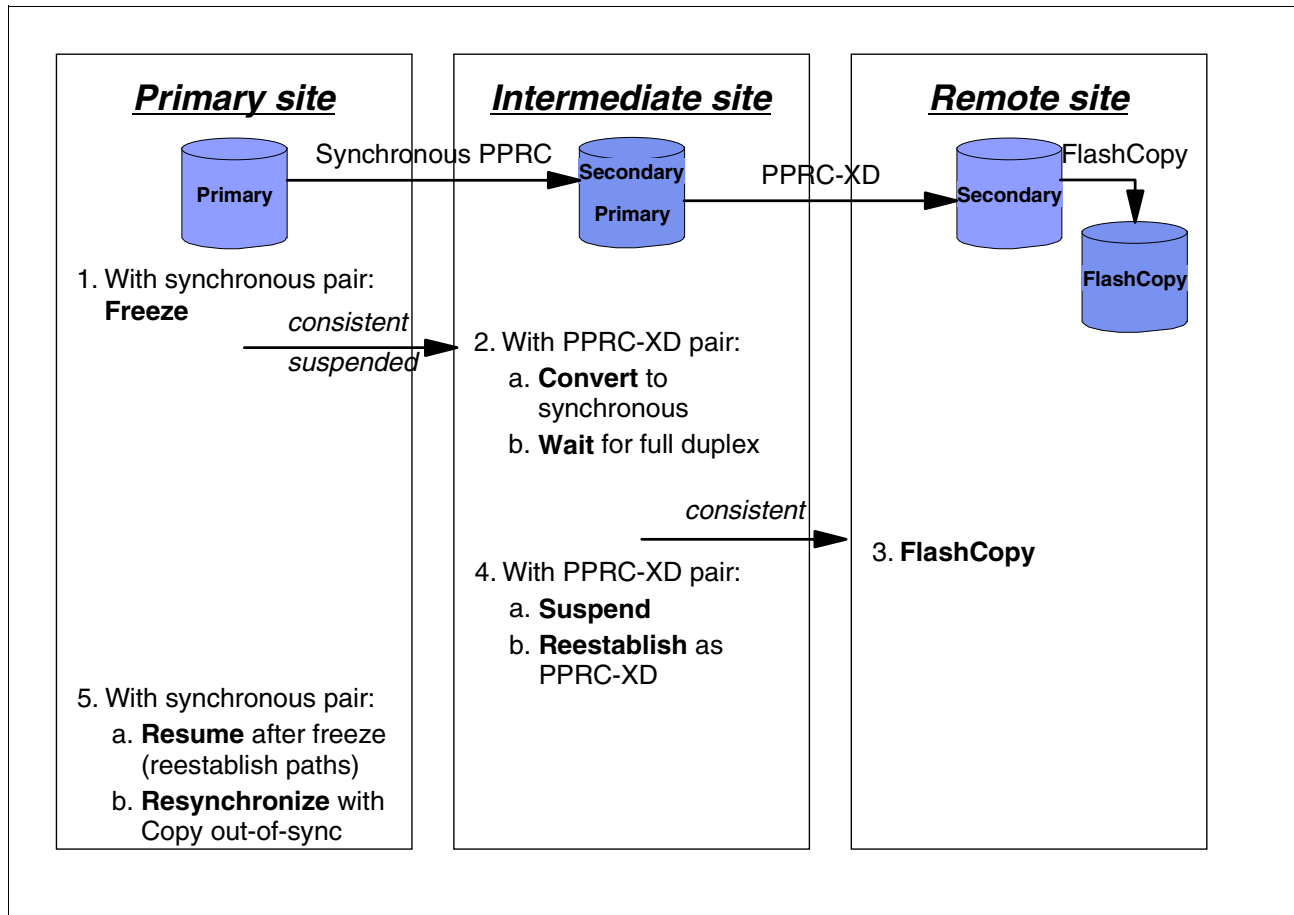
Figure 4-12   Create consistent copy with Asynchronous Cascading PPRC

Reviewing this procedure, you should notice the following details:

► Synchronous PPRC Consistency Groups need to be implemented in the standard way for the synchronous PPRC relation (step 1 of the procedure).

► Once the PPRC Freeze operation happens, the intermediate site volumes become static and consistent; therefore, as soon as the bandwidth permits, the remote site will catch up also to this static and consistent point. So step 2 is not really an action, it is only a process to watch for.

► The FlashCopy at the remote ESS (step 3) can be issued as Inband FlashCopy.

4.4.5 Outage scenarios with Asynchronous Cascading PPRC

The outage scenarios with Asynchronous Cascading PPRC depend on the configuration with two or three ESSs:

► Asynchronous Cascading PPRC with local and intermediate volumes in the same ESS leads to the same scenarios as a single PPRC configuration. After losing access to the local ESS, the remote ESS provides a fuzzy copy if the remote pair was in a PPRC-XD relationship. (Operating both PPRC relations in synchronous mode does not make sense in a two-ESS configuration.)

Recovery has to be made with the last consistent point-in-time copy in the remote ESS.

Therefore, when we compare configurations of two ESSs running Asynchronous Cascading PPRC or a single PPRC, we do not see increased availability or recovery. But

there is less impact to the applications with Asynchronous Cascading PPRC when creating a consistent point-in-time copy at the remote site (compared with single PPRC-XD).

► Asynchronous Cascading PPRC with three ESSs offers more options:

– If there is an outage of the local ESS, the intermediate ESS contains a consistent and up-to-date copy of the data. You will fail over production to the intermediate ESS without loss of transactions.

Shortly (depending on bandwidth), the intermediate ESS will complete sending all updates to the remote ESS. In this way, the remote site will automatically catch-up.

– If there is an outage of the intermediate ESS, data at the local ESS is not affected. Applications continue to run as normal.

There is only one copy of data available until the intermediate ESS has been repaired.

– If there is an outage at the remote ESS, data at the local and intermediate ESS is not affected. Applications continue to run as normal. The intermediate ESS maintains a consistent up-to-date copy.

You may ask the question what happens if you lose both the local and the intermediate ESS of a three-ESS Asynchronous Cascading PPRC configuration. The answer is that you go back to the most recent point-in-time copy that was made at the remote ESS. In other words, the situation is the same as losing the primary ESS in a two-ESS PPRC-XD configuration. If the outage is a disaster large enough to cause the loss of both the primary and the intermediate ESS, then there are probably problems affecting the whole region and the business. In such situations, the larger time needed to recover the data will not be the limiting factor for business contingency.

4.4.6 PPRC Failover and Failback with Asynchronous Cascading PPRC

Asynchronous Cascading PPRC has been enhanced to exploit PPRC Failover and Failback modes in a three-site solution. This can help reduce the time required to synchronize PPRC volumes after switching between sites during planned or unplanned outages.

This PPRC enhancement is supported on ESS Models 800, F20, and F10, with LIC level 2.3.0 or above, and is provided with PPRC Version 2 (PPRC Version 2 is an optional feature to the ESS).

PPRC Failover and Failback with Asynchronous Cascading PPRC is discussed in detail in 10.2, "Data consistency" on page 402.

4.5 PPRC data consistency

As mentioned in 4.3.3, "Creating a consistent point-in-time copy" on page 61, data consistency needs special attention with copy operations. Here we discuss consistency types and describe additional ESS features to control consistency.

4.5.1 Consistency types

With respect to operational context and scope, we have to distinguish several types of consistency. In the context of PPRC operations, there are three types of consistency guarantees to consider:

► **No guarantee**: This is the weakest type. We are ensured only that given sufficient time without updates to the primary volumes, the secondary volumes will eventually be equivalent to the primaries.

► **Power-failure consistency**: This guarantees that the secondary volumes are equivalent to what would have be seen at the primary site had there been a power failure and recovery.

► **Application consistency**: This ensures that the secondary volumes are equivalent to what would have been seen at the primary site had the application been given a chance to shut down normally.

The first type is our case with PPRC-XD: We know only that all updates will be transferred, but we cannot make assumptions about the ordering of dependent writes or update completion time.

In the description of power-failure consistency, you can include other types of hardware failures like loss of fibre connections. This is the case with synchronous PPRC. As noted in 4.3.4, "Synchronous PPRC and PPRC-XD comparison" on page 63, this ensures the ordering of dependent writes: If the write operation W_1 completes at the primary site before the write operation W_2 is initiated, then the secondary site never has modifications from W_2 without also having all modifications from W_1.

Typically, software recovery is required when secondary volumes with this type of consistency are used. Databases are able to handle that automatically as part of their startup routine using log files. Because such log files record successful and committed transactions (any partially committed transactions will be rolled back), this gives transaction-level consistency. File systems may behave differently, so it is important to perform a file system check on all PPRC secondary volumes before you start using them. Of course, rebooting the hosts will achieve the same result.

With power-failure consistency, you still may be exposed to so-called *lost writes*. These are the in-flight transactions that have not been committed. You should expect that uncommitted transactions will be lost. On the other hand, data that was transferred to the ESS and confirmed back as written into the NVS (of the secondary ESS in the case of PPRC) will be destaged to disk.

Application consistency gives the shortest recovery time. But to get an application consistent copy, you have to quiesce the application's updates and create a FlashCopy in a catch-up operation. As discussed in 4.3.4, "Synchronous PPRC and PPRC-XD comparison" on page 63, this is a point-in time copy that will not be reflective of the current updates, but rather, any updates up until the last catch-up operation.

With respect to scope, we have to distinguish between volume consistency and global consistency:

► **Volume consistency**: This means that some type of consistency is guaranteed for the specific secondary volume.

► **Global consistency**: This means that some type of consistency is guaranteed across all the volumes upon which the application does updates.

The scope of consistency needed will depend on the application write operations characteristics, as well as on how the data files are distributed over the production site volumes.

Synchronous PPRC ensures power-failure global consistency only as long as all volume pairs are in full duplex state. When error conditions affect some of the volume pairs (or different volume pairs at different time), global consistency may be lost. For example, if one of the secondary volumes cannot be updated because of path failures, then the corresponding primary volume will normally go into suspended state, still allowing updates. But these updates are no longer transferred to the secondary volume. Only the bitmap of changed tracks is created and maintained. So global consistency across volumes is lost, although the order of writes is still guaranteed for the other secondary volumes.

Consistency groups, as explained below, can help control global consistency.

4.5.2 PPRC consistency groups

To maintain consistency of data across volumes at the recovery site, volumes in a PPRC relationship can be collected into a *PPRC consistency group*. A PPRC consistency group is a set of volume pairs that have the same primary and secondary LSS for which the **PPRC consistency group** option has been activated. There are two possible ways to activate this option:

► During the establishment of the path from the primary to the secondary LSS (see 6.12, "PPRC setup examples using the Web user interface" on page 232). You select the check box **PPRC consistency group** in the path options panel.

  If the option is activated in that way, then all volume pairs in a PPRC relationship from this primary to this secondary LSS (sharing the same paths) belong to a consistency group. Other volumes pairs are not affected.

► At the LSS level. You select the check box **PPRC consistency group** in the LSS properties panel.

  If the option is activated in that way, then all volume pairs in a PPRC relationship from this primary to any secondary LSS (using any path) are affected.

This attribute changes the behavior of these volume pairs when an error occurs that affects any of these volumes:

► Without the **PPRC consistency group** option, the ESS causes the volume where the error is detected to enter a suspended state, but still allows updates to that volume.

► If the **PPRC consistency group** option was activated, the volume becomes suspended and additionally enters a long busy state. Updates are not possible in that state.

The *long busy state* is defined for System/390 as a unit check with a special code that tells the host error recovery program to stop the current I/O operation and requeue it until the control unit reconnects or two minutes have passed. Then the I/O will be redriven. If the control unit is still busy, the unit check is repeated.

For open systems (attached via SCSI or FCP), this state is presented as a *QUEUE FULL (QF)* status byte code to the host. A status byte is sent during the STATUS phase at the completion of each SCSI command (unless the command is terminated by a special event). The QF status signals to the host that the tagged SCSI command queue is full and that the I/O request has not been placed in the queue. Then the host system will attempt to requeue requests in the logical unit's device queue. So consistency grouping of volumes provides the ability to temporarily queue (at the host's level) subsequent write operations to all PPRC consistency group volumes on a single LSS pairing when an error occurs to one of the

volumes in the group (primary or secondary), or when a total link failure is detected between the primary and secondary LSS pair (see Figure 4-13 on page 76).

Details of QF handling may be modified by the operating system's and the host bus adapter's settings. For example, in some UNIX-type systems, the utility `scsi_ifsetcap` provides options to enable or disable SCSI command retry when a QF status is returned. Some HBA tools allow similar modifications. For example, the Emulex FC Port Driver Utility for Windows has a setting to translate QF to BUSY to work around a behavior of the disk class driver. You should check that the settings on your host systems are according to the *IBM TotalStorage Enterprise Storage Server Host System Attachment Guide,* SC26-7446.

> **Note:** With PPRC-XD, there is no need to enable the **PPRC consistency group** option. PPRC-XD always keeps a fuzzy copy at the recovery site. PPRC-XD pairs become consistent only when all updates are transmitted to the secondary by catch-up while the application writes are quiesced. When building a consistent point-in-time copy on the PPRC-XD secondaries, you will need to have the application writes already quiesced, which makes the QF time-out window unnecessary.

When the first copy pair in the consistency group becomes suspended and enters QF, the ESS error recovery program (ERP) issues an SNMP notification (trap 200). When subsequent pairs in the consistency group become suspended (during a total path failure), the information is recorded in a table. As other pairs suspend, they are added to the table. Five seconds after the second pair is suspended, the ERP issues an SNMP notification (trap 202) listing the further suspended pairs.

An automation program, which is triggered by the notification during the QF interval on the erroneous copy pairs, can issue a *freeze* (refer to "Freeze operation" on page 78) to all LSS pairs relating to the application, affecting the other set of primary volumes in the consistency group. The freeze operation causes the primary volumes to suspend, and go into QF, as well as terminating their associated paths. Therefore, all updates at the recovery site are halted, resulting in a consistent checkpoint. Notice that such an automation program is not part of the ESS software and must be supplied by the user.

> **Attention:** It is important to notice that the QUEUE FULL condition is presented only for that primary volume that is affected by the error (in the case of path failures, multiple volumes are often affected). But a freeze operation is performed at the LSS level, causing all volumes in that LSS to go into suspended state with QF condition and terminating all associated paths.
>
> Therefore, do not intermix volumes relating to different applications in an LSS pairing that is part of a consistency group, because the non-related volumes will be frozen too.

The consistency grouping timer function gives the automation software time to issue the automated commands. The QF time-out value can be displayed and changed from the Logical Subsystems panel of the ESS Copy Services Web user interface by selecting the desired LSS and then clicking the **Properties** button. The Logical Subsystems Properties panel is displayed, and there in the **Consistency group time out** field, you can check or modify the time-out value (refer to 4.5, "PPRC data consistency" on page 72 and 6.6.4, "LSS properties" on page 190). The default value for this parameter is set to two minutes.

Figure 4-13   Error sequence on consistency groups

The following examples illustrate the contents of trap 200 and 202 issued to the automation program by the ESS ERP.

Example 4-1 illustrates that the PPRC pair with LUN ID 0 in LSS 2 in the consistency group is in error.

Example 4-1   SNMP trap 200

```
1002: 2002/09/05 13:51:32 CDT
LSS-Pair Consistency Group PPRC-Pair Error
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI:  IBM 2105-800 13-ABC12  2  0  8
SEC:  IBM 2105-800 75-DEF34  2  0
```

Example 4-2 shows that the PPRC pair with LUN ID 2B in LSS 14 has become suspended.

Example 4-2   SNMP trap 202

```
2002/08/23 16:34:02 CDT
Primary PPRC Devices on LSS Suspended Due to Error
UNIT: Mnf Type-Mod SerialNm LS LD SR
PRI:  IBM 2105-800 13-ABC12 14 2B 09
SEC:  IBM 2105-800 75-DEF34 14 2B
Start: 2002/08/23 16:33:57 CDT
PRI Dev Flags (1 bit/Dev, 1=Suspended):
FFF0000000100000000000000000000000000000000000000000000000000000
```

Figure 4-14 on page 77 provides a summary of consistency groups.

```
    ■ Consistency grouping:
        ► Enables the *Queue-Full* (QF) condition.
        ► Activated.
            – When defining the paths between primary and secondary LSSs or as LSS property (path
              independent)
        ► Triggers the QF condition for SYNC pairs when suspended because of errors.
    ■ Queue Full:
        ► Errors in a volume pair result in the LSS pair being suspended and the QF instigated.
        ► While in QF, all write operations to the affected primary LSS are temporarily queued.
        ► Meanwhile, PPRC sends alert messages to host operating system for the suspended LSS
          pair.
        ► The alerts can be used to trigger automated operations to create a point of consistency.
        ► QF time-out (by default, two minutes) causes the QF to end, and the writes are resumed.
            – The thaw operation allows writes to resume before QF time-out expires.
    ■ QF time-out can be modified from the Logical Subsystems panel when
      using the ESS Copy Services Web user interface.
    ■ Consistency grouping is valid to create point of consistency when
      volumes are in SYNC mirroring. Not necessary when volumes are in
      PPRC-XD relationship.
        ► PPRC-XD always keeps fuzzy copy at the recovery site.
        ► PPRC-XD pairs become consistent only when all updates are transmitted to secondary by
          catch-up while application writes are quiesced.
        ► Gives time for multiple freeze operations in order to reach global consistency.
```

Figure 4-14   Consistency group summary

CRIT attribute

The CRIT attribute is used to determine the behavior of the PPRC pairs or consistency
groups after a failure in communication between the primary and secondary ESS, when all
the paths between a pair are lost. It is used to maintain the consistency of the replicated data.

Note: This parameter only applies to System/390 volumes. This section is included here
for the sake of completeness.

The ability to use the CRIT parameter is set in the ESS Vital Product Data (VPD). In the event
that this parameter needs to be altered, a power cycle of the ESS is necessary. The use of
this parameter is set using ESS Web CopyServices or through TSO commands.

CRIT (NO) specifies that following an I/O completion error to the secondary volume, PPRC
suspends the copy pair, allowing subsequent write requests to the PPRC primary volume to
be satisfied. All writes subsequent to the suspension of the copy pair are recorded for update
on the resolution of the problem. This is the default setting.

CRIT (YES) specifies that if an I/O error to the secondary volume occurs, PPRC either allows
or does not allow subsequent writes to the primary, depending on how the storage subsystem
is configured. The PPRC pair then remains in a suspended state until you correct the problem
and either issue a command to resynchronize the PPRC pair or delete the PPRC pair.

The implementation of CRIT (YES) on the ESS is similar to the implementation on the IBM 3990. There is an option that can be set by the CE in the VPD of the ESS that determines how this CRIT(YES) setting will behave in an error situation:

► CRIT=YES - Paths (light version)

– Suspend the pair and do not accept any further writes if the control units can no longer communicate.

– Suspend the pair and accept further writes if the control units still can communicate with each other. The reason for not being able to copy the data to the remote volume is probably only a device problem on the secondary site and not a disaster. Therefore we continue with write operations to the primary volume. The ESS records which cylinders have changed. After investigating the problem and after it has been solved, you can re-synchronize the source and target volume again.

► CRIT=YES - All (heavy version)

– Suspend the pair and do not accept any further writes to the primary volume if data cannot be sent to the secondary volume.

Freeze operation

With a *freeze* operation, you can stop the write activity on all the active PPRC primary and secondary volumes of a given source and target LSS pair. This function enables you to maintain secondary volume update consistency. It affects all the volumes in the consistency group that are in a PPRC active copy process: duplex, duplex pending synchronous, or duplex pending PPRC-XD states. It does not affect the suspended and simplex volumes or volumes that are not part of the consistency group.

The freeze function operates on LSS pairs that have been defined with (or without) the **PPRC consistency group** option enabled. The freeze operation has three effects:

► The paths that connect the pair of LSSs being frozen are terminated.

► The active consistency group volumes under the frozen LSS pair are suspended. This state transition (to suspended) is then communicated to the host with SNMP alert messages. These alert messages can be used by automation routines to trigger other recovery operations.

► If the **PPRC consistency group** option was enabled at path definition time or in the LSS properties panel, then, additionally, the QF condition is instigated, so the write activity to the primary LSS is temporarily queued. During this QF interval, other automated operations can be triggered, for example, freezing other application-related LSS pairs.

When several freeze operations are executed within the QF time-out interval for all the application-related LSS pairs, then you can insure a global consistent checkpoint at the recovery site. For this to be possible, you should have defined all these related LSS pairs associations with the **PPRC consistency group** option.

When you suspend a pair of volumes, whether specifying the source or the target LSS, you will observe that both volumes of the pair appear as suspended after the operation. Both volumes are still linked by existing logical paths. But this is not the case when you do a freeze operation. For this situation, only the primary volumes show as suspended. The secondary volumes will appear as an active secondary.

When freezing volumes in a duplex state, the secondaries will be a consistent copy. When freezing volumes in duplex pending PPRC-XD or duplex pending synchronous state, the secondaries will be in an indeterminate point of consistency.

Once the error has been resolved, the consistency group enabled LSSs are *thawed* to remove the resulting QF condition, and to reset the suspended volume count table. The thaw operation is started when you run a **Consistency group created** task at LSS level. This operation closes the time-out window, and allows writes to the primary volumes to resume. However, the affected pairs remain suspended, and the paths will not be automatically re-established. This has to be done manually, so it is advisable to create and save these tasks.

Before re-establishing the paths, in order to re-establish the pairs, you may perform a FlashCopy of the secondary volumes because at this moment they are holding a globally consistent point-in-time copy of the data. Figure 4-15 summarizes the characteristics of the freeze operation.

- The freeze operation works on an LSS basis:
  - Duplex, duplex pending PPRC-XD, and duplex pending SYNC volumes are frozen.
  - Non-PPRC (simplex) and suspended pairs in the same LSS are not frozen.
  - PPRC pairs not part of the consistency in the same LSS remain unaffected.
- The primary LSS stops propagation of updates to the secondary LSS:
  - Removes the paths between the primary and secondary LSS pair
  - Suspends the volumes pairs
    - Secondary volumes that were in duplex state will have a consistent point of data.
    - Secondary volumes that were in duplex pending XD or duplex pending SYNC will have an indeterminate consistency.
- If LSSs were consistency group enabled:
  - The LSS is put in the QF condition when the first write comes.
  - The QF prevents updates to the primary volumes during the default two minute time-out interval. Updates are temporarily queued.
    - The time-out interval will end if consistency-group-created task (*thaw*) is invoked.
      - PPRC pairs remain suspended.
    - The time-out interval value (two minutes) can be changed with the Web user interface.
- To resume the mirroring activity to the secondary volumes:
  - The paths must be re-established.
  - The pairs must be re-established.

Figure 4-15   Freeze and resume summary

4.6 Planning for PPRC

For using PPRC on an Enterprise Storage Server, some requirements must be fulfilled. Besides the planning of volumes and Copy Services tasks, you have to check the following items:

► **PPRC hardware requirements**

PPRC is only supported between the same type of storage subsystems. Therefore, it is only possible for an Enterprise Storage Server to perform PPRC replication to another Enterprise Storage Server. It is possible to carry out PPRC between different models of the ESS. An ESS F20 can be in a PPRC relationship with an ESS 800, for example.

You will need to purchase ESCON adapters and cables, possibly with additional equipment, depending on the distance between the primary and the secondary ESSs. ESCON connections have to be configured between the primary and secondary units (see 4.7.1, "ESCON links" on page 82) over which the replication will occur. Figure 4-16 on page 81 illustrates a Copy Services configuration.

▶ **PPRC software requirements**

PPRC is a part of Copy Services, which is a combination of software and licensed internal code that runs on each ESS cluster in a Copy Services domain.

Although PPRC is part of the ESS-internal software, it requires you to purchase a usage license. It needs to be activated on the Enterprise Storage Server clusters that will be participating, using an activation key applied by the IBM customer engineer (IBM CE).

▶ **Network requirements**

All the ESS clusters should be configured to communicate over Ethernet using TCP/IP. Each cluster within each ESS will have their own TCP/IP network address.

The Copy Services servers and their client clusters need to have their IP addresses resolvable, either using DNS or a standard /etc/hosts file during installation. The method for IP-Address resolution can be configured by the IBM CE during installation.

▶ **Optional SNMP configuration**

The ESS has the ability to raise SNMP alerts when an error condition has been encountered. When an error is encountered, an alert is posted to a systems management console, such as Tivoli Netview. During installation, the TCP/IP network address of the systems management console, along with the community to which the ESS alerts will be posted, will need to be configured. This becomes more pertinent for capturing the new traps that are raised during ESCON link loss (see 4.7.4, "PPRC path failure alerts" on page 88) and errors on volumes belonging to *consistency groups* (see 4.5, "PPRC data consistency" on page 72). It is the systems management console that will perform the automation on the receipt of an alert.

▶ **Management stations**

You have to plan which machines will be used to control PPRC operations. These machines must be connected to the ESS through the IBM TotalStorage Enterprise Storage Server Network (*ESSNet*). The ESSNet is the ESS access facility that IBM installs when they install your ESS, consisting of a switch (hub) and other networking components and the dedicated IBM TotalStorage Enterprise Storage Server Master Console (*ESS Master Console*).

With respect to disaster recovery, you should plan for at least two management stations (one in each site) that do not depend on the ESS for proper operation.

▶ **Web browser**

A browser for the ESS Specialist has to be installed on the management stations. For recommendations about Web browser versions see *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7448.

▶ **ESS Copy Services CLI**

If you plan to use the CLI, you have to install the CLI also on the management stations. For details, see the *IBM TotalStorage Enterprise Storage Server Command-Line Interfaces User's Guide,* SC26-7494.

▶ **Task structure**

To avoid human mistakes and automate routine actions, you will create and save tasks in advance. We highly recommend that you use an extendable task naming schema (see Appendix E, "Sample naming convention" on page 499). Prepare tasks for PPRC establishment with several options. Table 4-4 on page 81 summarizes the combinations of options available for establishing PPRC relationships. Also, prepare the corresponding tasks to suspend and terminate the PPRC relationships to be scheduled to the primary as well as to the secondary LSS.

Table 4-4   Task options for establishing PPRC relations

Establish PPRC	Permit read from secondary	Suspend PPRC after establish complete	Asynchronous Cascading PPRC
Copy entire volume	SYNC + XD	SYNC only	SYNC + XD
Do not copy volume	SYNC + XD	Not available	SYNC + XD
Copy out-of-sync cylinders only	SYNC + XD	SYNC only	SYNC + XD
PPRC Failover (SYNC only)	Not available	Not available	SYNC only
PPRC Failback	SYNC + XD	SYNC only	SYNC + XD
Notes: SYNC and XD mean that the combination is available for synchronous or XD mode. The options **Critical volume mode** and **Permit establish if target is online** are not available for open-systems volumes.			
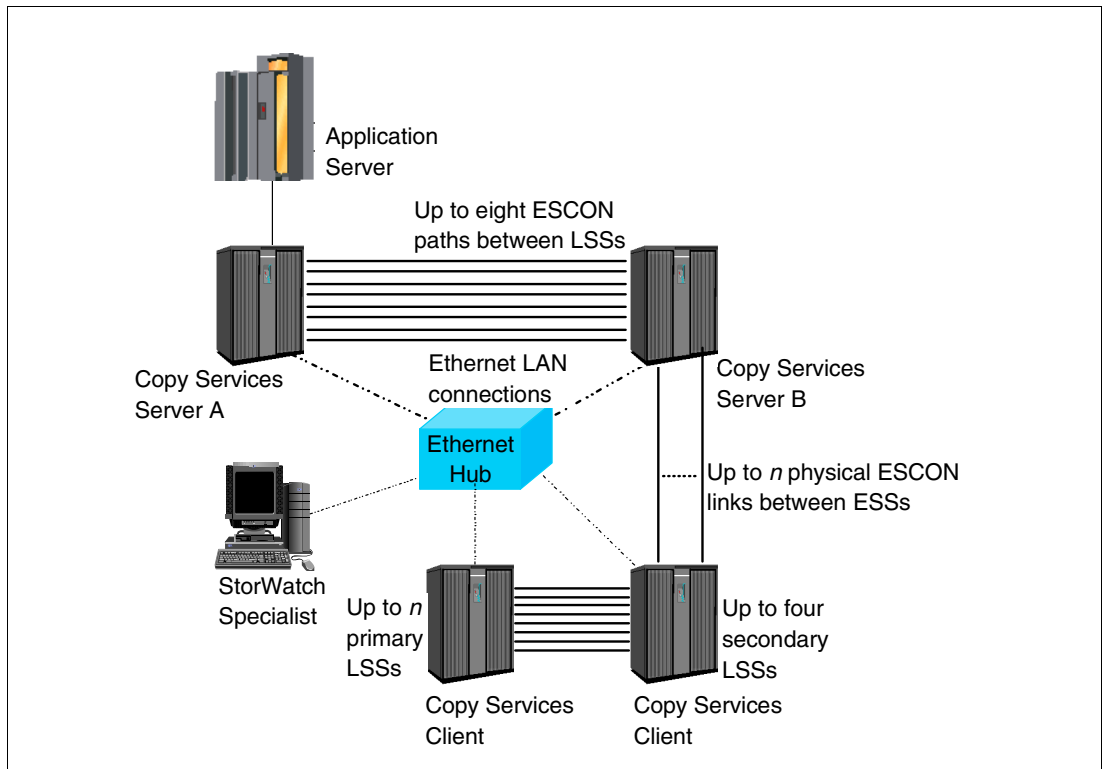
Figure 4-16   ESS Copy Services configuration

When planning resources for your PPRC environment, you need to consider which data is critical and requires protection in the event of a disaster. Knowing this, it is possible to plan the throughput and capacity requirement of the secondary ESS. It is important to realize that the capacity needed on the secondary ESS for disaster recovery may not have to be initially as large as the primary ESS. Some applications and data will be more critical than others. You need to size your secondary ESS based on your critical business requirements, possibly with some headroom for applications of intermediate importance.

4.7 PPRC connectivity

When implementing a PPRC configuration, one requirement is the definition of the paths that PPRC is going use to communicate the primary and secondary LSSs. These logical paths are defined upon existing physical links. The physical links can either be ESCON links or Fibre Channel links.

In this section, we discuss the physical connectivity and the logical definitions of paths between primary and secondary ESSs.

> **Note:** The present section discusses PPRC links and paths when using ESCON links. Because many of the discussions are valid to both ESCON and Fibre Channel links, we recommend that you first read the present section and then proceed to read 4.8, "PPRC using Fibre Channel links" on page 92 for the Fibre Channel specifics.

4.7.1 ESCON links

ESCON (Enterprise Systems Connection) is an IBM architecture with fiber optic technology. ESCON was the first implementation of storage networking before there was a Fibre Channel standard. Introduced in 1989, ESCON is a 200 Megabit per second (Mb/sec.) serial protocol that can achieve effective instantaneous data transfer rates up to 18 MB/sec. ESCON is used to connect S/390 and later generation IBM "mainframe" computers to peripherals such as disk systems. It allows a data center installation to replace electrical-copper cables with fiber optic cables. Beyond this physical aspect, ESCON provides the capabilities to connect ESS units at distances beyond 103 km. Similar to Fibre Channel, ESCON is based on message exchanges.

An optical fiber functions as a kind of wave guide for light. It is usually made of silica glass. The fiber itself has a central core and a surrounding cladding of slightly different glass material. The physical size of an optical fiber is determined by the diameter of the core and cladding, expressed in micrometers (μm). A fiber optic cable having a core diameter of 62.5 μm and a cladding of 125 μm is designated as 62.5/125 μm optical fiber. Other fibers commonly used for ESCON are 50/125 μm and 9/125 μm. The ESCON cables consist of two fibers.

Two modes can be used to send light signals through an optical fiber: single mode (mono mode) or multimode. The optical fibers used are called accordingly single-mode or multimode fibers. These fibers have different physical dimensions and light transmission characteristics. The multimode fibers supported by IBM are either 62.5/125 μm or 50/125 μm. The light source used for multimode fiber is usually a light emitting diode (LED). The LED jumpers are orange, and their duplex connectors are black. The maximum distance for a multimode fiber link is 3 km if 62.5/125 μm fibers are used and 2 km if 50/125 μm fibers are used. The maximum distance for a multimode fiber link is 2 km if both 62.5/125 μm and 50/125 μm fibers are used.

Single-mode fiber usually has a core diameter of 8 to 10 μm and a cladding diameter of 125 μm. The light source used for single-mode fiber is a laser. IBM supports single-mode 9/125 μm fibers for use in an ESCON environment. Single-mode fiber must be used on links between devices with ESCON-Extended Distance Feature (XDF) adapters installed. These adapters are available only on ESCON Directors and on remote channel extenders.

Fiber jumper cables of 9/125 μm for IBM ESCON Directors are supplied by IBM in standard lengths up to 122 m (400 ft.), if XDF ports are ordered. The number of jumper cables supplied at no charge depends on the number of XDF ports installed. XDF jumpers are yellow, and their duplex connectors are grey. The maximum distance for a single-mode fiber link is 20 km.

The number of ESCON links between the primary and secondary units is dependent on the performance requirements and the availability of slots in the hostbays (possibility of 32 ESCON links). ESCON channels provide 160 Mbps point-to-point theoretical throughput.

The ESCON host adapter ports have the capability of operating in two ways, but when implemented for PPRC transmission, they operate in either one of those two ways:

► The ESCON host adapter port operates in *control unit* mode, when it is talking to a host. In this mode, an ESCON port can also receive data from a primary ESS, when the ESS port is connected to an ESCON director.

► An ESCON port is operating in *channel mode*, when it is used on the primary ESS for PPRC operations onto a secondary ESS.

Therefore, ESCON links for PPRC are unidirectional. The primary unit ESCON port (the one in channel mode) has to be dedicated for PPRC. The ESCON port on the secondary unit can also be used for S/390 host attachment, provided the ESCON director is used, and the host is connected to it. This is illustrated in Figure 4-17.
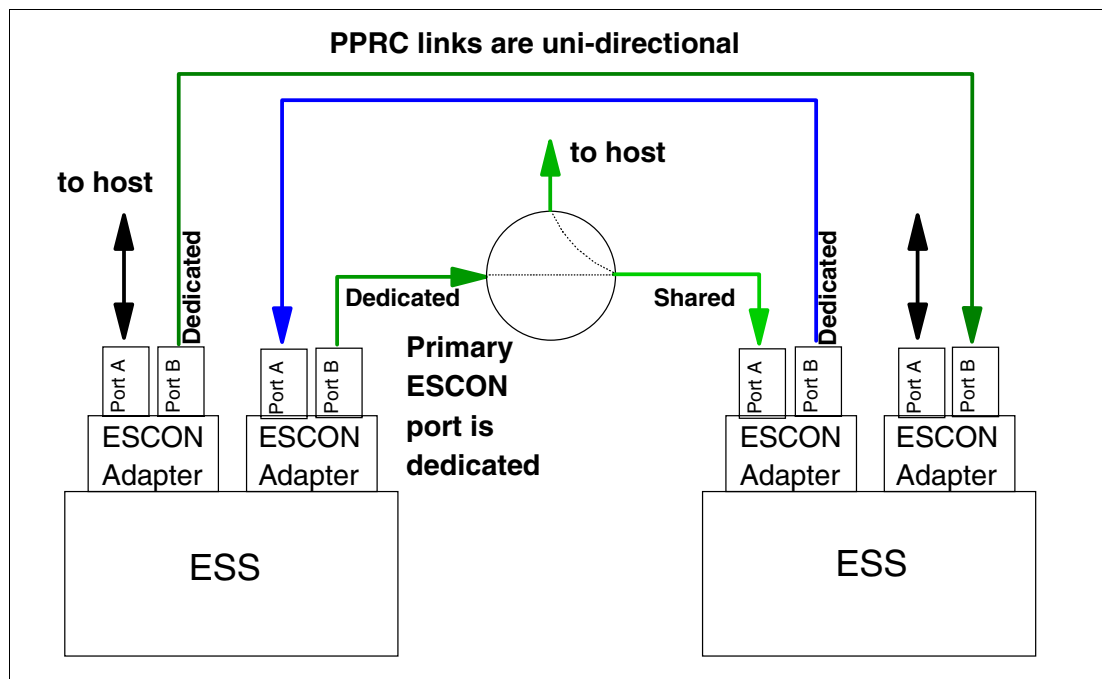
Figure 4-17   PPRC ESCON connectivity

Direct ESCON connection

Direct ESCON link connections between primary and secondary ESS units can be established without the need of any intermediary devices. Standard ESCON adapters are used in the ESS; this connection can only be a multimode connection. Direct connections can be established for up to 2 km with 50/125 µm multimode fiber ESCON cables and 3 km with 62.5/125 µm multimode fiber ESCON cables. The ESCON channel performance is a function of distance. The direct connection is shown in Figure 4-18 on page 84.
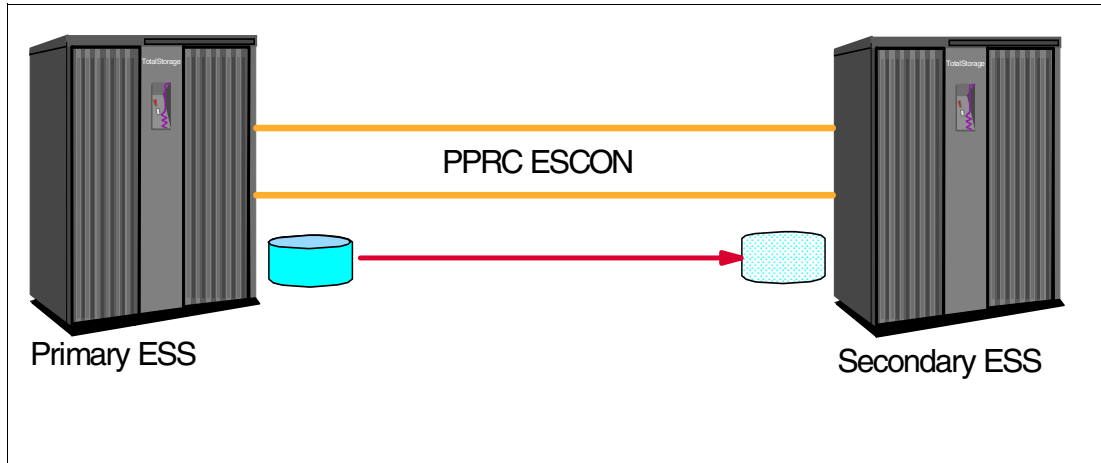
Figure 4-18   Point-to-Point PPRC configuration

ESCON Directors

The distance between primary and secondary ESS units can be extended using an ESCON Director as a signal relay. The maximum distance between each ESS and the ESCON Director can be 2 km (50/125 µm fiber) or 3 km (62.5/125 µm fiber), resulting in a maximum distance between both ESSs of 4 km or 6 km, respectively (see Figure 4-19).
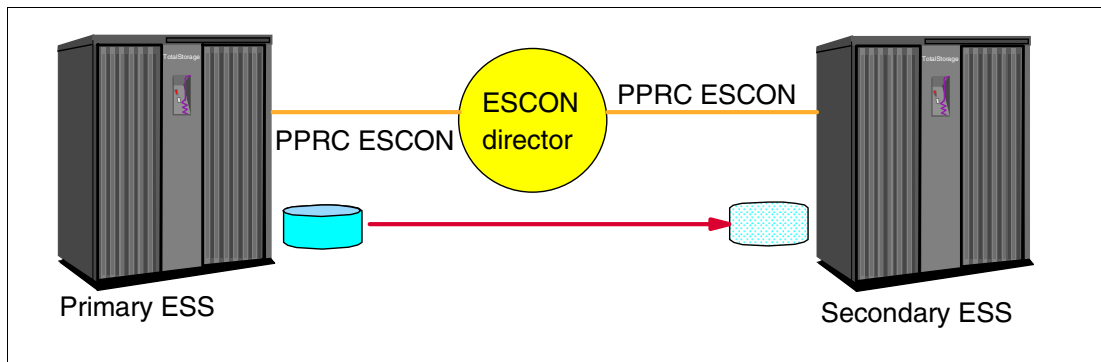
Figure 4-19   Configuration with one ESCON director

Adding a second ESCON Director is the next step to increase the distance. The distance between the two ESCON Directors can be 2 km for 50/125 µm fiber and 3 km for 62.5./125 µm fiber.

You can use Extended Distance Feature (XDF), which installs ESCON-XDF adapter ports on the ESCON Directors with 9/125 µm single mode fiber optic cables. This allows the maximum distance between the XDF Director ports to extend to 20 km, providing an overall distance of 26 km between ESS units. This can be seen in Figure 4-20 on page 85.
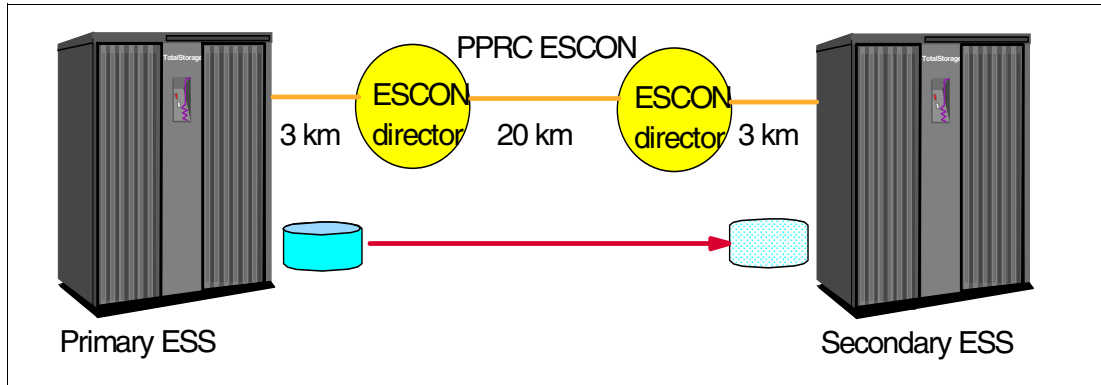
Figure 4-20   Configuration with two ESCON directors

Channel extenders

Various channel extender vendors can also be used to increase the distance between ESS servers via a Wide Area Network (*WAN*) using a variety of connections, such as Fibre Channel, Ethernet/IP, ATM-OC3, and T1/T3

When using channel extender products with PPRC, the channel extender vendor will determine the maximum distance supported between the primary and secondary ESS. The channel extender vendor should be contacted for their distance capability, line quality requirements, and WAN attachment capabilities. Evaluation, qualification, approval, and support of PPRC configurations using channel extender products are the sole responsibility of the channel extender vendor.

Dense Wavelength Division Multiplexor (DWDM)

Wave® Division Multiplexing (*WDM*) and Dense Wave Division Multiplexing (*DWDM*) are used for connectivity via Metropolitan Area Networks (*MAN*), and are the basic technology of fibre optical networking. These techniques are used for carrying many separate and independent optical channels on a single dark fibre.

A simple way to envision DWDM is to consider that at the primary end, multiple fibre optic input channels such as ESCON, Fibre Channel, FICON, or Gbit Ethernet are combined by the DWDM into a single fibre optic cable. Each channel is encoded as light of a different wavelength. You might think of each individual channel as an individual color: the DWDM system is transmitting a *rainbow*. At the receiving end, the DWDM fans out the different optical channels. DWDM by the very nature of its operation provides the full bandwidth capability of the individual channel. As the wavelength of light is, from a practical perspective, infinitely divisible, DWDM technology is only limited by the sensitivity of its receptors, as the total aggregate bandwidth possible.

The IBM 9729 Optical Wavelength Division Multiplexer (*MuxMaster*) enables a 50 km distance between MuxMaster units. ESCON is used to attach the ESS to it.

The IBM 2029 Fiber Saver, also known as Dense Wavelength Division Multiplexer (DWDM), supports ESCON, FICON, Fibre Channel, and many more protocols, enabling up to a 50 km distance between Fiber Saver units that are ESCON attached to the primary and secondary ESS. You can use it for Fibre Channel, network, and telephone links between the sites as well. Figure 4-21 on page 86 shows how you can reach a maximum PPRC-SYNC distance of 103 km using the IBM 2029 Fibre Saver.
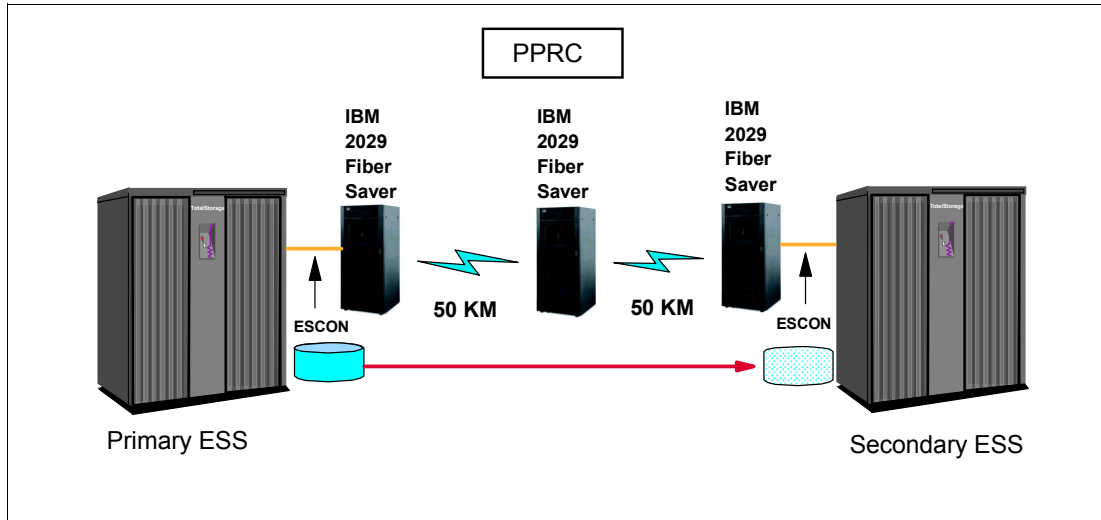
Figure 4-21   Configuration using IBM 2029 Fibre Saver

Other Channel Extension, DWDM, and Network connectivity options are:

► CNT UltraNet Storage Director (USD)
► CNT (INRANGE) 9801 SNS
► Cisco ONS 15530/15540
► Nortel Networks OPTera Metro 5200/5300

A complete and current list of PPRC supported environments, configurations, networks, products, and required ESS LIC levels is available at:

http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Attention: The DWDM vendor should be contacted regarding evaluation, qualification, approval, and hardware and software support prerequisites when using their products in an ESS PPRC configuration.

4.7.2 Configuring PPRC links

The ESCON protocol has been streamlined with less handshaking and larger frames transmitted between ESSs. ESCON connections have to be configured between the primary and secondary units (see 4.7.1, "ESCON links" on page 82) over which the replication will occur.

When establishing ESCON connections between two ESS subsystems, there are some requirements to observe:

► As many ESCON links as required (depending on performance requirement and host bay slot availability) can be used to connect the ESSs. However, only eight logical ESCON paths may be configured between the primary and secondary LSSs.

► A primary LSS can be connected via ESCON links to up to four secondary LSSs.

► A secondary LSS can be connected to any number of primary LSSs, limited by the number of ESCON links available.

► PPRC links are unidirectional, because the ESCON port at the primary ESS is reconfigured to act like an ESCON channel in a host S/390 processor. The primary ESCON port is dedicated to PPRC.

An ESCON PPRC link can be used only to transmit data from the primary storage control to the secondary. If you want primary and secondary volumes on each of two ESSs, you need ESCON PPRC links in each direction. The number of links needed in each direction depends on the total write activity to all the primary devices in each ESS.

4.7.3 PPRC paths definition

The ESCON links, as discussed above, are the physical layer of the PPRC connections. The logical layer consists of *PPRC logical paths*, defined over the physical links. Prior to setting up volume copy pairs, PPRC requires logical paths to be established between the primary and the secondary ESS logical subsystem (LSS). Each LSS with primary volumes requires at least one path to be set to the LSS that holds the secondary volumes (Figure 4-22 on page 88).

> **Note:** In the open systems world, we refer to a logical subsystem or LSS. This equates to a logical control unit (LCU) in the mainframe world. LSS and LCU are sometimes freely interchanged.

Up to eight paths can be defined per LSS pair. For each LSS pair, each path is defined over a different physical link. The path definition by identifying the source and the target LSSs implicitly sets the direction of the link. A link will operate in only one direction. If you want to have a path in the opposite direction, then you will have to use a different link. Note that if you delete all the paths over a physical link, then you can reuse that link for paths in the opposite direction.

When you establish a path:

► You identify the source and target LSSs.

► You identify the physical link that will be used between the primary and secondary LSSs.

► Then, optionally, you can set the **Consistency Group** option (discussed in 4.5, "PPRC data consistency" on page 72).

> **Important:** Be aware that each path definition you make will totally replace any existing path definitions for that same pair of LSSs. So, if you are adding a path, then you should redefine all previous path definitions between these two LSSs explicitly.

Each ESCON link supports 64 logical paths, so even with 16 LSS defined you are able to set up a logical path from each LSS to each LSS with only four ESCON PPRC links. We always recommend that you use all available links for each LSS pair used for PPRC. That gives you maximum protection against ESCON link failure.

For a pair of ESSs, the physical links between them can be shared with many path definitions from different pairs of LSSs within those ESSs. In practice, many installations configure paths to share the same links for easier management.
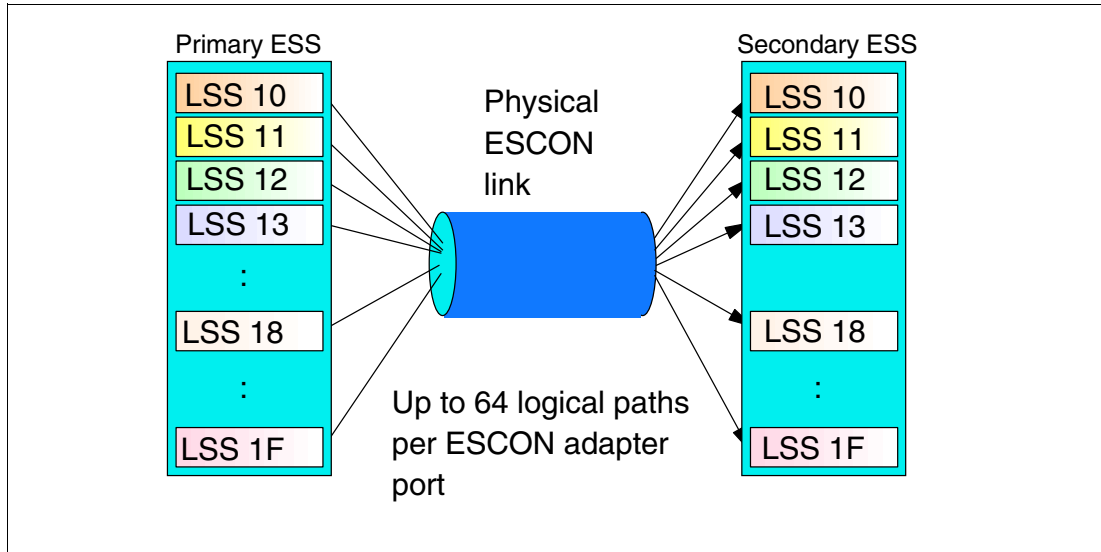
Figure 4-22   Logical paths

For redundancy, and for performance, it is recommended to maximize (up to eight) the number of paths between a pair of source and target LSSs. When a pair of volumes is initially established or at the re-synchronization time of a suspended pair of volumes, PPRC is able to use all the defined paths in parallel. Especially with long-distance links between data centers, you should distribute these paths evenly over so many ESCON adapters, directors, and DWDM links as possible to keep paths functional when some of the links are temporary unavailable.

4.7.4 PPRC path failure alerts

When a path fails, a notification is presented in the form of an update to the WUI in the Path Origin panel of the Paths tab or as an SNMP trap 100 to a systems management console (SMC). The activity for that pair of LSSs continues if more paths are available, rerouting the updated transmissions over alternate paths between that LSS pair.

Example 4-3 illustrates the trap information for a failed path. It contains identification of a degraded transfer rate, along with the serial numbers of the primary and secondary ESSs involved and SAID (`PP` and `SP`) details of the path that has failed, with the return code (refer to Appendix D, "System Adapter ID (SAID)" on page 493 for a description).

Example 4-3   SNMP trap information of path failure

```
2002/09/05 16:03:51 CDT
PPRC Links Degraded
UNIT: Mnf Type-Mod SerialNm LS
PRI:  IBM 2105-800 13-ABC12 02
SEC:  IBM 2105-800 75-DEF34 02
Path: Type PP   PLink  SP   SLink   RC
1:    ESCON 0001 00     0008 00      OK
2:    ESCON 0020 00     0020 00      08
```

When the last existing path, between a pair of source and target LSSs, fails due to an error condition, this event is considered by PPRC to be a major error condition. In an open systems environment, an SNMP trap 101 is sent to the SMC. Under this error condition, the pairs will be suspended and all volumes in a consistency group exhibit Queue Full condition (refer to 4.5, "PPRC data consistency" on page 72). Example 4-4 illustrates a total link-down event.

Example 4-4   SNMP trap information for total path failure

```
2002/09/05 16:03:42 CDT
PPRC Links Down
UNIT: Mnf Type-Mod SerialNm LS
PRI:  IBM 2105-800 13-ABC12 02
SEC:  IBM 2105-800 75-DEF34 02
Path: Type  PP    PLink  SP    SLink   RC
1:    ESCON 0001 00     0008 00      08
2:    ESCON 0020 00     0020 00      08
```

Example 4-5 illustrates the contents of trap 102 on the resolution of the error event and the reconstitution of the paths.

Example 4-5   SNMP trap information for path recovery

```
2002/09/05 16:03:58 CDT
PPRC Links Up
UNIT: Mnf Type-Mod SerialNm LS
PRI:  IBM 2105-800 13-FCA36 02
SEC:  IBM 2105-800 75-18592 02
Path: Type  PP    PLink  SP    SLink   RC
1:    ESCON 0001 00     0008 00      OK
2:    ESCON 0020 00     0020 00      OK
3:    ESCON 0001 00     0008 00      OK
4:    ESCON 0020 00     0020 00      OK
5:    ESCON 0001 00     0008 00      OK
6:    ESCON 0020 00     0020 00      OK
7:    ESCON 0001 00     0008 00      OK
8:    ESCON 0020 00     0020 00      OK
```

4.7.5 Performance considerations

In the following sections, we discuss the considerations involved when setting up the Copy Services of the Enterprise Storage Server (ESS) in order to achieve better performance. This should help you understand the performance impact of ESS Copy Services. As there are many different parameters that have an influence on performance, such as applications, type of workload, and configuration of the Enterprise Storage Server, the information should serve as a guideline when planning ESS Copy Services.

As a rule of thumb, the number of ESCON paths between ESS and the type and placement of ESCON adapters have the most significant influence on PPRC performance. Other factors are processors, caches, and path grouping.

Keep in mind that the general ESS performance considerations, such as volume placement, type of hard disk drives (rotational speed and number of drives per RAID array), or the amount of storage per host adapter still apply when planning for PPRC.

Optimized PPRC communication

There were certain modifications made to the ESCON protocol used for PPRC communication of the Enterprise Storage Server, in particular:

► A larger frame size, which results in less overhead during PPRC communication.

► Less handshaking between the two communicating ESSs, which makes transfer of data more efficient. The handshake was reduced from six down to three exchanges.

Cache and NVS sizes, number of processors

One of the major considerations about capacity planning for PPRC used to be the cache and NVS sizes of both application and secondary ESSs. Today, there are multiple cache sizes available, depending on the model of the ESS configurations. So you will need to take these into account when planning a PPRC configuration. PPRC requires extra I/O activity to the back-end of the ESS as compared to non-PPRC boxes, so it is important to consider larger cache sizes, or the cache backstore ratio when using PPRC.

With the ESS F20, there were diminishing returns from using more than 8 PPRC links in terms of 4 KB operations. For the ESS 800 with turbo option (two additional and faster processors in each ESS cluster), the per link throughput scales well to 16 links. In environments with a very high number of I/O operations per second, the ESS 800 turbo model provides better PPRC performance and keeps the I/O response time stable in a wider range of workloads than the ESS 800 standard model. This effect increases with the number of ESCON links. Therefore, when a large number of links is used, the ESS 800 with turbo option is recommended for the best possible performance.

You can call your IBM representative to help you determine the resources that are needed to implement PPRC, and the impact of PPRC on the application system. There are some modeling tools, like DiskMagic, that can predict the PPRC effect on the application system before the real implementation takes place.

Number of ESCON paths between ESSs

The number of ESCON channels needed between PPRC primary and secondary LSS depends on the response time requirements of the application servers in a synchronous environment and the available bandwidth between primary and secondary ESS subsystems in an asynchronous environment.

Always make sure that you are using an appropriate number of physical ESCON links for PPRC between the primary and the secondary ESS. Increasing the number of the physical ESCON links will increase the maximum overall bandwidth for updating the secondary volumes. Using multiple physical ESCON links for a PPRC pair (maximum of eight logical paths per LSS) will improve the response time of an I/O request from the host and minimize the out of sync tracks in a PPRC-XD environment. Keep in mind that too few physical links may result in a bottleneck. A minimum of four links between the primary and secondary ESS are recommended.

Type and placement of the ESCON adapters used for PPRC

Until 2002, the ESCON host adapters on the ESS have used a 32-bit wide bus. With the introduction of 64-bit ESCON Host Adapters for the ESS 800, ESCON speed, efficiency, and throughput are improved. Use of 64-bit ESCON Host Adapters matches the 64-bit internal buses within the ESS 800 and therefore the hardware capability may be fully exploited. The 64-bit adapter has about 10% less overhead than the 32-bit adapters used previously, resulting in lower response times and additional throughput. The throughput for PPRC establish will improve roughly 10% compared to the 32-bit ESCON adapters.

All models of the ESS 800 now come standard with 64-bit ESCON adapters. If you have older ESSs with 32-bit adapters, then you may consider upgrading to 64-bit adapters. Configurations with 64-bit ESCON adapters on one end of a PPRC link and old 32-bit adapters at the other end work also. As might be expected, performance will be somewhere in between what would be seen for a link using two 64-bit adapters and a link using two 32-bit adapters.

Distribute the ESCON adapters used for PPRC and PPRC-XD evenly across the two clusters and the host adapter bays of the ESS. This will distribute the PPRC workload over multiple buses and both clusters.

For example, if there are four ESCON adapters used for PPRC between two Enterprise Storage Servers, place one ESCON adapter in each of the host adapter bays.

Grouping of physical and logical paths

A physical path describes the physical ESCON connection between two Enterprise Storage Servers. A logical path is the connection used for the PPRC copy pair, either between two volumes or two logical subsystems. There could be multiple logical connections established over a single physical connection. This will most likely be the case in a real environment.

Also, consider that multiple logical paths using the same physical ESCON links(s) will share the bandwidth of these ESCON links(s). If there are performance-critical PPRC pairs, we recommend that you separate them on dedicated physical paths so that the I/O traffic of the data copy from the primary to the secondary side will not interfere with I/O traffic of less critical PPRC pairs.

When you plan to use synchronous PPRC as well as PPRC Extended Distance between a pair of ESSs, IBM recommends that you establish separate logical paths over separate physical paths for the copy pairs managed by each of the modes. In other words, for your synchronous PPRC copy pairs, use one set of logical and physical paths between the source and target LSSs; for your PPRC Extended Distance copy pairs, use another set of logical and physical paths between the source and target ESSs.

By keeping the paths separate for the two copy modes, the updates to PPRC Extended Distance target volumes minimize the effect on the I/O performance of the synchronous PPRC pairs. This recommendation only applies to environments where the distance between primary and secondary ESSs does not exceed the synchronous range. This recommendation does not apply to the following situations:

► Using the PPRC Extended Distance mode to establish copy pairs that you convert to synchronous pairs after the bulk copy is complete.

► Converting PPRC Extended Distance copy pairs to synchronous PPRC copy pairs on an LSS-to-LSS level.

► Re-establishing PPRC Extended Distance pairs that were previously synchronous PPRC copy pairs on an LSS-to-LSS level.

When using Asynchronous Cascading PPRC, different physical ports should be used on the intermediate ESS for the primary–intermediate and the intermediate–remote connections. You should keep the distance between the local and intermediate sites as close a possible to minimize the performance impact of the synchronous PPRC operation.

Setup of the secondary ESS

For disaster recovery reasons, you may be doing PPRC between two or more different Enterprise Storage Servers. Under normal operating conditions, you always have a source (primary side) and a target (secondary side) of a PPRC pair.

One single ESS could have up to four secondary ESSs. However, the number of primary servers of a single secondary server is only limited by the number of available ECSON links. So it may be the case that different primary storage servers are connected to the same secondary ESS. In that case, the I/O traffic of multiple primaries has to handled by a single secondary ESS.

Furthermore, it may be possible that secondary volumes from different primary storage servers are placed on the same disks within the same Array (rank). In that case, the response time of each primary storage server will increase if other primaries are doing I/O at the same time, as all requests are handled simultaneously.

Therefore, when planning your Enterprise Storage Server network, keep in mind how many primary storage servers are connected to the same secondary. Distribute the I/O load evenly across the secondary storage server.

Try to distribute the secondary volumes belonging to PPRC pairs from different primaries across all available RAID arrays within the secondary ESS.

4.8 PPRC using Fibre Channel links

Fibre Channel can be used as the communication link for PPRC between primary ESS and secondary ESS. This allows exploitation of existing Fibre Channel infrastructure and capacity.

> **Note:** The present section addresses the specifics for Fibre Channel connectivity. This information should be complemented with the general discussion on PPRC paths and links presented in 4.7, "PPRC connectivity" on page 82.

PPRC over Fibre Channel requires PPRC Version 2 (feature #85xx) and is supported on the ESS Model 800 at LIC level 2.3.0 with Fibre Channel/FICON host adapters (features #3024 and #3025).

When compared with ESCON links, Fibre Channel reduces the link infrastructure by at least 4 to 1 with equivalent or better performance.

When implemented over Fibre Channel links, PPRC can be managed using:

► ESS Copy Services Web User Interface, ESS API, and ESS Copy Services CLI, for all environments
► ANTRQST API, TSO commands, and ICKDSF, for z/OS® environments
► ICKDSF, for z/VM® and VSE/ESA™ environments

4.8.1 Configuration guidelines

An ESS Model 800 Fibre Channel port can simultaneously be:

► Sender for a PPRC primary
► Receiver for a PPRC secondary
► Target for FCP host(s) I/O from Open systems and Linux on zSeries.

A PPRC Fibre Channel link is full duplex. This means that one link can have a PPRC path established in one direction and simultaneously have a PPRC path established in the opposite direction. This contrasts with ESCON PPRC links, which are uni-directional. ESCON ports are dedicated and so cannot be used for host I/O.

As for ESCON PPRC links, PPRC Fibre Channel links support both synchronous PPRC and PPRC extended distance (XD).

Although one PPRC Fibre Channel link would have sufficient bandwidth for most environments, we recommend configuring two PPRC Fibre Channel links between each primary and secondary ESS, as shown in Figure 4-23 on page 93.

.
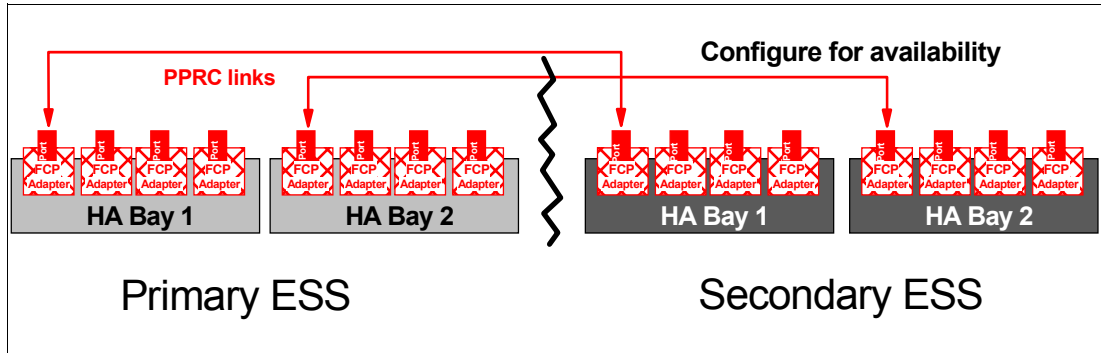
Figure 4-23   Configuring for availability

PPRC Fibre Channel links can be direct-connected (dedicated) or use switches (up to two switches).

Figure 4-24 shows direct PPRC Fibre Channel links. The PPRC Fibre Channel links are direct connected between two Fibre Channel ports. There is no other sharing of these Fibre Channel ports; they are used exclusively for PPRC. Both ports can be sending and receiving at the same time due to the full duplex capability of Fibre Channel. Dedicated Fibre Channel ports guarantee no interference from host I/O. This is recommended, as synchronous PPRC is time critical and should not be impacted by host I/O activity. The PPRC ports can be used to provide connectivity for all LSSs within the ESS and can contain multiple PPRC paths.
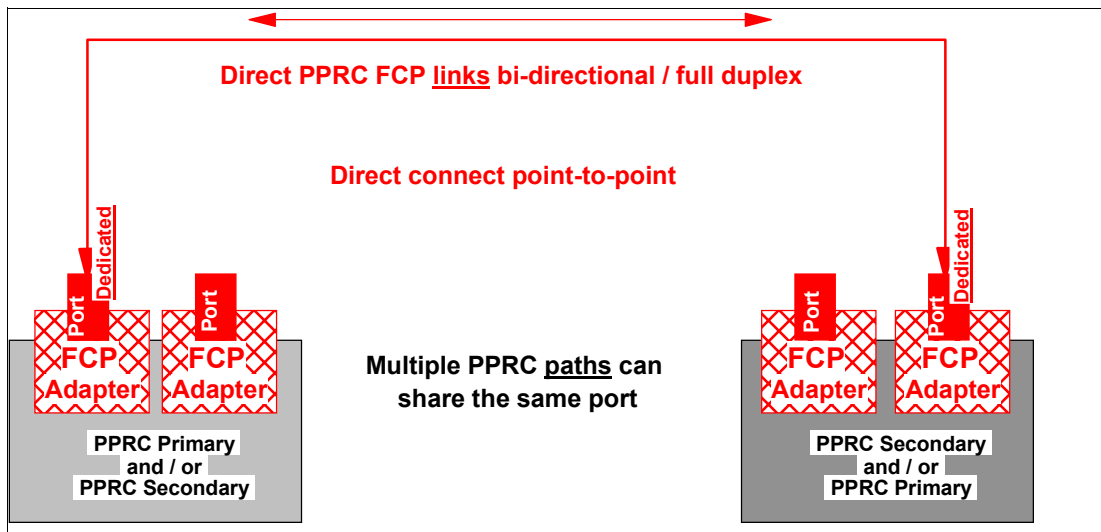
Figure 4-24   Direct PPRC FCP links

Figure 4-25 on page 94 shows a typical configuration for an open systems environment using two switches. As in the direct-connected configuration, we recommend that the PPRC Fibre Channel ports be dedicated and not used for host I/O.
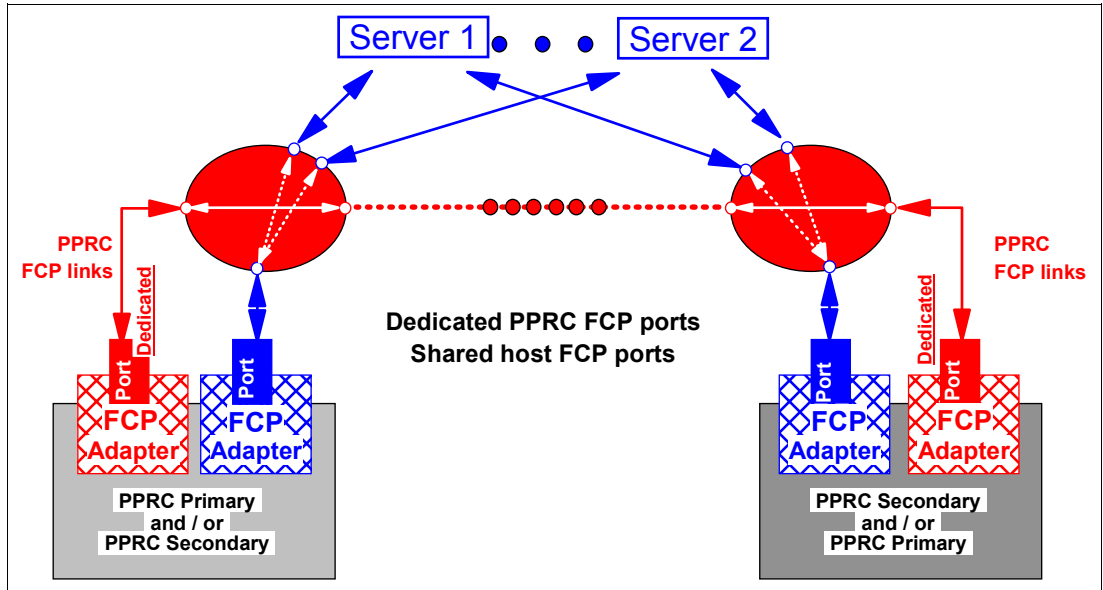
Figure 4-25   Typical open systems host and PPRC connectivity using Fibre Channel

Logical paths

A logical path is a logical (zoned) connection between ESS Fibre Channel Host Adapter and the server I/O Fibre Channel Adapter card. This constitutes the medium for host to ESS disk I/O. This differs from a physical connection, which is the collection of elements (links, ports, adapter cards, and switch ports) that provide the physical path for the information flow. Figure 4-26 shows how multiple logical paths can use the same physical path.
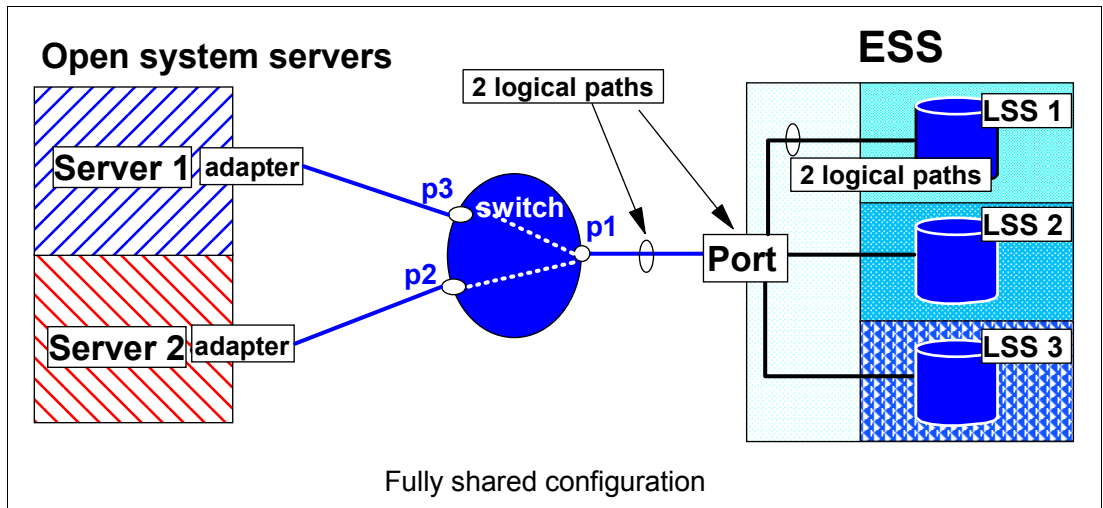
Figure 4-26   Logical paths between servers and LSSs

Logical paths also apply to PPRC connections. A PPRC logical path is a logical connection between the sending LSS and receiving LSS. A PPRC link can accommodate multiple PPRC logical paths.

Figure 4-27 on page 95 shows that if we have a 1:1 ESS mapping, we would have three logical paths over the PPRC link: LLS1 to LSS1, LLS2 to LSS2, and LSS3 to LSS3. If the volumes in LSS1 mapped to all three secondary LSSs and similarly for LLS2 and LLS3, there

would be nine logical paths over the PPRC link. Note that we recommend a 1:1 LSS mapping.
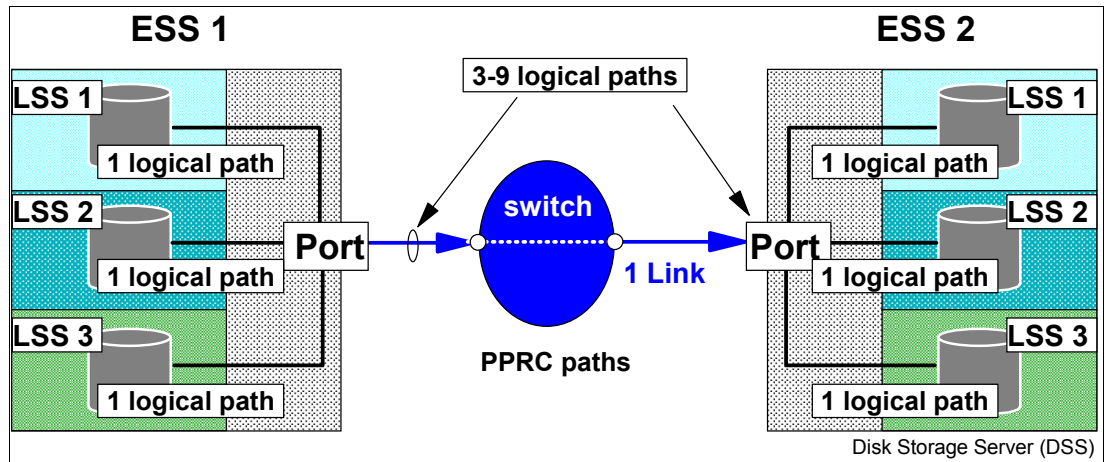
Figure 4-27 Logical paths for PPRC

PPRC Fibre Channel has not changed certain PPRC architectural characteristics:

▶ A primary LSS can maintain paths to a maximum of four secondary LSSs. Each secondary LSS can reside in a separate ESS.

▶ Up to eight logical paths per LSS-LSS relationship can be defined. Each PPRC path requires a separate physical PPRC link. This is illustrated in Figure 4-28.

▶ The ESS Copy Services Server (CSS) can manage up to 4096 volumes/LUNs (2048 pairs). This is the total number, including primary and secondary volumes/LUNs, as well as source and target FlashCopy volumes/LUNs.
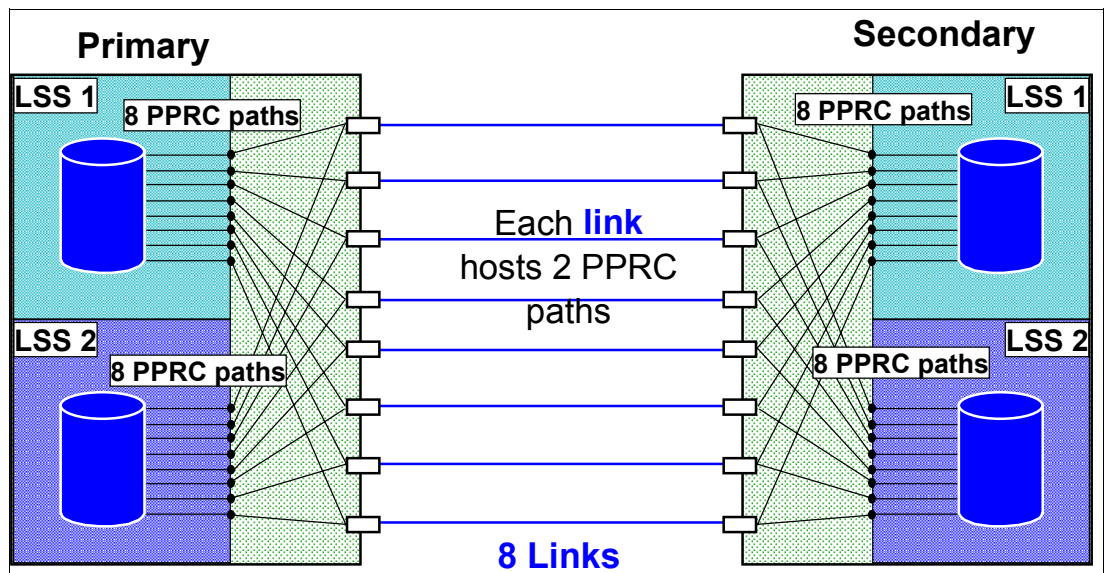
Figure 4-28 Up to eight paths per LSS-LSS relationship

Some PPRC implementation characteristics have been relieved when using Fibre Channel links. This is illustrated in Figure 4-29 on page 96 and they are listed here:

▶ A PPRC Fibre Channel port can host up to 1024 logical paths. The ESCON limit was 64.

▶ A physical link can host up to 128 PPRC logical paths. The ESCON limit was 64.

► A PPRC Fibre Channel port can accommodate up to 126 physical paths (PPRC links). These 126 physical paths are mapped to the ESS Fibre Channel ports through a corresponding SAN fabric.
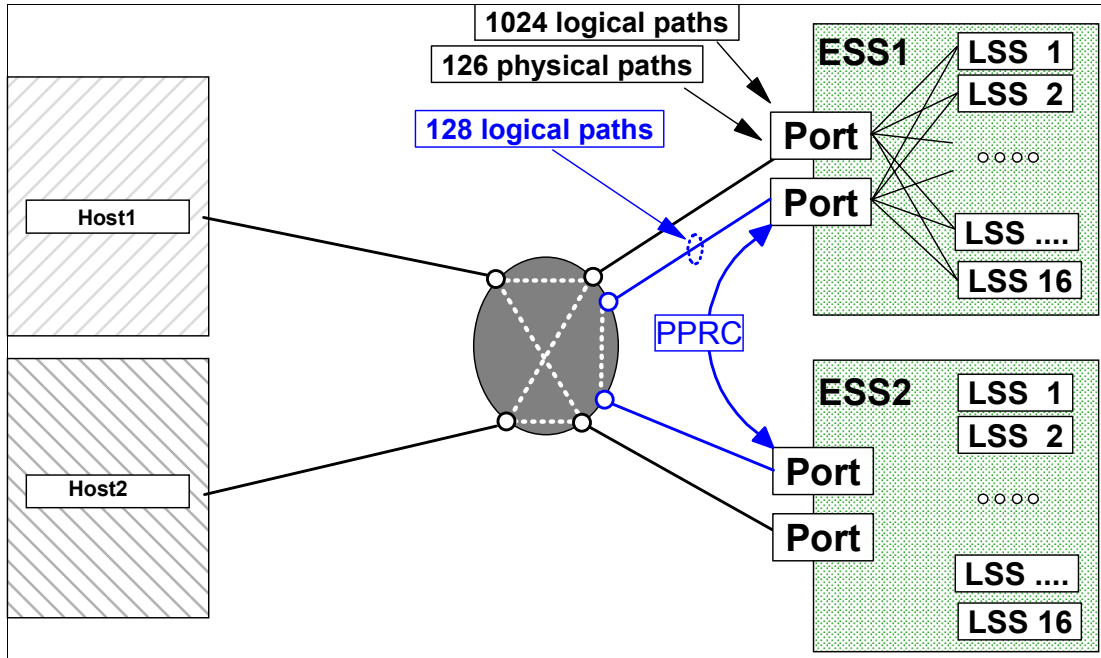
Figure 4-29   Logical path/physical path limits

Table 4-5 compares FCP versus ESCON architectural characteristics.

Table 4-5   Fibre Channel versus ESCON characteristics

FCP compared to ESCON	FCP	ESCON
PPRC primary port accepts host I/O.	Yes	No
PPRC secondary port accepts host I/O.	Yes	Yes
Link is full duplex (have paths established in both directions simultaneously).	Yes	No
Support synchronous PPRC.	Yes	Yes
Supports PPRC Extended Distance.	Yes	Yes
Number of secondary LSSs per primary LSS.	4	4
Number of PPRC logical paths per LSS.	8	8
Number of logical paths from a primary ESS.	1024	1024
Number of logical paths per port.	1024	64
Number of logical paths per physical path.	128	64
Number of physical paths per port.	126	64

4.8.2 Distance considerations

The distance between Fibre Channel ports without any extending infrastructure is dependent on the Fibre Channel port type. The ESS uses 2 Gbit GBICs on its Fibre Channel ports. For

short wave GBICs using 50 micron multimode fiber, the maximum distance is 200 meters. With 62.5 micron multimode fiber, the maximum distance is 250 meters. For long wave GBICs over single-mode fiber, the maximum distance is 10 km.

Figure 4-30 illustrates the Fibre Channel maximum distances when no networking devices are used to extend those distances.
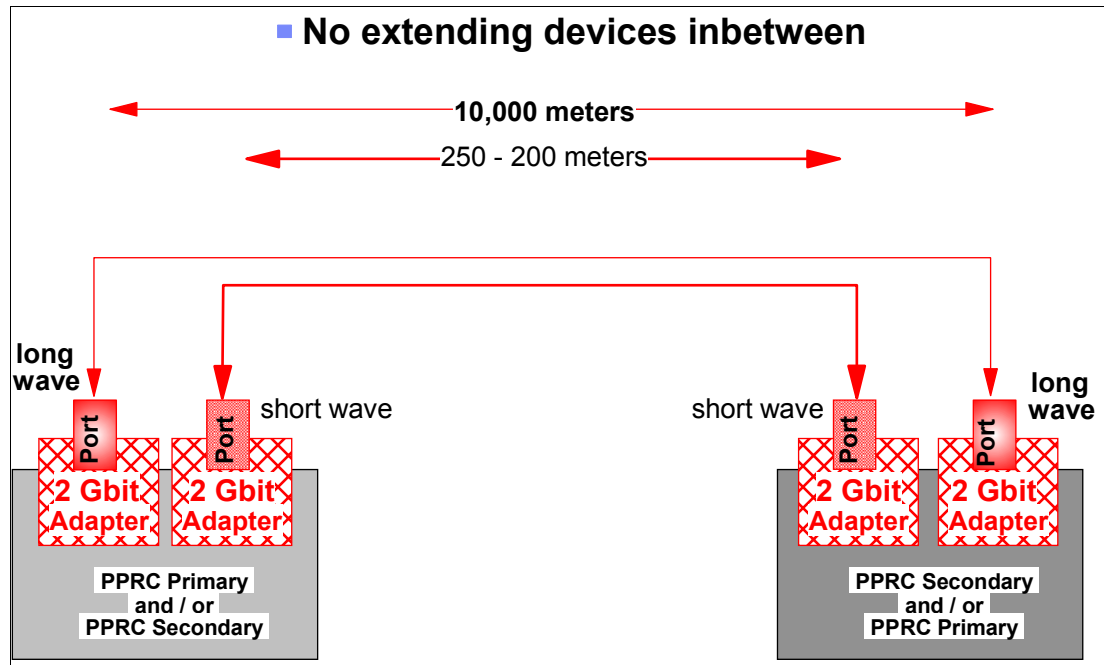
Figure 4-30   Fibre Channel distances without extending devices

4.8.3 SAN fabric and networking

In addition to Fibre Channel link direct connection, PPRC over Fibre Channel is also supported with the use of SAN fabric products and DWDMs. All Fibre Channel switches and directors supported by the ESS today are also supported for PPRC over Fibre Channel.

For extending the PPRC Fibre Channel distances, the following DWDMs are supported:

► CIENA CN 2000 Storage Extension Platform
► Cisco ONS 15530/15540
► CNT Edge Storage Router
► Nortel Networks OPTera Metro 5200/5300

A current list of supported environments, configurations, networks and products is available at:

http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

4.8.4 WWNN and WWPN

The WWNN (World Wide Node Name) is a single name given to an entity that is in a Fibre Channel network. A WWPN (World Wide Port Name) is the name of one of the ports into that entity. If, for example, you look at a Windows server with multiple HBAs, the code in the first adapter to get enabled will assign the WWNN of the server, and the second, third, and fourth adapters will all 'inherit' the WWNN but will represent themselves as the WWPN that is burned into the HBA.

In the case of ESS, there is a single WWNN of the ESS that is reported in the ESS Specialist Welcome panel (see Figure 4-31).
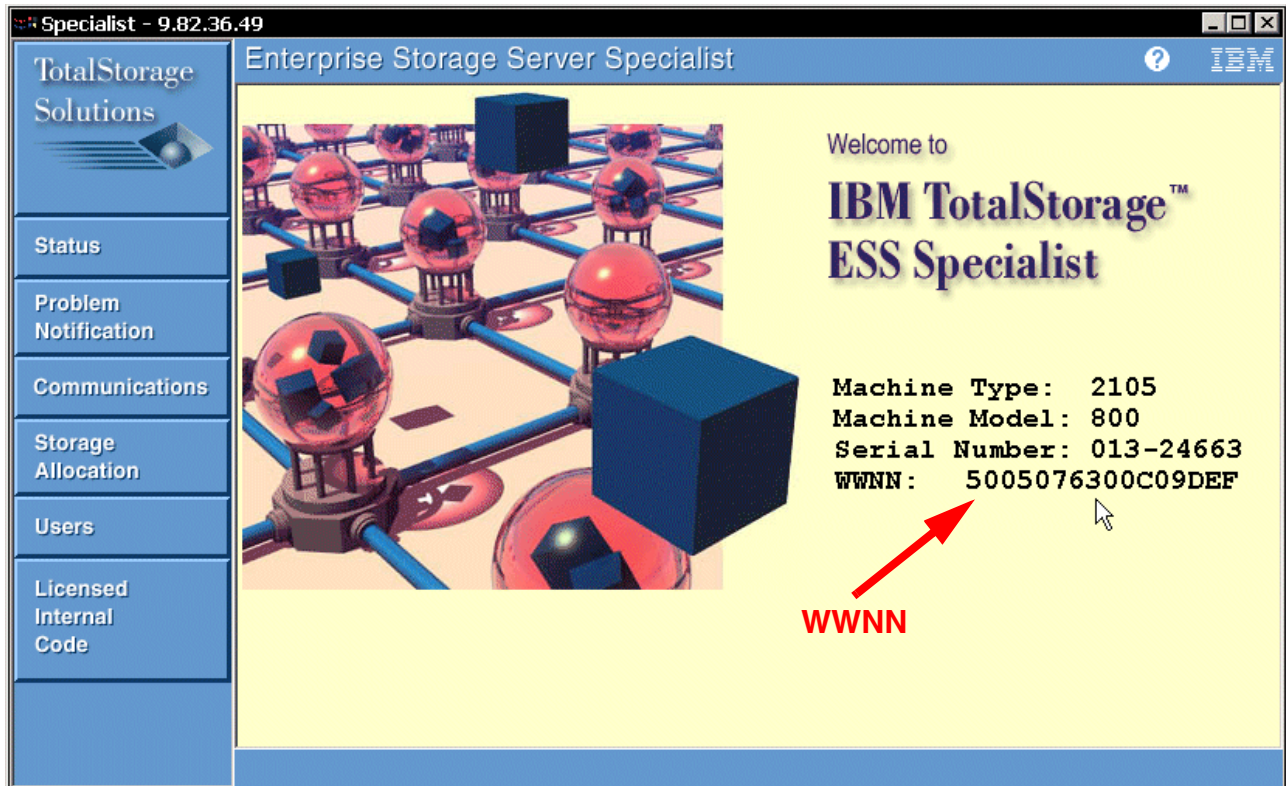
Figure 4-31   ESS WWNN information in the ESS Specialist Welcome panel

If you go into the ESS Specialist and click on each of the Fibre Channel host adapters (see Figure 4-32 on page 99), you will see that they all have a unique WWPN, but that the last four or six digits are all the same. In this case, instead of the burned-in WWPN, the ESS dynamically assigns the WWPN for each adapter. This way, if there is a need to replace a host adapter, then it will retain its previous identity in the Fibre network.
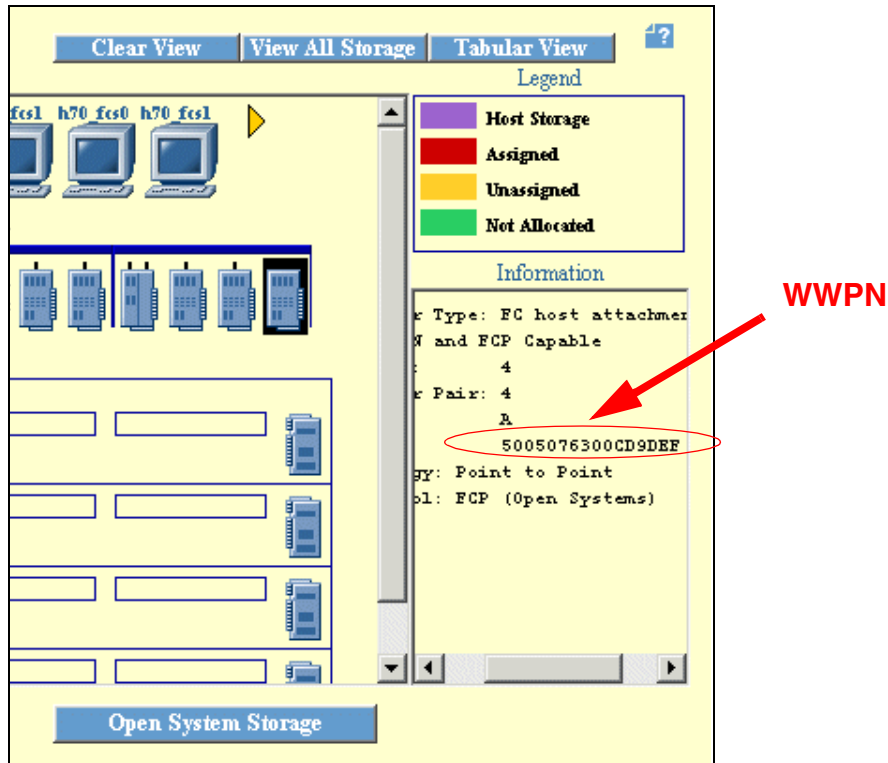
Figure 4-32   WWPN of ESS Fibre Channel Port 4-4-A

> **Note:** When looking at the Information Panel for a selected Fibre Channel port from the ESS Copy Services Web user interface Paths panel (see Figure 6-138 on page 276), you will see the WWNN and WWPN of the Fibre Channel switch to which the ESS FCP port is connected to.

4.8.5 Recognizing the ESS ports within the SAN fabric

In the ESS, there is a function called "locally administered WWPNs". The WWPN (World WIde Port Name) of a Fibre Channel port is calculated by a byte-level operation from the ESS node WWNN (World Wide Node Name, the one shown by the entry panel of ESS Specialist) and the adapter's port location in the ESS bay. Thus, the adapter's WWPN remains unchanged when the adapter is replaced.

A nice consequence of this characteristic is that you can easily recognize your ESS ports in the SAN: all Fibre Channel switch GUIs have some windows displaying the device's WWPNs attached to the fabric ports (with Brocade switches, it is the Name Server panel). You see something there similar to "IBM 2105 ... 50:05:07:63:00:c1:84:c8 ...". When you know how the WWPNs are identified in the ESS, you then can verify that the ESS ports are attached to the fabric ports as you want. This knowledge will also help you in other fabric management activities, for example, verifying that the correct WWPNs appear in zoning lists.

The way to determine how the FCP ports are assigned their WWPN in the ESS is as follows: take the ESS node WWNN (you can get if from the ESS Specialist Welcome panel, as shown in Figure 4-31 on page 98). Byte 6 in the WWNN is always c0 (in hexadecimal code). You get an FCP adapter's WWPN by replacing the x'c0' with the location code (c1...d0) for that adapter as per the diagram and example illustrated in Figure 4-33 on page 100.

```
ESS node or port id:          xx:xx:xx:xx:xx:yy:xx:xx
                                             ||
                                             c0 or port id

Example:

  ESS node WWNN:              50:05:07:63:00:c0:84:c8

  ESS port WWPN:             50:05:07:63:00:yy:84:c8
                                             ||
                                       replace Port id from diagram below


ESS Port ids:
+---------------+---------------+---------------+---------------+
| Adapter Bay 1 | Adapter Bay 2 | Adapter Bay 3 | Adapter Bay 4 |
|  C4 C3 C2 C1  |  CC CB CA C9  |  C8 C7 C6 C5  |  D0 CF CE CD  |
+---------------+---------------+---------------+---------------+

For the last FCP port in Bay 4, the WWPN will be:

  ESS port WWPN:              50:05:07:63:00:cd:84:c8
```
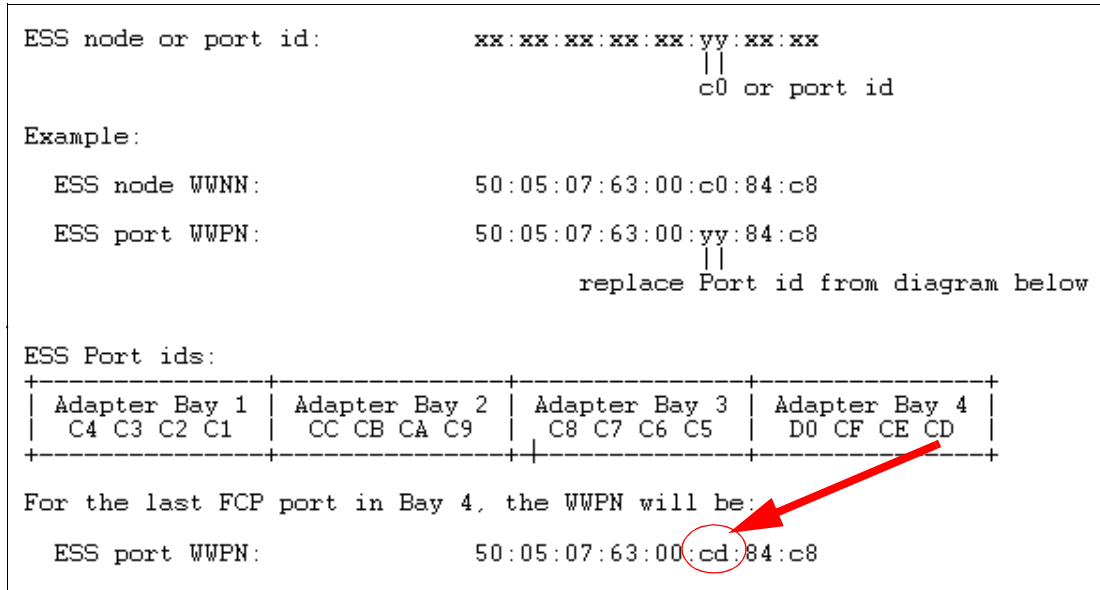
Figure 4-33   Determining the WWPN of an FCP port in the ESS

4.9 Practical examples of PPRC

This section outlines some of the practical uses of PPRC in both modes of operation, synchronous and extended distance.

4.9.1 Site migration

Site migration without remote copying facilities must be performed often using backups onto tape volumes and then restoring them onto a remote disk subsystem. Incremental backups need to be performed and applied regularly to the remote site. This method provides a PIT-like recovery scenario. When the time comes to do the actual site migration, a final backup and apply would need to be performed to get an up-to-date copy of the application data. The time elapsed for migration would be entirely dependent and elongated by the backup/apply process.

Site migrations can now be done with PPRC in a much shorter time with the minimum possible disruption to the production process. The application site volumes to be migrated can be receiving updates from the application, while keeping a copy at the remote site. For asynchronous mirroring using PPRC-XD, an appropriate application checkpoint needs to be defined to allow application writes to be quiesced for a very short time, to let the secondaries catch up.

Once the decision has been made to swap sites, all that is needed is to quiesce the application writes and allow them to propagate to the remote site. Once all the volumes are in duplex, the relationship can be terminated, and the application can be started again on the remote site. The amount of application downtime for this procedure depends on the availability of the application servers on the remote site.

- ■ Start Up PPRC-SYNC/XD
  - ► Establish all volumes
- ■ At beginning of data center cutover window
  - ► Quiesce applications
  - ► Allow PPRC-XD to catch up (time would be short)
  - ► Once PPRC-XD caught up, *data is ready at remote location*

Production Servers

Backup / Admin Servers

Significant reduction in work and effort required to migrate data centers over long distance

ESS 1: Production Data

B
L
X

C
M
Y

ESS 2: Mirrored Data

Figure 4-34   Site migration

4.9.2 Synchronous PPRC using static volumes

Remote data migration can be achieved with synchronous PPRC using static volumes. A static volume is a volume that is not receiving any write updates. This particular implementation allows the use of synchronous PPRC over long distances, beyond the supported 103 km, while taking advantage of the excellent throughput of PPRC when it does the initial copy or the re-synchronization of volume pairs (duplex pending state).

This powerful way of PPRC data copying, combined with the latest microcode improvements in the supported channel extenders, makes the static volume implementation ideal for data migration, data copying, or remote backup over long distances.

This PPRC implementation can also be used for disaster recovery solutions over long distances when:

► The applications in consideration can be quiesced.
► Solutions are based on switching database log files, and transmitting when inactive.

Additional information on this particular implementation can be found in the redbook *IBM TotalStorage Enterprise Sorage Server PPRC Extended Distance,* SG24-6568.

4.9.3 Database log transmission

Database recovery solutions based on switching active database logs and transmitting the then inactive log and the bootstrap data set (BSDS) can be efficiently implemented with PPRC. Once at the remote site, the log is applied onto the shadow database (see Figure 4-35 on page 102).
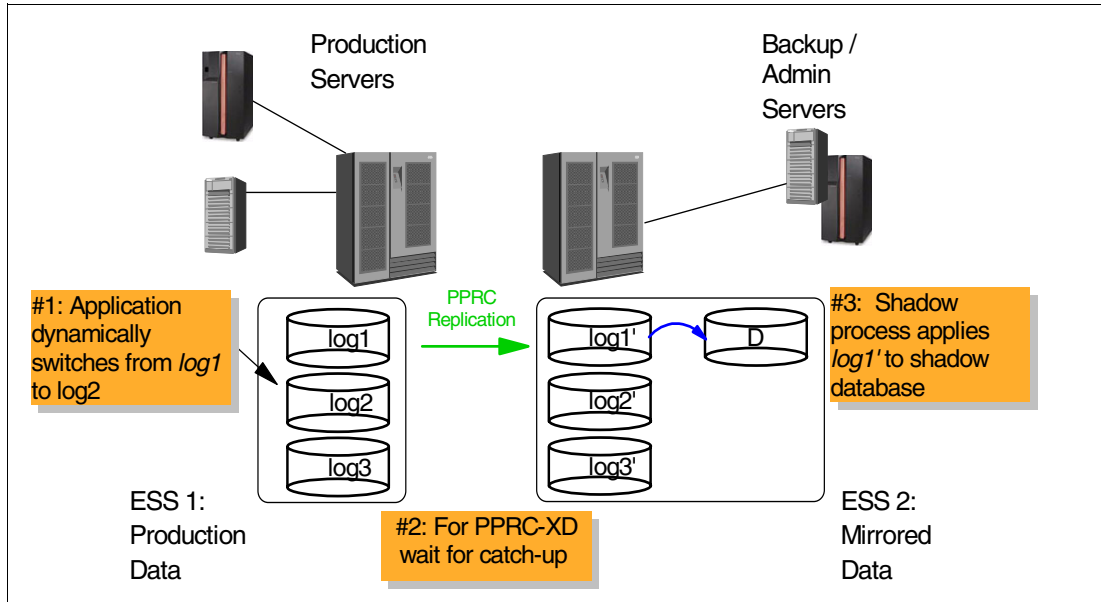
Figure 4-35   Database log transmission

4.9.4 Off-site backups

Using PPRC to perform off-site backups, there is no need for transporting media to the vault site. Off-site backups can be performed using split-mirror implementations using PPRC with FlashCopy (see Figure 4-35).

Sequence of operation

1. Establish volume pair A-B in PPRC-XD relationship
2. Run in duplex/duplex-pending XD mode until PiT consistent copy needed
3. When consistent copy needed, quiesce the application
4. If PPRC-XD, wait for catch-up and then freeze (or suspend)
5. With the pairs suspended, the application can resume
6. FlashCopy secondary B to tertiary C (no impact to host)
7. When FlashCopy is initiated, resume mirroring SYNC/XD
8. Go back to Step 2.

Figure 4-36   Split mirror implementation

5

ESS Copy Services Web User Interface prior to LIC 2.2.0

There are three different methods of using the ESS Copy Services in the open systems environment:

- ► A Web-based Interface
- ► A Java-based Command Line Interface (CLI)
- ► And, starting with LIC level 2.3.0, the ESS API has been enhanced to support Copy Services configuration and use for PPRC and FlashCopy

Very important to note is the fact that the Web-based user interface panels and icons have changed with LIC 2.2.0 first and later when LIC 2.3.0 became available.

In this chapter we explain how to use and set up the ESS Copy Services Web User Interface (WUI) when the ESS is at a LIC level prior to 2.2.0. The descriptions and panels correspond to what users will have if their LIC level is below 2.2.0.

The Web user interface for ESSs with LIC level 2.3.0 is discussed in Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169. The usage of the Command Line Interface is described in Chapter 7, "ESS Command Line Interface" on page 301. The ESS API is discussed in Appendix A, "ESS Application Programming Interface (API)" on page 471.

5.1 LIC level considerations within this chapter

All the procedures, screens and panels, PPRC options, and FlashCopy options described in this chapter refer to the LIC level Version prior to 2.2.0.

When you upgrade the ESS LIC level to 2.2.0 or higher:

► The Copy services WUI volume panel presentation will change.

► The definition of the Copy services domain will change.

► The FlashCopy will not change until you upgrade the FlashCopy feature to Version 2.

► The PPRC options will not change until you upgrade the PPRC feature to Version 2.

You may refer to Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169 to get more information about the new volume panel presentation and the new FlashCopy and PPRC options.

5.2 Overview and requirements

ESS Copy Services runs within the Enterprise Storage Server. When the ESS is running LIC prior to 2.2.0, one of the ESS clusters in the Copy Services domain has to be defined as the Primary Copy Services server and is responsible for maintaining all Copy Services related information. Optionally, you can define a second ESS cluster in the Copy Services domain to act as the Backup Copy Services server to provide availability of service. On each ESS cluster that is intended to use Copy Services, there is a Copy Services client running that communicates with the Active Copy Services server. For further information, refer to 2.1, "Copy Services terminology" on page 14.

Access to Copy Services is provided through a Web browser. Using a Web browser offers the possibility to easily control the ESS copy functionality over the network from any platform for which the browser is supported.

A Web server running in your ESS provides a Web interface that you can use to manage the ESS through a Web browser. The Web browser must be running on a workstation connected to the ESS through the IBM TotalStorage Enterprise Storage Server Network (*ESSNet*). The ESSNet is the ESS access facility that IBM installs when they install your ESS. The ESSNet consists of a switch (hub) and other networking components and the dedicated IBM TotalStorage Enterprise Storage Server Master Console (*ESS Master Console*).

The ESS Master Console replaces the ESSNet Console, the workstation included in some earlier versions of the ESSNet package. For additional information on the ESS Master Console, see the *IBM TotalStorage Enterprise Storage Server User's Guide*, SC26-7445, on the ESS Web site. For additional information on the ESSNet, see the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, GC26-7444, on the ESS Web site.

In addition to using a Web browser on the ESSNet Master Console to connect to your ESS, you can also use a Web browser running on your own workstation, either by connecting your workstation directly to the ESSNet hub, or by connecting your workstation to your intranet and connecting your intranet to the ESSNet hub. If you use your own workstation, IBM recommends that it have at least 128 MB of memory.

The ESS Copy Services require one of the following Internet browsers:

► Netscape Communicator
► Microsoft® Internet Explorer (MSIE)

For supported versions of Netscape Navigator and Internet Explorer, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

5.3 Using a browser to access the ESS

You can either use the ESS Master Console or a Web browser installed on a workstation connected through the ESSNet to access the ESS.

The ESS Master Console differs from other workstations in the way you access your ESS. The desktop of the ESS Master Console has a Netscape icon with the label ESS Specialist Launcher. That icon accesses a panel, which provides a list of IP addresses of your ESS clusters. Clicking on the address of the cluster that you wish to access will cause the ESS Launch panel resident in the selected cluster to open.

If you use a workstation other than the ESS Master Console to access your ESS, you must enter the desired cluster address in the Address field of your Web browser. The address that you click in the *Specialist Launcher* of your ESS Master Console or that you type in the Address field of your Web browser is the host name alias or the dotted decimal IP address of one of the clusters of your ESS. IBM configures these addresses in your ESS at installation time, based on the addresses you enter in the Communications Resources work sheet that you submit for your ESS. For details, see the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, GC26-7444.

The first page to open, the ESS Launch panel, is the same for either of the two clusters in the ESS, so access to either cluster allows you to configure the entire machine. The ESS Launch panel contains buttons that you click to access ESS Copy Services (Figure 5-1 on page 106). This connects the browser to the ESS that is specified as the Copy Services server. If you have not previously selected the **Copy Services** button, you will be prompted for the user name and the password before starting the Copy Services Web User Interface.
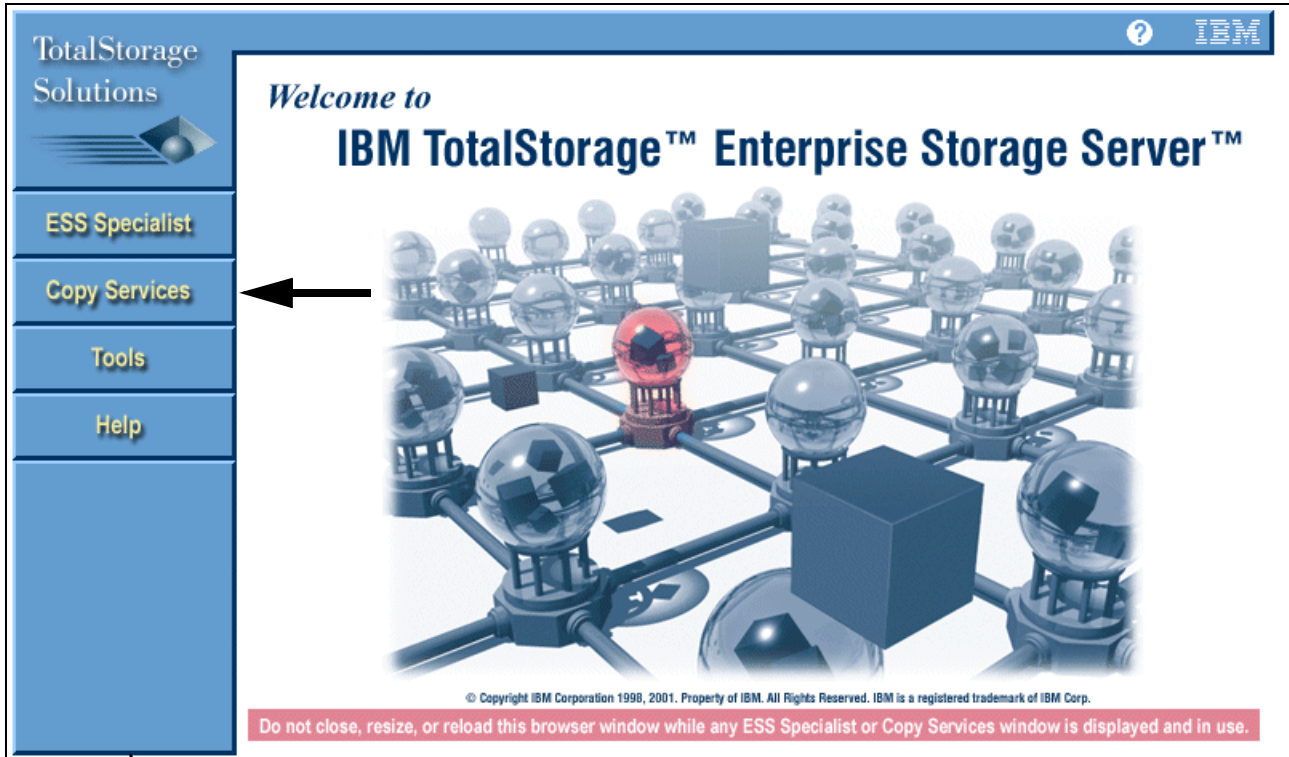
Figure 5-1   ESS launch panel

The message panel shown in Figure 5-2 will be displayed while connecting to the Copy Services Server.
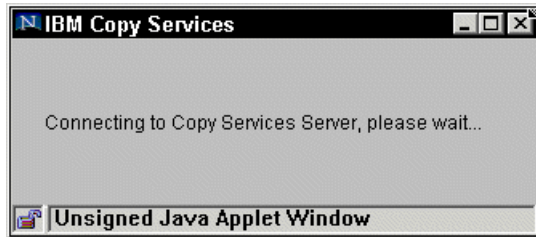
Figure 5-2   Copy Services start message

Once the connection to the Copy Services server is successful, the main menu of the *Copy Services Web User Interface* will be displayed (Figure 5-3 on page 107). From here, you can access all Copy Services menus by selecting one of the buttons on the left side.

Figure 5-3   Main menu of ESS Copy Services Web User Interface

5.3.1 Failure to connect to ESS Copy Services

If you receive the following message when you click Copy Services in the navigation frame of the ESS Launch page:

```
Failed to connect to Copy Services Server. Server may be down or not configured and cannot
access ESS Copy Services
```

the failure is most likely the result of at least one of the following three conditions:

► The Primary Copy Services server has not been defined. If you have successfully accessed ESS Copy Services before, this would not be the source of the problem. If this is the first time that you are connecting to ESS Copy Services, see 5.3.2, "Location of ESS Copy Services server" on page 108.

► The primary Copy Services server is not running. See 5.3.3, "Restarting ESS Copy Services" on page 108 for instructions.

► Your Web browser has a temporary operational or connection problem.

Note: For all three conditions noted above, if you have the CLI active on one of your host systems, you can verify that there is a connection to the active Copy Services server with the `rsTestConnection` command. For details, see 7.1.4, "Copy Services CLI Command description" on page 303.

For Web browser problems, take these corrective actions:

1. Close any open ESS Copy Services browser panel.

2. Access ESS Copy Services again by clicking **Copy Services**.

3. If step 2 is unsuccessful, close all browser panels (if you have ESS Specialist running, first complete any unfinished task definition), and then restart your browser.

4. Again, access the ESS Launch page and click **Copy Services**.

5. If Copy Services still does not launch, consider restarting the workstation that is running your browser.

5.3.2 Location of ESS Copy Services server

You can verify that the Copy Services server is defined by using the Tools panel. From the ESS Launch panel, click **Tools**. The Tools panel displays the topic "Tools: Copy Services Trouble Shooting and Disaster Recovery". To determine which Copy Services server is the active server, click the corresponding topic. Click **Find the active Copy Services server**. The Current Configuration Table is displayed (Figure 5-4).

Current Configuration:

IP Address	Description
9.43.232.236	Currently configured Primary Server.
9.43.232.222	Currently configured Backup Server.
9.43.232.222	Currently active Copy Services server.

NOTE: to refresh the above display, select the reload/refresh menu option from your browser's context menu for this browser frame (i.e. right-click in this frame and select *Refresh* or *Reload Frame* from the context menu).

Available Actions:

Action	Description
Reset to Primary	Restart Copy Services with the primary server as active server.
Reset to Backup	Restart Copy Services with the backup server as active server.
Disable	Disable Copy Services.
Cancel	Return to the Tools main page, without performing any action.

Figure 5-4   Configuration table of ESS Copy Services

Verify that there is an IP address displayed for Currently Configured Primary Server. If there is an IP address, the primary Copy Services server is defined. If not, the primary Copy Services server has not been defined.

5.3.3 Restarting ESS Copy Services

If ESS Copy Services is not available, you might want to restart ESS Copy Services. Be aware that you need to have administration authority to do this.

To restart Copy Services, click the **Tools** button (as in the previously described procedure) and the Tools panel displays. On that panel, review the message `Recovering from a Failure to Connect to the Copy Services Server` to link to the restart procedure. Be sure you understand the result of implementing a restart before proceeding. When you use this function, be aware that:

► You will lose any PPRC or FlashCopy tasks that have not completed.

► Established PPRC and FlashCopy relationships are maintained.

► You cannot submit any additional tasks through the CLI until ESS Copy Services has reinitialized.

► You must exit the Copy Services panel and re-enter it to receive valid updates from the Copy Services server.

If you decide to restart ESS Copy Services, click **Reset ESS Copy Services**. The message `Resetting the ESS Copy Services` appears.

Attention: Use this function only if you are certain that no other recovery options are available.

In the **Available Actions** section (Figure 5-4 on page 108), you can click **Reset to Primary** to restart Copy Services with the primary Copy Services server as the active Copy Services server.

Alternatively, you can click **Reset to Backup** to restart Copy Services with the Backup Copy Services server as the active Copy Services server. (For an explanation of these terms, see 2.1, "Copy Services terminology" on page 14.).

Note: You would take this action if the primary Copy Services server was unavailable because of a disaster or emergency situation. For more information, see 10.5, "Copy Services server failover" on page 406.

5.3.4 How to switch to the Backup server

To test your disaster recovery process, you can create a simulation of a disaster at your production site and restart the Copy Services Server at the backup site.

Let us assume that you have two ESSs: ESS1 and ESS2. The primary Copy Services Server is defined on ESS1, and the Backup Copy Services server is defined on ESS2. Perform the following steps to transfer control of Copy Services from the primary to the Backup server:

1. On the ESS launch panel, click **Tools**. The Tools panel opens, as shown in Figure 5-5.
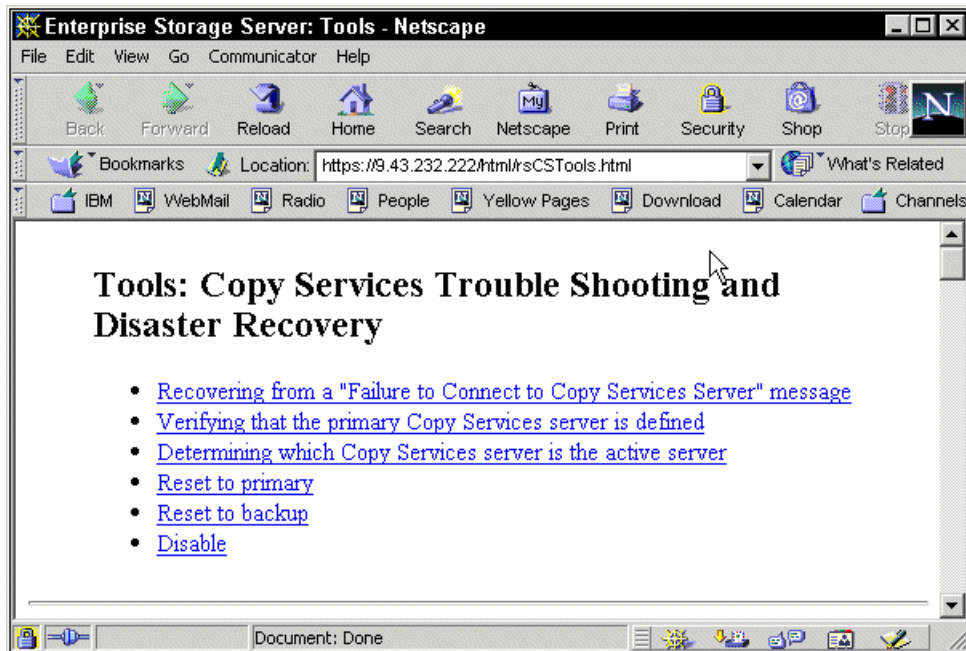
Figure 5-5   The Tools panel

2. On the Tools panel, click **Reset do backup**.

3. Within the **Reset to backup** section, click **Reset ESS Copy Services** and the Configuration Table of ESS Copy Services displays (Figure 5-6). This will show the Primary Copy Services server, the Backup Copy Services server, and the current Active Copy Services server.

Current Configuration:

IP Address	Description
9.43.232.236	Currently configured Primary Server.
9.43.232.222	Currently configured Backup Server.
9.43.232.222	Currently active Copy Services server.

NOTE: to refresh the above display, select the reload/refresh menu option from your browser's context menu for this browser frame (i.e. right-click in this frame and select *Refresh* or *Reload Frame* from the context menu).

Available Actions:

Action	Description
Reset to Primary	Restart Copy Services with the primary server as active server.
Reset to Backup	Restart Copy Services with the backup server as active server.
Disable	Disable Copy Services.
Cancel	Return to the Tools main page, without performing any action.

Figure 5-6   Configuration table of ESS Copy Services

4. Click **Disable** to deactivate the current Active Copy Services server.

5. Click **Reset to Backup** to make the Backup Copy Services server at the recovery site the Active Copy Services server.

> **Attention:** You can create and save tasks when the active server is the Backup server, but these tasks are lost when you perform the Reset-to-Primary function, which re-establishes the primary server as the active server.

5.3.5 Steps following a disaster at the production site

If a disaster occurs at the production site and your Primary Copy Services server is running on this site, you must invoke ESS at the recovery site as the active Copy Services server. For the following procedure to work, you must have created the proper recovery tasks (see 10.5, "Copy Services server failover" on page 406).

To switch to the recovery site, establish a Web browser connection to the ESS at the recovery site and perform the steps in 5.3.4, "How to switch to the Backup server" on page 109 to enable the Backup server as the active Copy Services server.

Open the Copy Services Web User Interface and terminate all PPRC pairs across all ESSs. Direct the commands to the secondary volumes, because the primary LSSs may not be accessible. See 5.11.8, "Terminating a PPRC copy pair" on page 157 for instructions on terminating pairs.

Perform any necessary host specific tasks required to revise the operating system configuration with the data held on the secondary volumes. Then bring up the applications on the disaster servers.

Once the production site has been restored, establish PPRC paths from the recovery site to the production site and copy the entire contents of a PPRC volumes at the recovery site to the volumes at the production site.

Once all the PPRC volumes are in duplex state and you have decided to switch back to the production site, perform this procedure in the reverse direction.

5.4 Volumes panel

Volumes are defined with the ESS Specialist in order to provide a fixed storage capacity to the connected host system. They are the base components of each data copy task. The ESS assigns each volume a unique eight digit identifier (ID). This identifier is used to address each volume within the ESS.

From the Volumes menu, you will be able to:

► Establish and withdraw FlashCopy pairs.
► Establish, suspend, and terminate synchronous PPRC copy pairs.
► Establish, suspend, and terminate PPRC Extended Distance copy pairs.
► Convert a PPRC Extended Distance copy pair to synchronous PPRC.
► Establish multiple volume-pair relationships for FlashCopy and PPRC.
► Find a volume.
► Display volumes based on a filter.
► View information about a volume.

Figure 5-7 shows the entry panel of the Volumes panel.
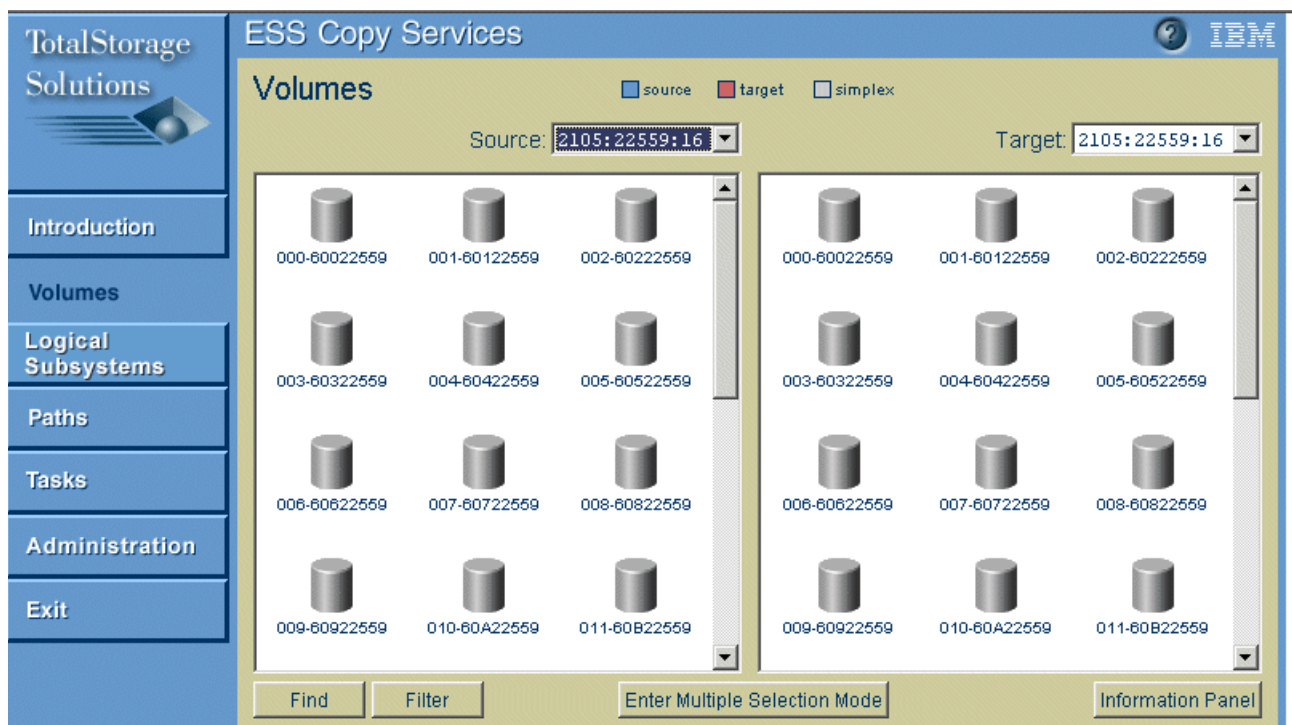
Figure 5-7   Volumes panel

The source and target logical subsystem are specified in the following way: Device type (4 digits):ESS Serial number (5 digits):LSS number (2 digits). An example would be a logical subsystem that is addressed by 2105:2259:16.

The Volumes menu shows all volumes defined within the LSS. Below each volume you will find its unique serial number. The volume icon indicates if it is used in any copy relationship (source or target), or if it is not part of a copy pair at all.

Table 5-1 explains the meanings of the volume icons that appear on the Volumes panel.

Table 5-1   Meaning of icons on Volume panel

Volume icon	Appearance	Meaning
	Gray solid (simplex volume)	The volume is in simplex state, meaning that it is not in a copy pair created by FlashCopy or PPRC
	Gray over blue Primary (source) volume	Synchronous PPRC, initial copy in progress (copying entire volume).
	Gray over red secondary (target) volume	
	Blue solid Primary (source) volume	Synchronous PPRC volume in a non-suspended state.
	Red solid Secondary (target) volume	
	Blue wedge in gray cylinder Primary (source) volume	PPRC Extended Distance volume in a non-suspended state.
	Red wedge in a gray cylinder Secondary (target) volume	

Volume icon	Appearance	Meaning
	Blue and white stripes Primary (source) volume	PPRC (synchronous or Extended Distance) in a suspended state.
	Red and white stripes Secondary (target) volume	
	Blue lightning bolt on gray cylinder Source volume	FlashCopy in progress. The FlashCopy icon on the target (the small white square) indicates that this is a duplicated volume.
	Red lightning bolt on gray cylinder Target volume	
	White square on gray cylinder I-am-a-target Icon (FlashCopy indicator)	After a successful FlashCopy, the FlashCopy icon remains until you issue a Withdraw FlashCopy command to the target volume (see 5.10.3, "Withdrawing a FlashCopy pair" on page 135) or until you use the volume as a source.

5.4.1 Volume Information Panel button

You can get more detailed information about a single volume by selecting the volume and clicking the **Information Panel** button, as shown in Figure 5-8.
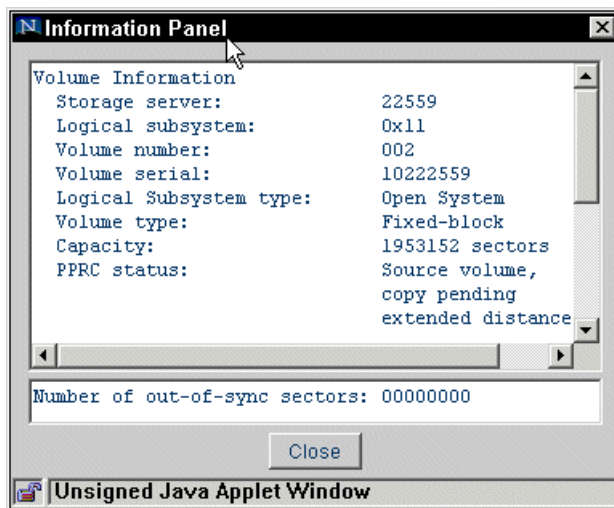
Figure 5-8   Volume Information Panel

The following information is displayed for a volume:

► The storage server (ESS) and logical subsystem (LSS) designations

► The volume number

► The volume serial number

► The LSS type

► The volume type (fixed block for open system; 9337, 3390, and other types for S/390)

► The volume capacity

► Peer status (indicates whether the volume is a source or target volume and the copy mode)

► Any active FlashCopy operations

► Whether Read from Secondary is enabled

► Information about a companion volume if a copy pair exists

5.4.2 Finding volumes

With the **Find** button, you can search for a specific volume. The volume is specified with its eight digit ID. See Figure 5-9.
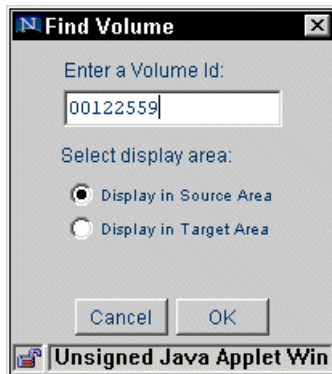
Figure 5-9   Find Volume panel

If the desired volume is found, its designator is displayed in the upper-left hand corner. All other volumes in the same LSS are also displayed.

5.4.3 Filtering volumes

In addition, you can filter the output of the volume display to a selected range by clicking the **Filter** button and selection the **Filter volumes** option.

In the example shown in Figure 5-10 on page 115, we want to display Open System volumes only. The volumes should be in simplex state; that means they are currently not in a copy relationship.
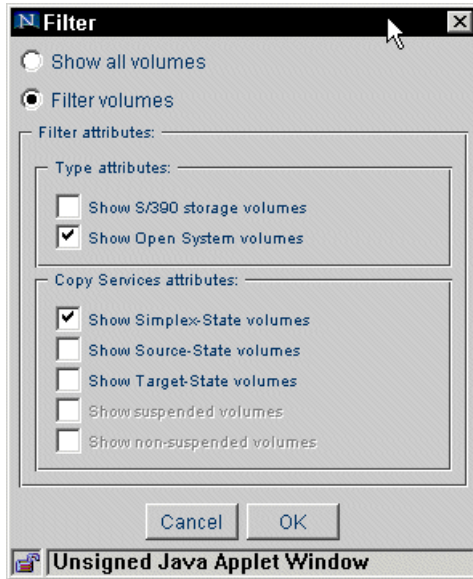
Figure 5-10   Filter volumes panel

5.4.4 Multiple Selection Mode

When you enter the **Multiple Selection Mode**, you can select multiple volume pairs when you create tasks for both FlashCopy and PPRC. This mode is very useful to save time when creating the tasks.

You might use this mode when selecting the volume for FlashCopy and PPRC tasks (for further information, refer to 5.10.1, "Establishing a FlashCopy pair" on page 130, 5.11.3, "Establishing a synchronous PPRC pair" on page 147, and 5.11.4, "Establishing a PPRC Extended Distance (PPRC-XD) copy pair" on page 153). Rather than defining one copy pair at a time, you can define all the copy pairs at one time by selecting multiple sources and their targets in the same task definition.

To enter the **Multiple Selection Mode**, click the **Multiple Selection Mode** button at the bottom of the panel. This allows multiple pairs to be defined. Select the desired source volume, and right-click the desired target volume. Repeat the selection of source and target volumes until you have defined all volume pairs. Right-click one of the target volumes again and the Task Wizard opens. You can decide whether to run one of the following operations:

► Establish synchronous PPRC copy pairs
► Establish Extended Distance PPRC copy pairs
► Suspend PPRC copy pairs
► Terminate FlashCopy copy pairs
► Establish FlashCopy copy pairs
► Withdraw FlashCopy pairs

To exit the **Multiple Selection Mode**, click the **Exit Multiple Selection Mode** button at the bottom of the Volume Panel.

Restriction: You can only perform the multiple selections on volumes within the same LSS (all sources must be in the same LSS and all targets must be in the same LSS).

Attention: If you click **Exit Multiple Selection Mode** before you perform any operations on them, you lose all your previous selections.

5.5 Logical Subsystems panel

The **Logical Subsystems** panel displays the Enterprise Storage Servers and the logical subsystems (LSSs) within the storage network. The storage network includes all Enterprise Storage Servers that are configured to use the same Copy Services Server. Each of the logical subsystems is specified by the serial number of the ESS it belongs to and its 2-digit LSS number within the ESS.

With the **Logical Subsystems** panel, you will be able to:

► Establish PPRC copy pairs for LSSs.

► Convert PPRC Extended Distance copy pairs to synchronous pairs for all the volumes in a source LSS and target LSS.

► Filter a list of LSSs.

► Find an LSS.

► Freeze a PPRC consistency group.

► Modify the PPRC consistency group time-out value.

► Remove orphaned paths.

► Resynchronize PPRC copy pairs for LSSs.

► Run a consistency-group-created operation.

► Suspend PPRC copy pairs for LSSs.

► Terminate PPRC copy pairs for LSSs.

► View information about an LSS.

In Figure 5-11 on page 117, you can see the **Logical Subsystems** panel for a selected ESS. The color indicates the state of the LSS, whether it contains volumes that are currently in a copy relationship (source, target, or mixed), or that they are not part of a copy pair at all.
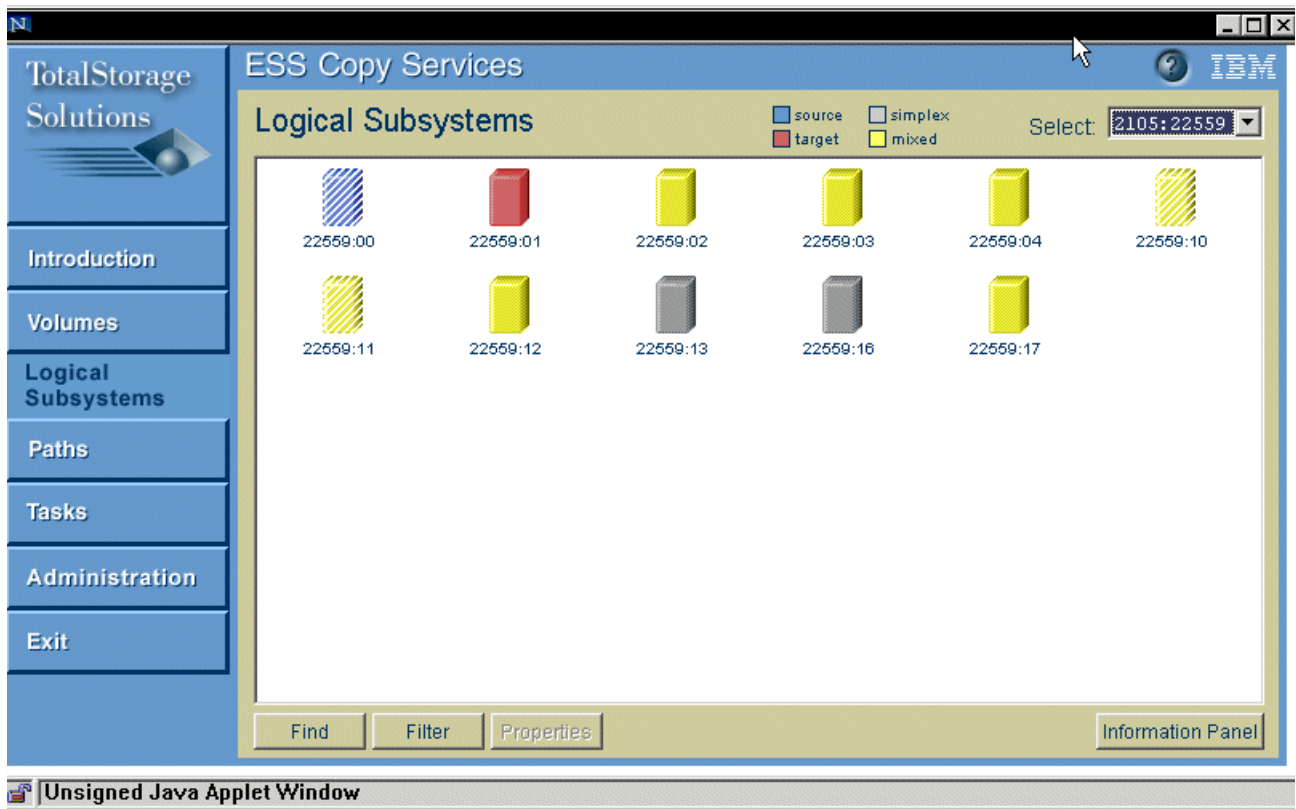
Figure 5-11   Logical Subsystem panel

Figure 5-2 explains the meanings of the icons that appear on the Logical Subsystems panel. The Volumes panel displays similar icons for individual volumes (see Table 5-1 on page 112).

Table 5-2   Meaning of icons on Logical Subsystems panel

LSS Icon	Appearance	Meaning
	Gray solid ("Simplex" state)	All the volumes in the LSS are in simplex state, meaning that they are not in a copy relationship.
	Blue solid (PPRC sync. source)	All the volumes in the LSS are in PPRC pairs (synchronous or Extended Distance) in a non-suspended state.
	Red solid (PPRC sync. target)	

LSS Icon	Appearance	Meaning
	Blue and white stripes (PPRC suspended source)	All the volumes in the LSS are in PPRC pairs (synchronous or Extended Distance) in a suspended state.
	Red and white stripes (PPRC suspended target)	
	Yellow solid	Not all volumes in the LSS are in the same type of Copy Services relationship (this apply for PPRC and FlashCopy relationships), and none of the volumes are in a suspended state.
	Yellow and white stripes (mixed types and states)	The volumes in the LSS are either the same type but some are in a non-suspended state or not all of the volumes are the same type but some are in suspended state.

5.5.1 LSS Information Panel button

You can get more detailed information about a single logical subsystem by selecting the LSS and clicking on the **Information Panel** button. In our example in Figure 5-12 we have selected LSS 16, which is specified for open systems and contains 25 volumes. None of these volumes are currently part of a FlashCopy or PPRC pair.
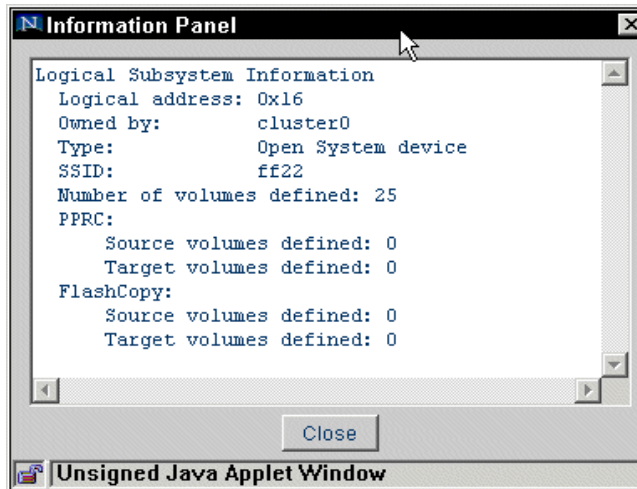
Figure 5-12   LSS Information panel

5.5.2 Finding LSSs

You can search for a specific LSS based on its address by selecting the **Find** button of the **Logical Subsystems** panel. The LSS must be referred to using the following syntax: ESS serial number:LSS number.

Figure 5-13 on page 119 shows an example in which we want to find the logical subsystem 12 of the Enterprise Storage Server with the Serial Number 22559.
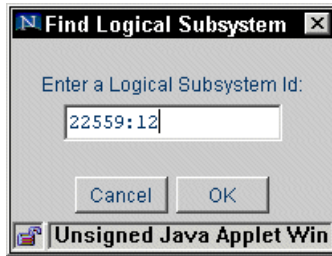
Figure 5-13   Find LSS

5.5.3 Filtering LSSs

In addition, you can limit the output of the volume display to a selected range by clicking the **Filter** button and selecting the **Filter devices** option.

In our example (Figure 5-14), we want to display physical and logical storage servers only. In addition, we only want to display open systems devices that contain volumes that are currently part of a Copy relationship.
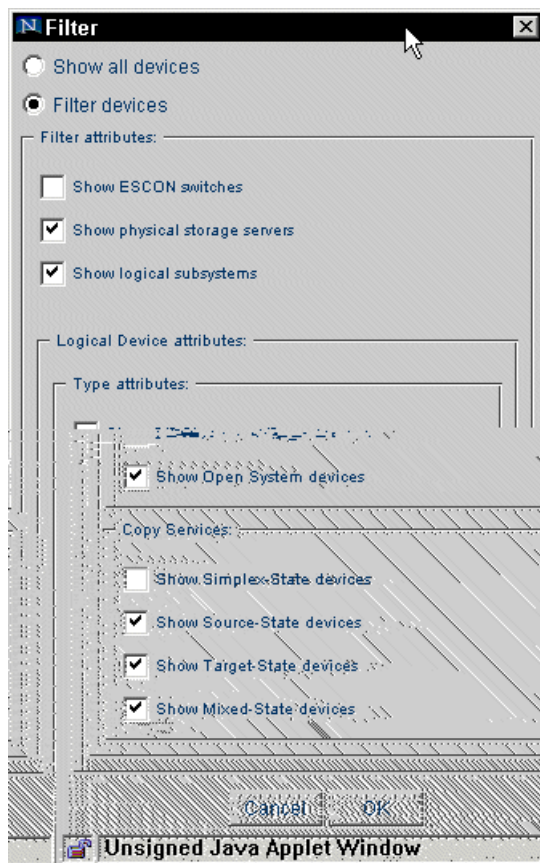
Figure 5-14   Filtering logical subsystems

5.5.4 Properties

By selecting one LSS and clicking the **Properties** button, you can view or change the copy properties of the entire LSS (see Figure 5-15 on page 120).

Figure 5-15   LSS properties panel

PPRC consistency group

To maintain consistency of data across volumes at the recovery site, volumes in a PPRC relationship can be collected into a *consistency group*. If you mark this check box, than all the volumes in that LSS will be considered as part of one Consistency group. Creating a consistency group of volumes provides the ability to temporarily queue subsequent write operations to all PPRC consistency group volumes on an single LSS pairing when an error occurs on one of the volumes in the group (primary or secondary), or when a total link failure is detected between the primary and secondary LSS pair. The other way to establish a consistency group is to select the PPRC Consistency group option when establishing a path between two LSSs. For further information about PPRC Consistency groups, refer to 4.5, "PPRC data consistency" on page 72.

XRC Session Time Out

This parameter does not apply to open systems.

CONCOPY Session Time Out

This parameter does not apply to open systems.

PPRC Consistency Group Time Out

This parameter indicates the amount of time that an I/O is withheld from updating a primary volume of a consistency group in an case of an error event. This timeout enables automation software to detect that an error has occurred and to issue commands to *freeze* all other members of the consistency group. When the consistency group is created, this parameter is set to a default of two minutes. For further information about PPRC consistency groups, refer to 4.5, "PPRC data consistency" on page 72.

5.6 The Paths panel

A path is used to send data between the primary and secondary volumes of PPRC pairs. The physical path consists of the ESCON connection between two Enterprise Storage Servers, while a logical path describes the connection of the PPRC source and targets.

Be aware that before you can create PPRC copy pair relationships, you must use the Paths panel to establish logical paths between logical subsystems between the source and the target LSS. After you establish the paths, the Paths panel displays the current status of the

paths for the selected LSS. There could be multiple logical paths established over a single physical path. This is recommended for availability and performance reasons. For further information, refer to 4.7, "PPRC connectivity" on page 82.

From the Paths panel, you will be able to:

- ► Establish paths.
- ► Add paths.
- ► Remove a group of established paths.
- ► Remove one or more paths from a group of established paths.
- ► Create a PPRC consistency group.
- ► View information about paths.

Figure 5-16 shows the entry panel of the Paths panel.
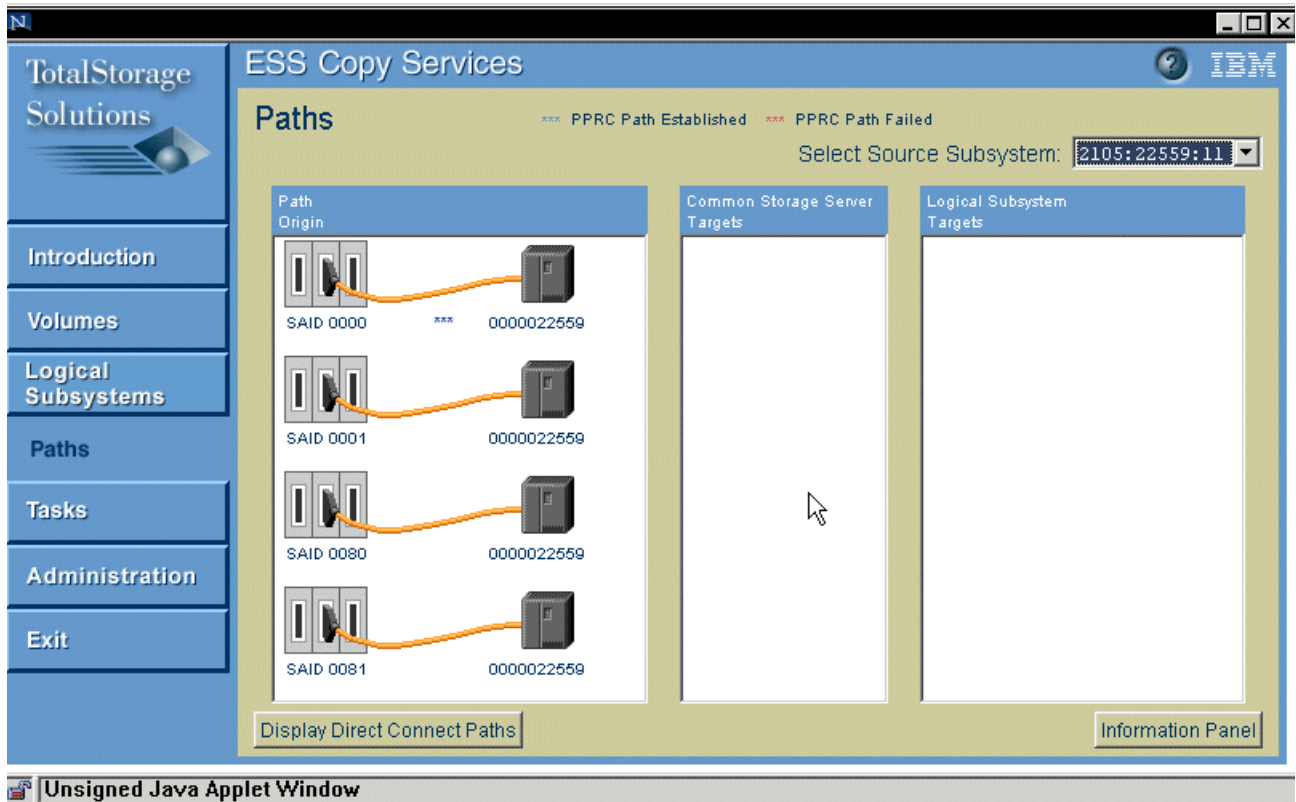
Figure 5-16   The Paths panel

In the upper right corner of the path panel from the drop-down menu, **Select Source Subsystem**, you select the primary LSS, which is denoted by an ESS serial number and LSS number. In this example, we use LSS 11 of the ESS with the serial number 22559. This shows you all configured ESCON adapters of the primary ESS in the Path Origin column of the Paths panel and their targets. Table 5-3 on page 122 gives you an overview about the connection symbols used in the "Local port / remote port" column located on the left side of the panel.

Table 5-3   Meaning of the Paths panel icons

Connection icon	Description
	Host adapter port, with the system adapter identification number (SAID) provided under it.
	ESCON director.
	Host server.
	Storage server.

The ESCON adapters are specified by their System Adapter ID (SAID). Figure 5-17 shows the SAID of all ESS ESCON adapters (see 4.7, "PPRC connectivity" on page 82 for further information).
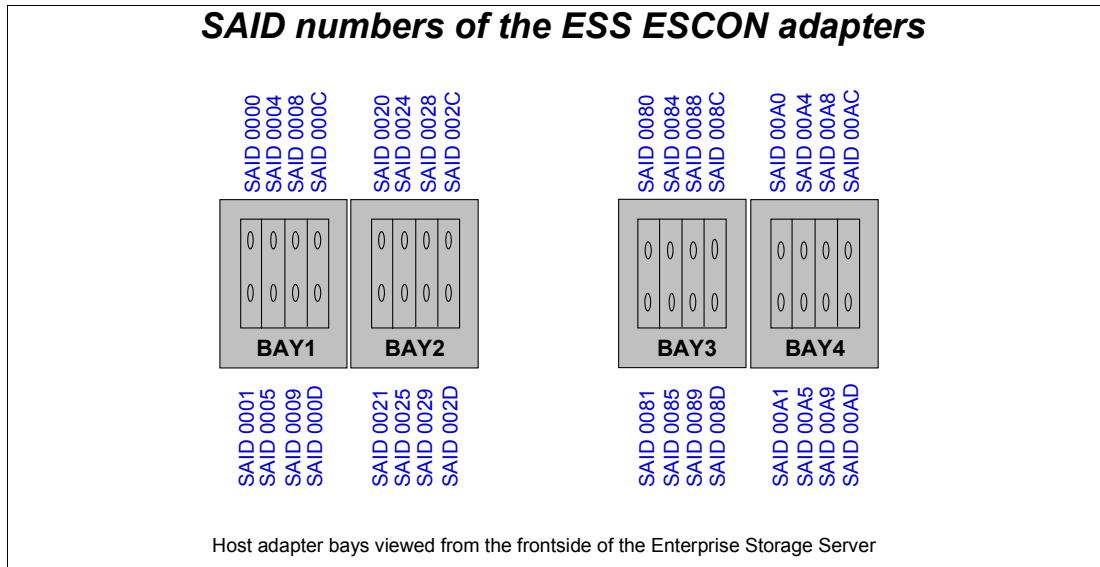
Figure 5-17   SAID numbers of the ESS ESCON adapters

Once an ESCON adapter is selected, all Enterprise Storage servers that are connected to this adapter are displayed in the Common Storage Server Targets column. All logical subsystems that are available on a particular ESS will be listed in the Logical Subsystem Targets column if one of the Storage Servers is selected.

In the example shown in Figure 5-18, we have selected the ESS 2105:14146:17, that is, the LSS number 17 of the ESS with serial number 14146. The path origin is the ESCON adapter with the SAID 0200. The selected target is ESS 2105:13350.
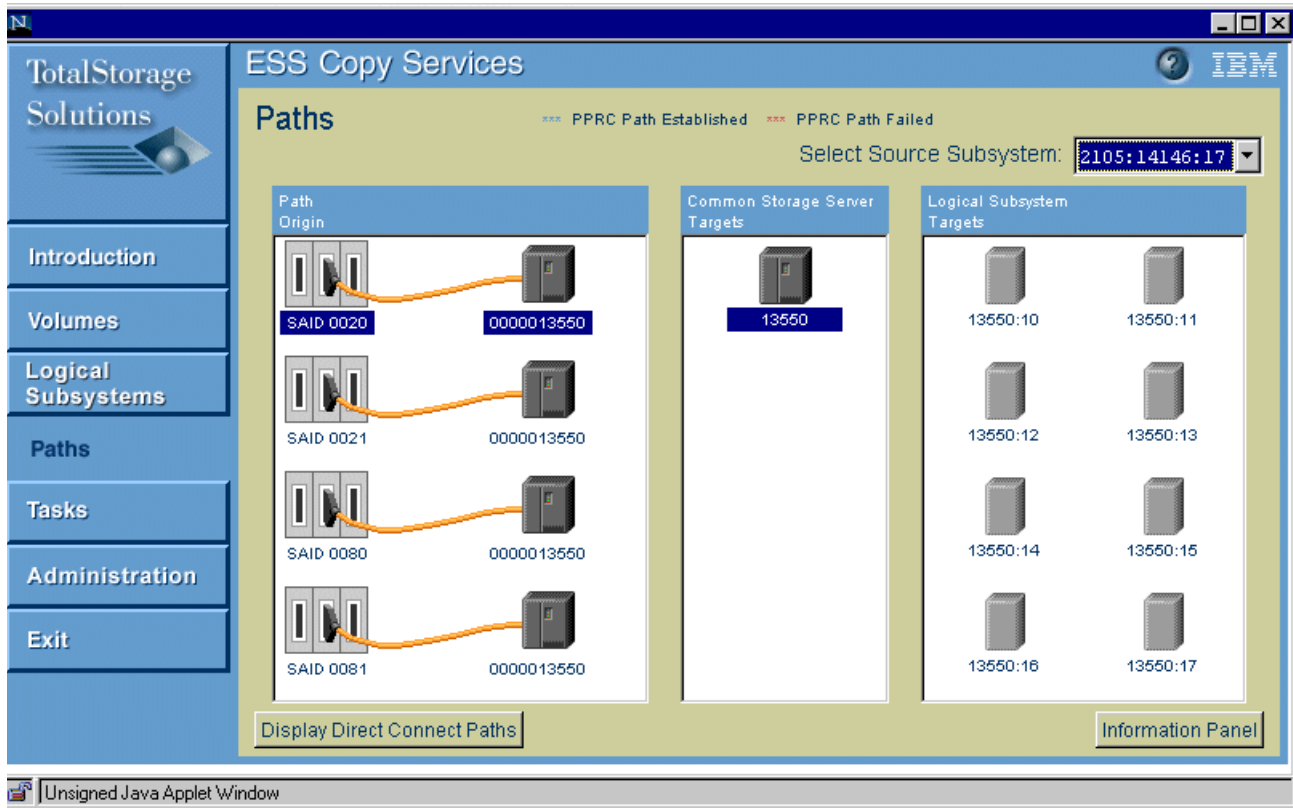
Figure 5-18   Example of the Paths panel

5.6.1 Display Direct Connection Paths

If you have two ESSs that are directly connected (that is, no switch between them), you can click **Display Direct Connect Paths** to display those paths.

> **Note:** You must perform a refresh function for the following conditions:
> - ▶ If this is the very first time anyone is accessing the ESS Copy Services WUI.
> - ▶ After an update of the ESS microcode.
> - ▶ After a concurrent service action.
> - ▶ If the cabling for the ESCON links is modified.

5.6.2 Path Information Panel button

Once an ESCON adapter is selected, you can get more information about the paths by clicking the **Information Panel** button at the bottom of the Paths panel.

The example shown in Figure 5-19 on page 124 shows a path defined between source LSS 10 and target LSS 11 using the ESCON adapters SAID0000 and remote SAID0080 on the same ESS.
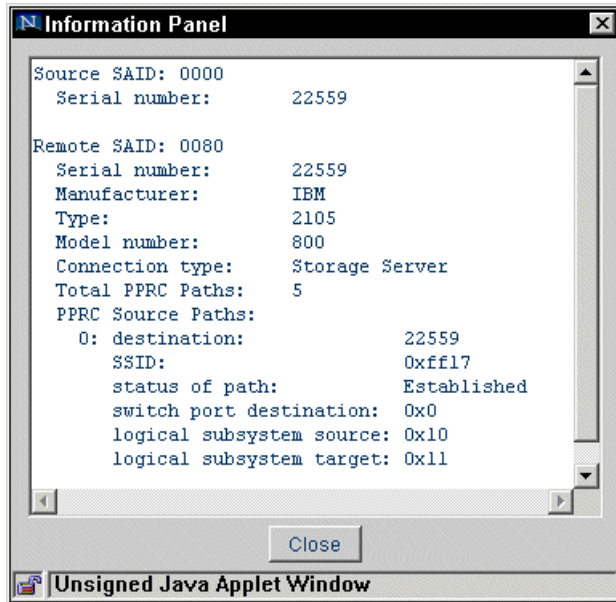
Figure 5-19   Path Information panel

If there are logical paths defined on an ESCON adapter, you will find three blue asterisks right below the adapter in the Path Origin column. Three red asterisks below the path-connection symbol in the Path Origin column means that the last attempt to establish the path failed.

Figure 5-20 shows an ESCON adapter SAID without, then with, the defined path, and finally with a failure to establish the path.
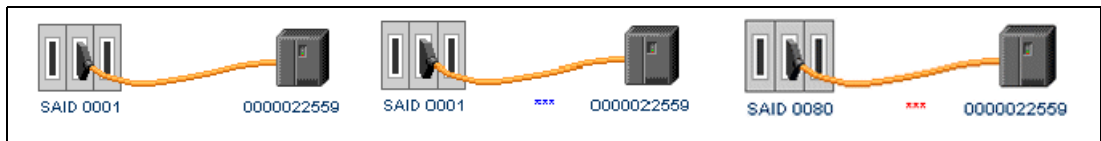
Figure 5-20   ESCON adapter without, with defined path, and failed establish path

5.7 Tasks panel

With the ESS Copy Services, you have the possibility to save the setup of any data copy action within a Task. This could be any kind of FlashCopy, PPRC, and path operation.

With the Tasks menu, you will be able to:

► Group tasks.
► Modify a task.
► Remove a task group.
► Remove a saved task.
► Run a saved task.
► Ungroup tasks.
► View error information about a failed task.
► View information about a saved task.

Figure 5-21 on page 125 shows the Tasks panel of the ESS Copy Services. For each task, the name, a description, and the last status of the execution is displayed.

You will be prompted to choose a task name when you create a task. Use the following guidelines:

► The name can only contain alphanumeric characters.
► The name cannot contain blanks.
► The name can contain the underscore (_) or hyphen (-) characters.
► The name cannot exceed 16 characters.

Note: In Appendix E, "Sample naming convention" on page 499, we present the task naming convention used in this chapter. You can refer to this appendix to find some hints and tips about naming conventions.
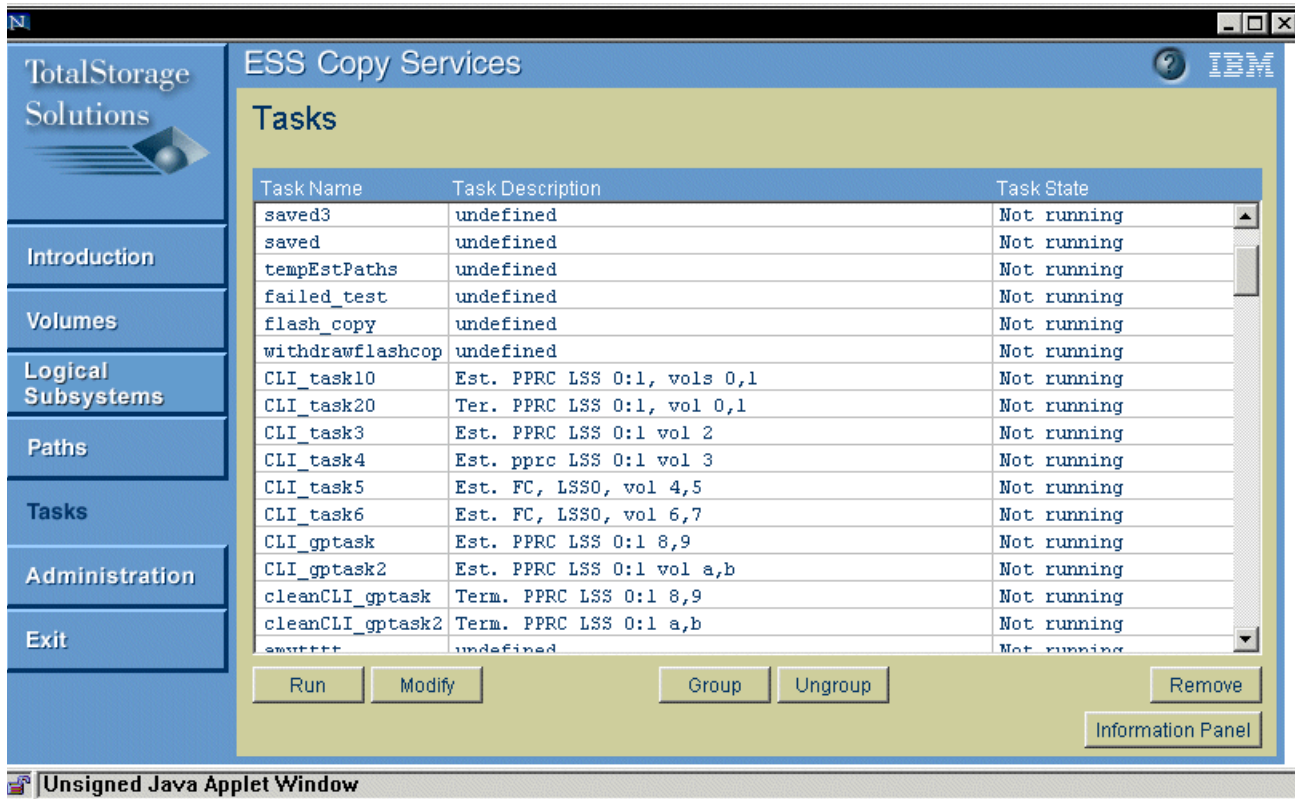
Figure 5-21   The Tasks panel

5.7.1 Grouping and ungrouping tasks

To create a group task, click the single tasks you want to group together while holding the Shift key or the Control key (mark single tasks). Once you are finished, click the **Group** button and specify the group name. It is not possible to include a group into another task group, all tasks within a task group will be processed in parallel.

An example for the usage of a task group would be multiple FlashCopy pairs from different logical subsystems that need to be issued all at the same time in order to do a backup.

In our example in Figure 5-22 on page 126, we have a grouped task named Group_task. This group contains two single tasks named CLI_task5 and CLI_task6. Both of the tasks establish a FlashCopy pair within LSS 00.

You can get detailed information about the setup of the task. Select the task and click the **Information Panel** button at the lower right.

Grouping tasks will cause the individual tasks to collapse into the overall group task.
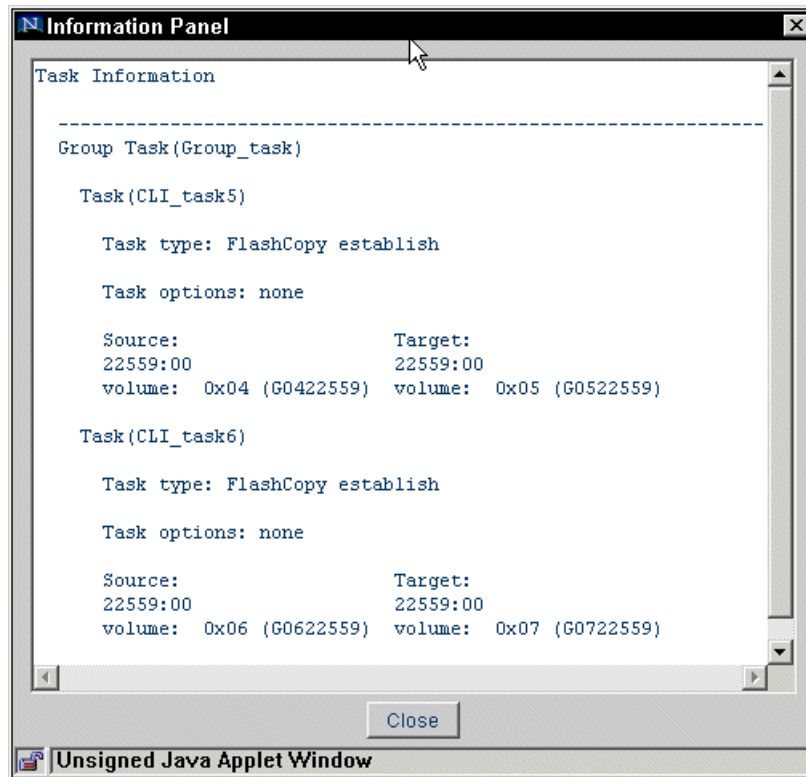
Figure 5-22   Task information panel

To ungroup a task, select the task and click the **Ungroup** button at the bottom of the Tasks panel. It is necessary to ungroup the overall group task in order to gain access and to execute the individual tasks.

5.7.2 Removing a task

To remove a task, select the task and click the **Remove** button at the bottom of the Tasks panel.

5.7.3 Running a task and viewing an error about a failed task

To run a saved task, select the task and click the **Run** button at the bottom of the Tasks panel. The task will be executed immediately.

If the task fails, a failure report is available through the Information panel. The failure report contains the error messages and the sense data. The messages are documented in the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448, which is available on the ESS Web site.

You can also view the logs in the administration panel.

5.7.4 Modifying a task

Use the Tasks panel to modify a previously saved task. You can change the task goals, task options, task name and description.

You will be prompted if you want to run, to replace the selected task, or create a new task.

Use the Tasks panel to modify a previously saved task. You can change a task from one that establishes a PPRC copy pair to one that suspends, re-synchronizes, or terminates a PPRC copy pair. For FlashCopy, you can change a task that establishes a FlashCopy pair to one that withdraws a FlashCopy pair. Or you can change the options of one task.

Note: Grouped tasks cannot be modified. The tasks must be ungrouped and then the tasks can be modified.

5.8 Administration panel

Use the Administration panel (see Figure 5-23) to manage the server logs, reports, and ESS Copy Services Command Line Interface (CLI) user IDs (user IDs) and passwords. You can also use this panel to refresh the volume and LSS information for an ESS.

Figure 5-23   Administration panel

You can perform the following actions from the Administration panel:

► Clear the server logs.
► Manage the CLI user ID and password for an open-systems host:
   – Define the user ID and password.

- Remove the user ID and password.

► Refresh the path and volume information for an ESS.

► Send the ESS Network Configuration report to an e-mail address.

► Send the ESS Resource Configuration report to an e-mail address.

► Send the server logs to an e-mail address.

► Specify the e-mail address that receives server logs, ESS Network Configuration reports, and ESS Resource Configuration reports.

► View the ESS Network Configuration report.

► View the ESS Resource Configuration report.

► View the server logs.

5.8.1 Working with the Copy Services logs and reports

► The Copy Services server maintains the following logs and reports:

- **Copy Services status log (copyservices.log):** This contains messages that are associated with user actions issued through ESS Copy Services. This log includes messages associated with such actions as saving tasks and executing tasks. The log also includes information about the success or failure of the tasks. You can click the error message to get a description of the problem and its possible solutions.

  Attention: The status log is useful for problem determination. Therefore, you should only clear it when no problems exist or when you have completed data collection for all failures.

- **Copy Services Timing log (rsCStiming.log):** This contains information about the time it took for data-copy functions (PPRC or FlashCopy) to complete on a specific volume.

- **ESS Network Configuration report:** The ESS Network Configuration report contains important information about your network-defined ESS resources:

  • IP addresses for your primary and Backup Copy Services servers
  • IP addresses and host names for the ESS clusters
  • IP addresses for the ESS clients

- **ESS Resource Configuration report:** The ESS Resource Configuration report contains the following important information about your ESS resources:

  • The 2105 connection information (SAIDs, port types, and values)

  • PPRC path information (SAIDs, LSSs, remote SAIDs, ports, and status)

  • S/390 or zSeries volumes (LSSs, labels, Concurrent Copy status, FlashCopy status, PPRC status, or XRC status)

  • Open-systems volumes information (LSSs, volume serial numbers, user defined label, FlashCopy status, and PPRC Status)

To clear a log, select the log in the **Log/Report** list and click the **Clear** button. Notice that you cannot clear a report. You can view the log or report by selecting the appropriate log/report and clicking the **View** button.

To send a report or a log to an e-mail address, select the appropriate log/report in the **Log/Report** list and click the **E-Mail** button. Notice that you have to specify the e-mail address in the **E-Mail** field.

5.8.2 Refreshing path and volume information for an ESS

Use the refresh function of ESS Copy Services to get the most up-to-date information about paths and volumes for an ESS. To perform this action, select the ESS for which you want this information in the **ESS** list and click the **Refresh** button.

A status message opens on the top of the Administration panel noting that:

```
A refresh request was sent to <the selected ESS>. It may take a few minutes before the
server can refresh the data.
```

5.8.3 Managing the CLI user ID and password for an open-systems host

Use the administration panel to manage the user IDs (user IDs) and passwords for the Copy Services Command Line Interface (CLI). Open-systems hosts that issue CLI commands to a Copy Services server must provide a valid user ID and password as part of the command string. If the **Password Protection** option is enabled and the hosts do not provide the user ID and password, the command fails.

> **Attention:** When you define the user ID and password, the active Copy Services server maintains them. The active Copy Services server does not send them to the Backup Copy Services server as it does for Copy Services tasks. Therefore, you must define the user ID and password at the Backup Copy Services server in a separate step. You must manually restart the Backup Copy Services server before you can define the user ID and password there. See 10.5, "Copy Services server failover" on page 406.

To enable the password protection for access to an ESS through the CLI, click **Enabled** in the **Password protection** field. To add users, click the **Users** button and the Authorized Users dialog box will appear. Click the **Add** button and enter the user name and password. If you have entered the user name and the password, click the **Add** button and the new user will be authorized. To remove a user, select the user in the Authorized Users dialog box and click the **Remove** button. After you have performed all actions, click **Done** and the Authorized Users dialog box will be closed.

5.9 Exiting ESS Copy Services

There are several legitimate ways to quit the ESS Copy Services Web User Interface:

► You can click the **X** at the top right corner of the ESS Copy Services panel.
► You can click **Exit** in the navigation frame of any of the ESS Copy Services panels.
► You can shut down the whole browser from another browser window.

In addition, if you simply want to return to the ESS Launch page without closing ESS Copy Services, such as to launch ESS Specialist, you can keep the ESS Copy Services browser window open and use standard navigation methods such as Alt-Tab to switch to the ESS Launch page.

The difference between the first two options above is nothing more than a warning message. If you click **Exit** in the navigation frame of any of the ESS Copy Services panels, the warning message opens on top of the browser window and tells you that the tasks that you have

submitted will continue running after you close the window. Click **OK**. The ESS Copy Services browser window closes and the ESS Launch page comes into focus. In either case, the next time you click Copy Services on the ESS Launch panel, the ESS Copy Services browser window opens without requiring you to log in. The browser loads a cached version of the applet.

> **Note:** Loading the cached applet is a problem in two situations:
>
> ► If you load a new version of the ESS microcode during that time while you have not yet quit your browser, the cached applet might not conform with the new code on the ESS.
>
> ► If unauthorized users have access to your workstation when you leave it, they can access the ESS through your cached user ID.

5.10 Performing FlashCopy with the Web User Interface

In this section, we explain how to set up FlashCopy relationships using the ESS Copy Services Web interface.

Be aware of these requirements of the FlashCopy functionality:

► The source and target volume have to be in the same LSS.
► The target volume must be the same size as the source volume or larger.
► A volume can be in only one FlashCopy relationship at a time.

> **Note:** Prior to establishing a FlashCopy, we recommend that you do the following:
>
> 1. Unmount the target volume from all host systems, since the FlashCopy process is a destructive to the target and will overwrite the data on the target volume.
>
> 2. Additionally, stop all active reading/writing operations to the target.
>
> 3. Also, quiesce your application and flush the data to the source volume.

There are two different ways of establishing a FlashCopy pair:

► From the Volume panel
► From the Tasks panel (once a task for a FlashCopy is created and saved)

5.10.1 Establishing a FlashCopy pair

Use the Volumes panel to establish a FlashCopy pair. Select the LSS within which you want to perform the FlashCopy. This can be either done in the source or target area of the volumes panel (Figure 5-24 on page 131).
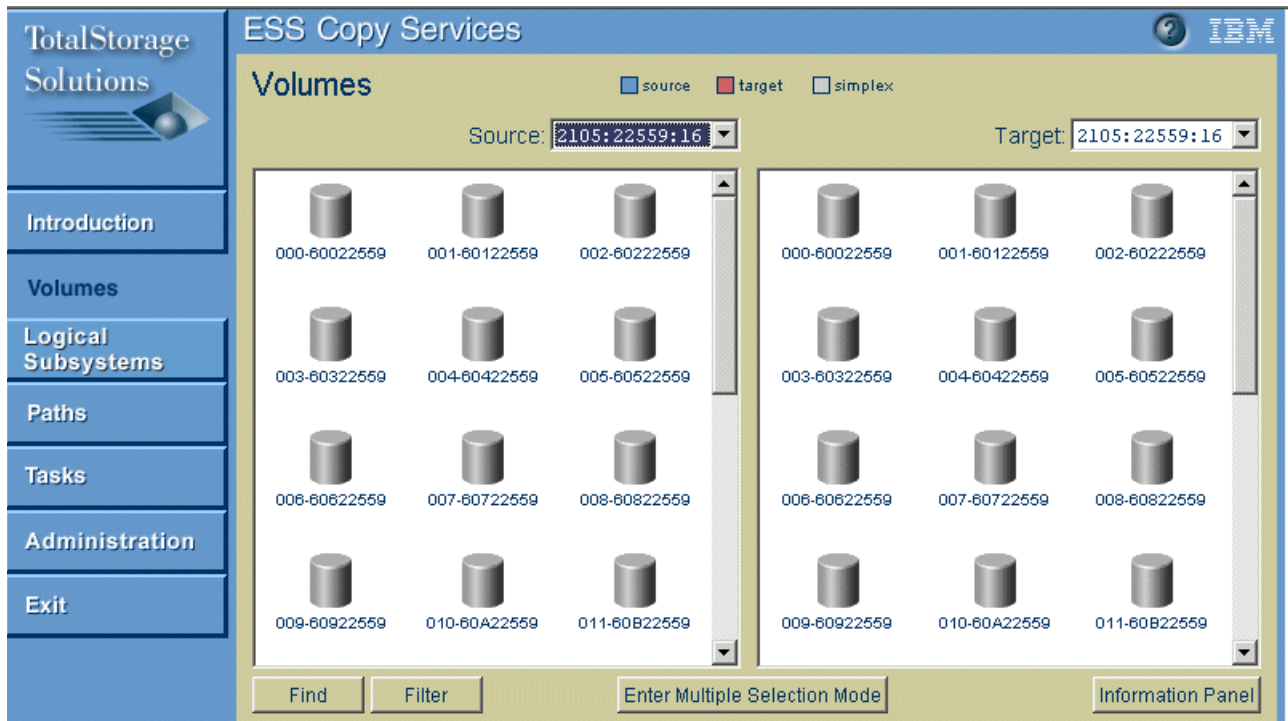
Figure 5-24   Volumes panel

You always need to have two components to establish a FlashCopy pair: a source and a target. Select the source volume, and, with a right-click, the target volume. If you have selected the wrong source or target volume, just click on the correct source volume again.

Once you have selected the source and the target, you do a second right-click on the target to bring up the Task Wizard (Figure 5-25). Select the **Establish FlashCopy pair** option and click **Next**.
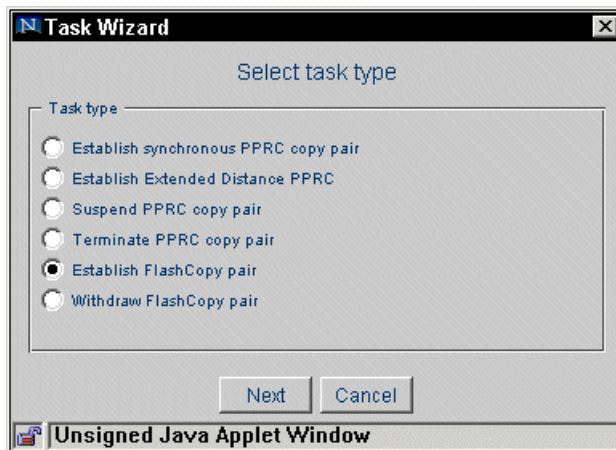
Figure 5-25   Task Wizard panel

Within the next window, you can specify the **Copy** option of the FlashCopy pair (Figure 5-26 on page 132). For further information about this option, refer to 3.4, "FlashCopy tasks and options" on page 39.
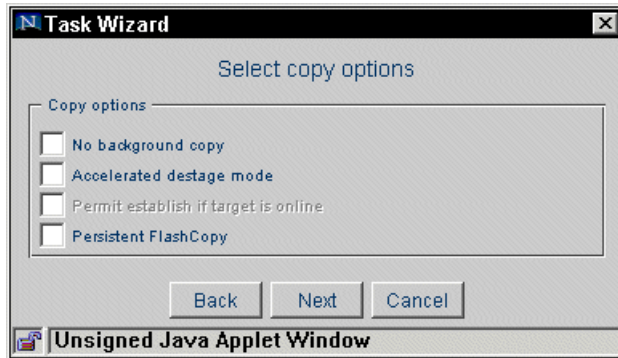
Figure 5-26   Select copy options panel

Selecting the FlashCopy Options

3.4, "FlashCopy tasks and options" on page 39 gives detailed descriptions of the different options for FlashCopy. Refer to this section to be sure to fully understand the different options. We present a summary here of the different options.

No background copy option (nocopy)

If this option is checked:

► Only the tracks that are modified on the source volume are copied to the target volume.

► This relationship between source and target volume remains forever and has to be broken manually.

► By default, this option is not selected and all data is copied from the source to the target volume of a FlashCopy pair (background copy). Once all data is copied, the relationship ends automatically (unless the Persistent FlashCopy option is selected).

► If two volumes are in a FlashCopy "nocopy" relationship, you can perform a FlashCopy Start Background Copy to create a physical background copy.

Accelerate destage mode

Select this option to cause a FlashCopy source volume track, which is being modified and is in FlashCopy relationship, to be destaged from cache sooner than it would be if normal cache algorithms were applied. This feature minimizes (for the source volume) the number of modified tracks that are resident in cache.

Note that:

► With FlashCopy Version 1, this option might (under specific circumstances) speed up the completion of the background copy process.

► With FlashCopy Version 2, This feature has no effect, since the cache algorithms have changed.

Permit establish if target is online

Select this option if you want to establish a FlashCopy relationship even if the target is online to the S/390 or zSeries host.

This option does not apply to open-systems volumes.

Persistent FlashCopy

If you want the FlashCopy relationship to remain after the FlashCopy operation completes you can use the **persistent FlashCopy** option. The FlashCopy relationship between source

and target volumes remains indefinitely and must be broken by a Withdraw FlashCopy task. By default, this option is not selected.

From the next panel you can either **Save**, **Run**, or **Cancel** the copy task, as shown in Figure 5-27. Once a task is saved, it can be executed from the Task panel at any time. Optionally, a name and description of the task could be specified. Even if you do not want to save the task you have created, we recommend specifying a name and description. This will help with the interpretation of the Copy Services log file later on. An example would be to retrieve the execution time of the background copy of a FlashCopy pair.
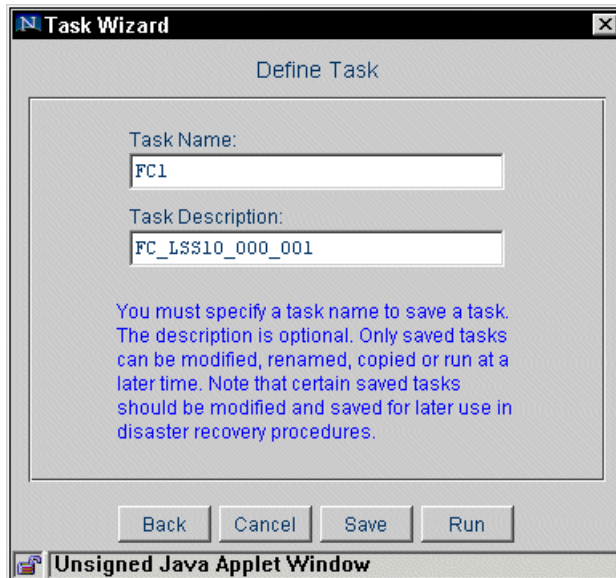
Figure 5-27   Define Task panel

If a FlashCopy is issued, a bitmap is created for the data copy from the source to the target. The time to establish the FlashCopy relationship is only a few seconds. After this period, the source is immediately available to the host system, and data will be copied from the source to the target in the background.

Once a FlashCopy is started, the display of the source and target volume from the Volumes panel changes:

► A blue lightning bolt appears on the FlashCopy source volume.

► A red lightning bolt appears on the FlashCopy target volume.

► In addition, a FlashCopy copy icon (a small white square on the lower left of the volume icon (see Table 5-1 on page 112) now also appears on the target volume at the time that the copy begins. The FlashCopy copy icon remains until you issue a command to withdraw it (see 5.10.3, "Withdrawing a FlashCopy pair" on page 135), or until you use the volume as a source volume.

► The lightning bolts disappear when the copy operation completes. If the pair was established without a background copy, the lightning bolts remain until the withdraw operation is performed or until the entire volume is rewritten.

► The Information panel shows the progress of the copy.

Figure 5-28   FlashCopy volume display

> **Tip:** After you have created and saved the task to establish the FlashCopy with the **No background copy** option, you need to create the corresponding task to withdraw the pair. We recommend that you modify the task that you just created so that it withdraws the pair, then name and save the task. This not only saves time, but it is also useful for disaster recovery. For further information, refer to 5.7.4, "Modifying a task" on page 127.

Establishing FlashCopy relationships using the LSS panel

To establish FlashCopy relationships between all the volumes in two LSSs, you can create a task at the LSS level in the Logical Subsystem panel.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 5-2 on page 117).

Click on the source LSS and the LSS ID turns blue. Right-click the target LSS and the LSS ID turns red. Right-click again on the target LSS and the Task Wizard opens. Select the options as you would do for volumes pairs.

5.10.2 Getting Information about a FlashCopy pair

By selecting one of the volumes of a FlashCopy pair and clicking the **Information button**, you get information about this particular pair. If you have selected the source volume, you can see how many tracks still have to be copied to the target volume (Figure 5-29 on page 135).
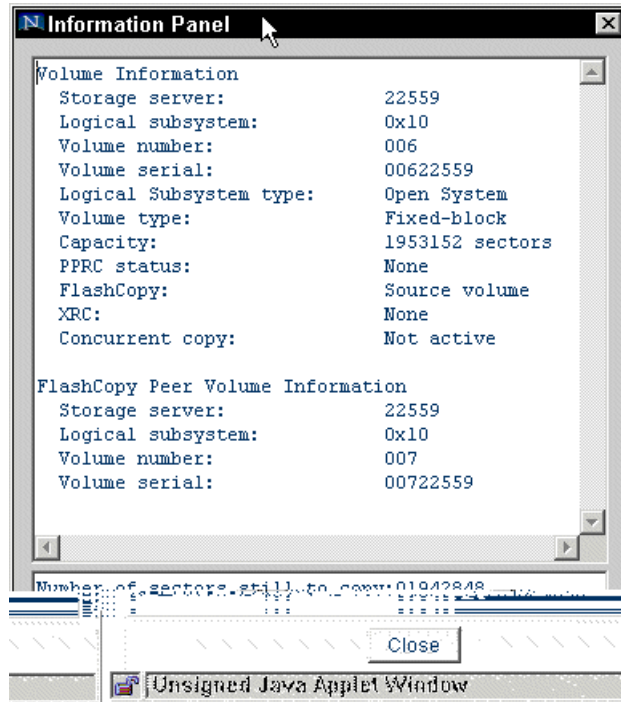
Figure 5-29   Information panel of a FlashCopy source

5.10.3 Withdrawing a FlashCopy pair

In the following cases, you have to *withdraw* a FlashCopy pair:

► If a FlashCopy pair is not needed anymore, but has not yet finished the background copy.

► If a FlashCopy pair that was created with the **Do not perform background copy** option is not needed anymore.

► If a FlashCopy pair that was created with the **Persistent FlashCopy** option is not needed anymore.

The command to **Withdraw FlashCopy pair** results in the termination of the FlashCopy operation. All the data on the target can be corrupted and unusable if it is inadvertently written on by another application.

Use the Volumes or the Logical Subsystem panel to withdraw a FlashCopy pair.

Withdrawing FlashCopy pairs using the Volumes panel

Open the Volumes panel and select the LSS within which you want to perform the withdraw FlashCopy. This can be either done in the source or in target area of the volumes panel (Figure 5-30 on page 136).
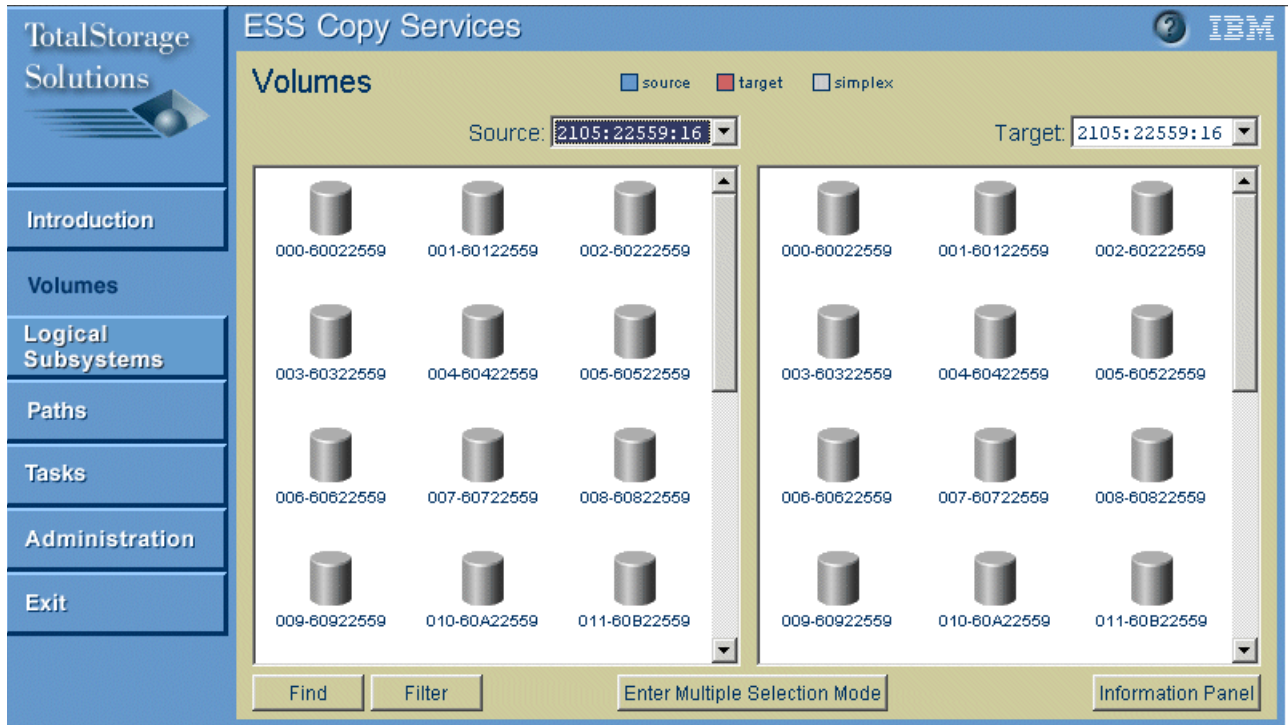
Figure 5-30   Volumes panel

You always need to have two components to withdraw a FlashCopy pair, a source and a target. Select the source volume and, with a right-click, the target. If you have selected the wrong source or target volume, just click on the correct source volume again.

To withdraw a FlashCopy pair, select the source volume, select the target volume with a right-click, and start the Task Wizard (Figure 5-31) by right-clicking the target volume again.
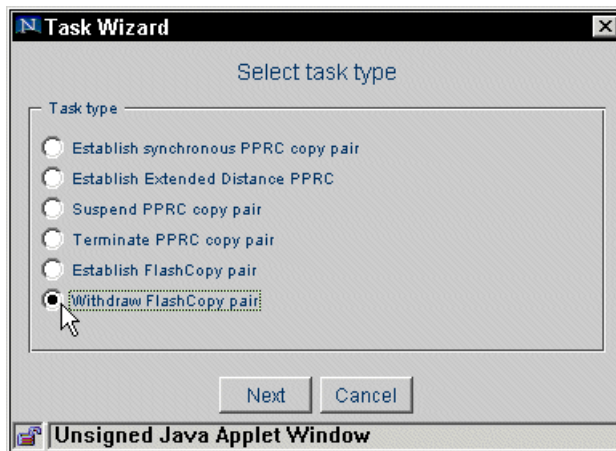
Figure 5-31   Withdraw a FlashCopy pair

Select the **Withdraw FlashCopy pair** option and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 5-32 on page 137).
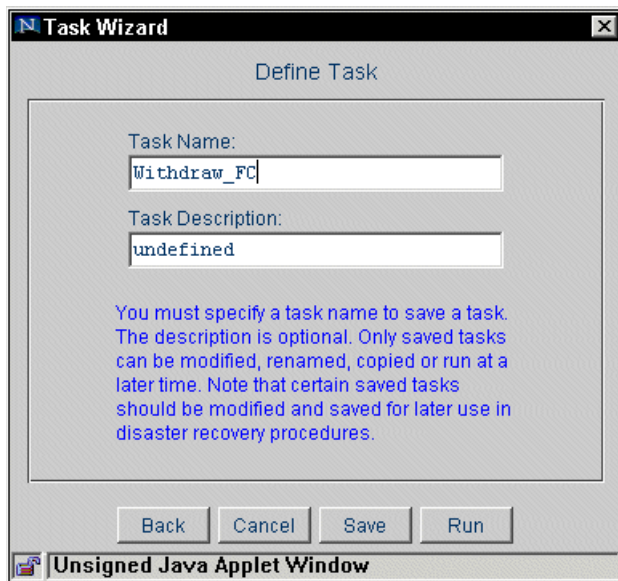
Figure 5-32   Define Task panel

If you save the task, it can be executed from the Tasks panel any time. Once you have run the task, the lightening bolts on the source and target volume disappear, and the FlashCopy relationship between source and target volume will be removed. Also, the FlashCopy icon (see Table 5-1 on page 112) on the target volume will disappear.

Withdrawing FlashCopy pairs using the LSS panel

To withdraw all FlashCopy pairs between two logical subsystems, you can create a task at the LSS level, in the Logical Subsystem panel.

> **Note:** All the FlashCopy relationships that have source volumes in the selected source LSS and their associated targets in the selected target LSS will be terminated.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 5-2 on page 117).

Click on the source LSS and the LSS ID turns blue. Right-click on the target LSS and the LSS ID turns red. Right-click again on the target LSS and the Task Wizard opens. Select the options as you would do for volume pairs.

Withdraw FlashCopy to the target

Use the option **Withdraw FlashCopy to the target** against the target volume when you want to remove its FlashCopy icon (represented by a white square at the lower left of the volume icon; see Table 5-1 on page 112) after the background copy has been successfully completed. Typically, you would remove the indicator because you have decided that it is no longer necessary to mark the volume as a copy of another.

> **Attention:** If you perform a Withdraw to target action to a target volume of an established relationship, it will withdraw the relationship and remove the "I-am_a_target" icon.

To perform this action, click on the target volume, right-click on the target volume (the volume ID turns gray), and right-click the target volume again to open the Task Wizard (Figure 5-33).
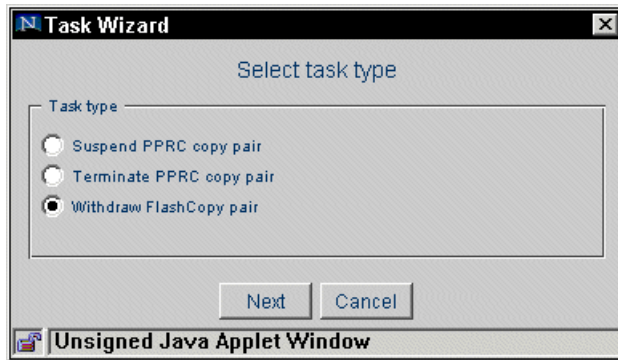
Figure 5-33   Task Wizard: Withdraw FlashCopy pair

Select the **Withdraw FlashCopy** pair option if you want to immediately terminate the FlashCopy relationship and click **Next**. In the next panel (Figure 5-34) select the **FlashCopy withdraw to the target** option. Click **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 5-32 on page 137). Once you have run the task, the FlashCopy icon (see Table 5-1 on page 112) on the target volume will disappear.
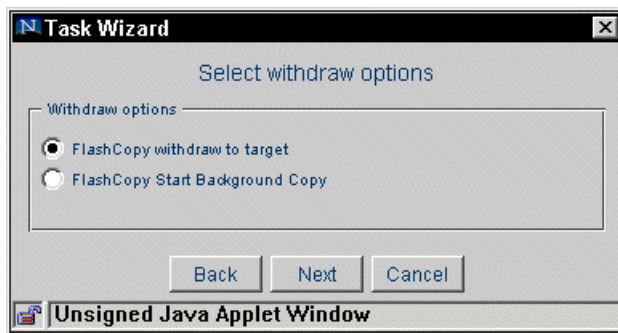
Figure 5-34   Task Wizard: Withdraw options

Note: If you perform the **Withdraw FlashCopy to the target**, make sure you select a target volume of a FlashCopy relationship in the volumes panel. You cannot perform this action by only selecting a source volume.

Withdrawing FlashCopy to the target using the LSS panel

To withdraw all FlashCopy relationships to the target for all target volumes in the same logical subsystem, you can create a task at the LSS level, in the Logical Subsystem panel.

Note: All FlashCopy relationships that have a target volume in the LSS you select will be terminated.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 5-2 on page 117).

Left-click on the LSS and the LSS ID turns blue. Right-click the same LSS and the LSS ID turns gray. Right-click again on the LSS and the Task Wizard opens. Select the options as you would do for volumes pairs.

5.10.4 FlashCopy Start Background Copy

If you have performed a FlashCopy pair with the **No background copy** option and you need to create a permanent physical copy of the data for backup or disaster recovery purposes, you can perform a **FlashCopy Start Background Copy**. All data will be copied from the source to the target. After the copy is complete, the FlashCopy relationship is terminated, unless the FlashCopy relationship is persistent, in which case, the relationship remains.

FlashCopy Start Background Copy using the Volumes panel

Open the **Volumes** panel and select the LSS within which you want to perform the **FlashCopy Start Background Copy**. This can be done in either the source or target area of the volumes panel (Figure 5-35).
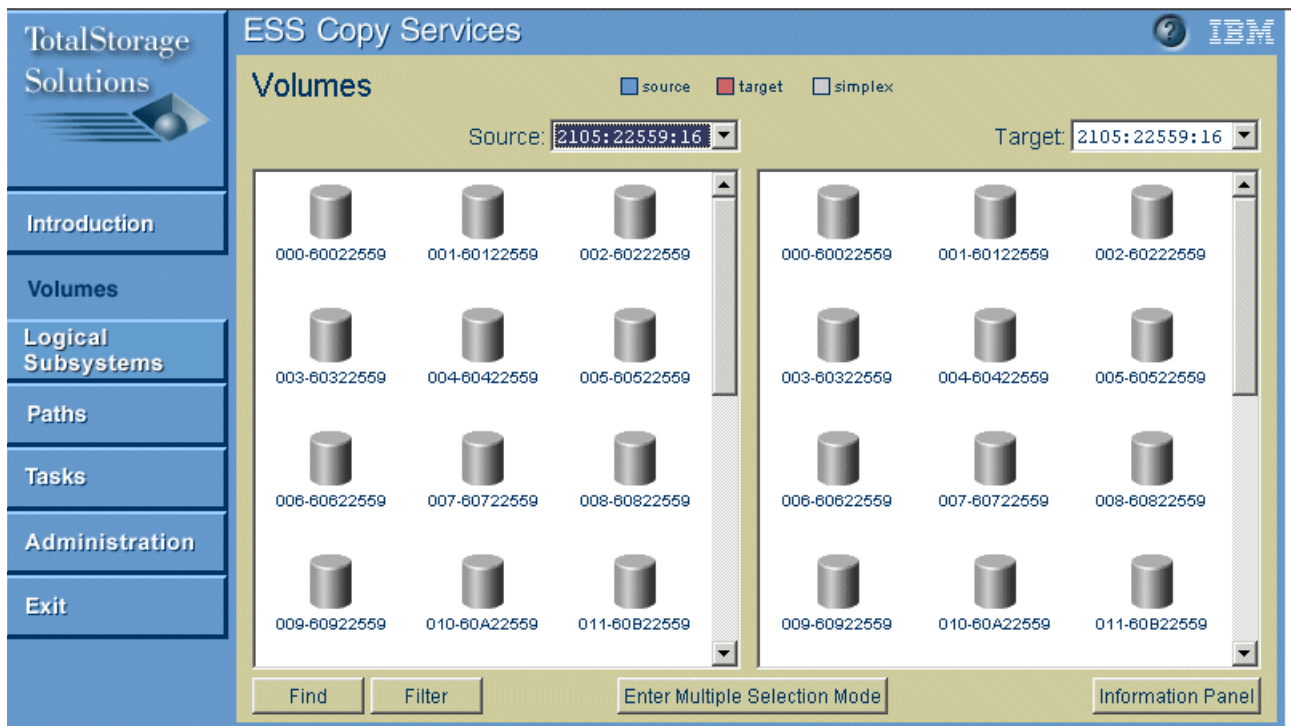
Figure 5-35   Volumes panel

To perform a **FlashCopy Start Background Copy**, click on the source volume of the FlashCopy pair. Right-click the source volume (the target ID of the volume turns gray) and right-click the volume again to open the Task Wizard (Figure 5-36 on page 140).
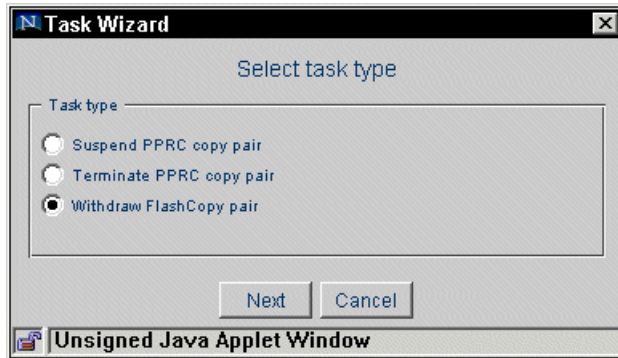
Figure 5-36   Task Wizard: Withdraw FlashCopy pair

Select the **Withdraw FlashCopy pair** option and click **Next**. In the next panel, select the
FlashCopy Start Background Copy option (Figure 5-37).
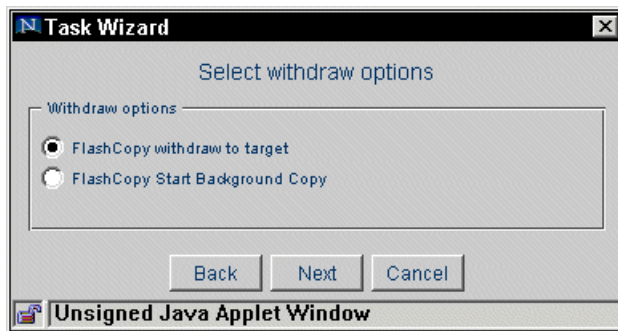
Figure 5-37   Task Wizard: Withdraw options

Click the **Next** button and decide whether to **Save**, **Run**, or **Cancel** the task. Once you run
the task, the background copy task copies the data from the source to the target. After the
copy is complete, the FlashCopy relationship is terminated and the lightning bolts disappear
on the source and the target volumes, unless the FlashCopy relationship is persistent, in
which case, the relationship remains. Refer to 5.10.1, "Establishing a FlashCopy pair" on
page 130 for further information.

Note: If you perform **FlashCopy Start Background Copy**, make sure you select a source
volume of a FlashCopy relationship in the volumes panel. You cannot perform this action
by only selecting a target volume.

FlashCopy Start Background Copy using the LSS panel

To start a background copy of all FlashCopy pairs that have a source volume on a logical
subsystem, you can use the Logical Subsystem panel.

Note: When you run the task at the LSS level, be aware that the task will run on all the
source volumes in that LSS.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right
corner. A panel showing all the LSSs is displayed, with the appropriate name and serial
number indicated below each LSS icon. The different icons on the Logical Subsystem Panel
do have different meanings (for further information, refer to Table 5-2 on page 117).

Click on the desired LSS. The LSS ID turns blue. Right-click the LSS and the LSS ID turns gray. Right-click the LSS again and the Task Wizard opens. Select the options as you would do for volumes pairs.

5.11 Performing PPRC with ESS Copy Services Web User Interface

In this section, we explain how to set up PPRC and PPRC-XD using the ESS Copy Services Web interface. For further information about PPRC and PPRC-XD, refer to Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53. In general, there are two steps needed to successfully establish PPRC:

► Setting up paths between the PPRC source and target.

► Establishing, suspending, and terminating the PPRC pairs, either single volumes or entire logical subsystems.

> **Note:** To set up all PPRC pair related functions, either the Volumes panel or the Logical Subsystems panel (all volumes in the source and the target LSS) can be used. In this book, only the usage of the Volume panel is explained. For further information about how to set up a PPRC pair related function on the Logical Subsystem panel, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

Be aware of these requirements of the PPRC functionality:

► Paths for PPRC have to be available and need to be defined first.
► All PPRC ESCON links are unidirectional.
► The target volume has to be the same size as the source or larger.
► You can have up to four ESSs in a Copy Services group.

5.11.1 Establishing paths

Before you can establish any PPRC pairs, you first have to set up the paths between the source and the target LSSs. The paths are needed for communication between the PPRC pairs and to copy data from the source to the target.

> **Note:** In our example we have used only one ESS to set up PPRC. That is possible, as the ESS contains both source and target volumes at the same time. However, for high availability and disaster recovery configurations, two or even more Enterprise Storage Servers are required.

Use the Paths panel of the ESS Copy Services Web Interface to set up paths for PPRC (Figure 5-38 on page 142).
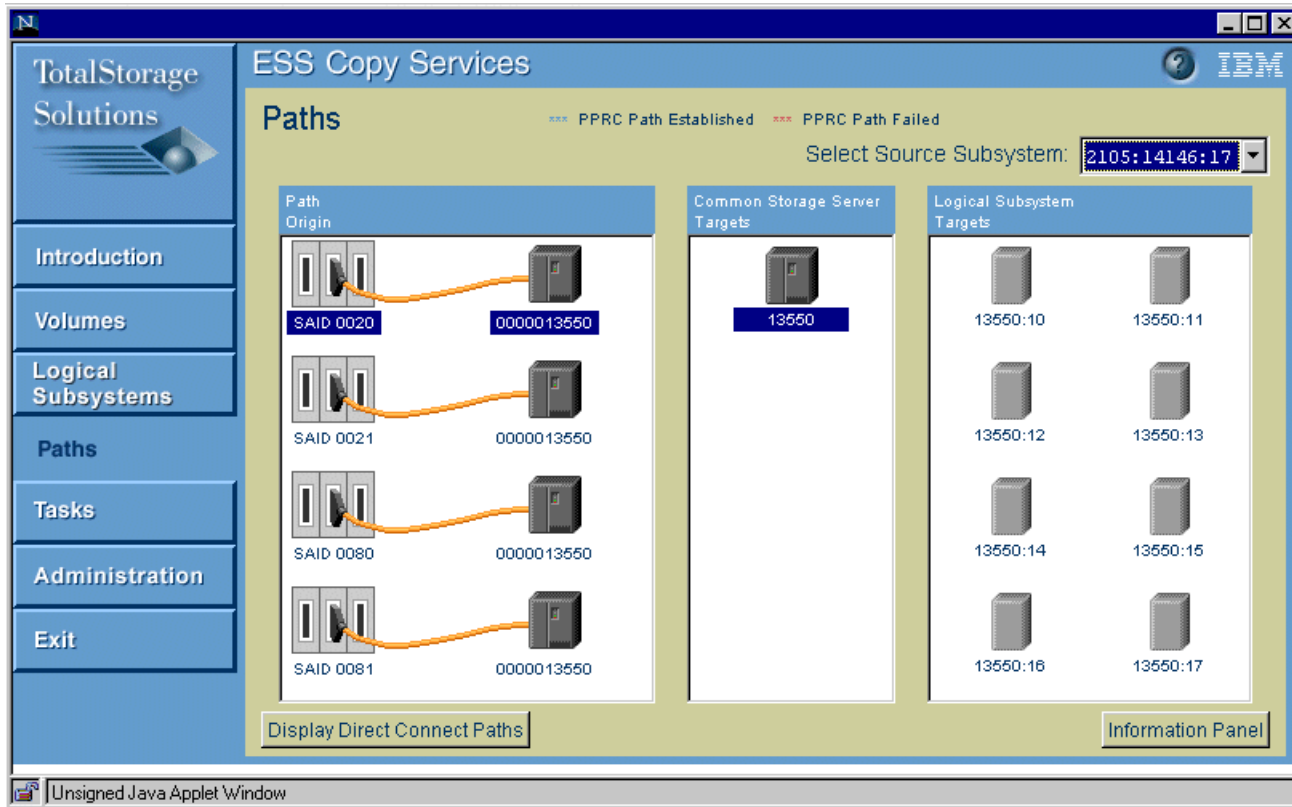
Figure 5-38   Paths panel of the Copy Services Web User Interface

Select the source of the PPRC relationship. This is done with the drop-down menu of the Select box. All available ESCON adapters for the source will automatically be displayed in the Path Origin area (see Figure 5-38).

Select the ESCON adapters you want to use for the PPRC.

Note: Multiple adapters could be selected with a right-click after the first ECSON adapter was selected. If you choose the wrong adapter(s), just click on the correct ESCON adapter again to delete the selection.

The ESSs that are connected to the adapter(s) will be automatically displayed in the Common Storage Server Target area.

Note: If you have some ESSs that are directly connected (that is, no ESCON switch between them), you can click **Display Direct Connect Paths** to display those paths.

You must perform a refresh function (see 5.6.1, "Display Direct Connection Paths" on page 123) for the following conditions:
► If this is the first time you are accessing ESS Copy Services
► After an update of the ESS microcode
► After a concurrent service action
► If the cabling for the ESCON links is modified

Next, click on the target ESS in the **Common Storage Server Target** area. All logical subsystems available on the target ESS will be displayed in the **Logical Subsystems Targets** area (see Figure 5-39).

Within the Logical Subsystems Targets area, select the target LSS of your PPRC path. If you have chosen the wrong target, just click on the correct LSS again to delete the selection.

In the example shown in Figure 5-39, we have selected LSS 11 to be the target logical subsystem.
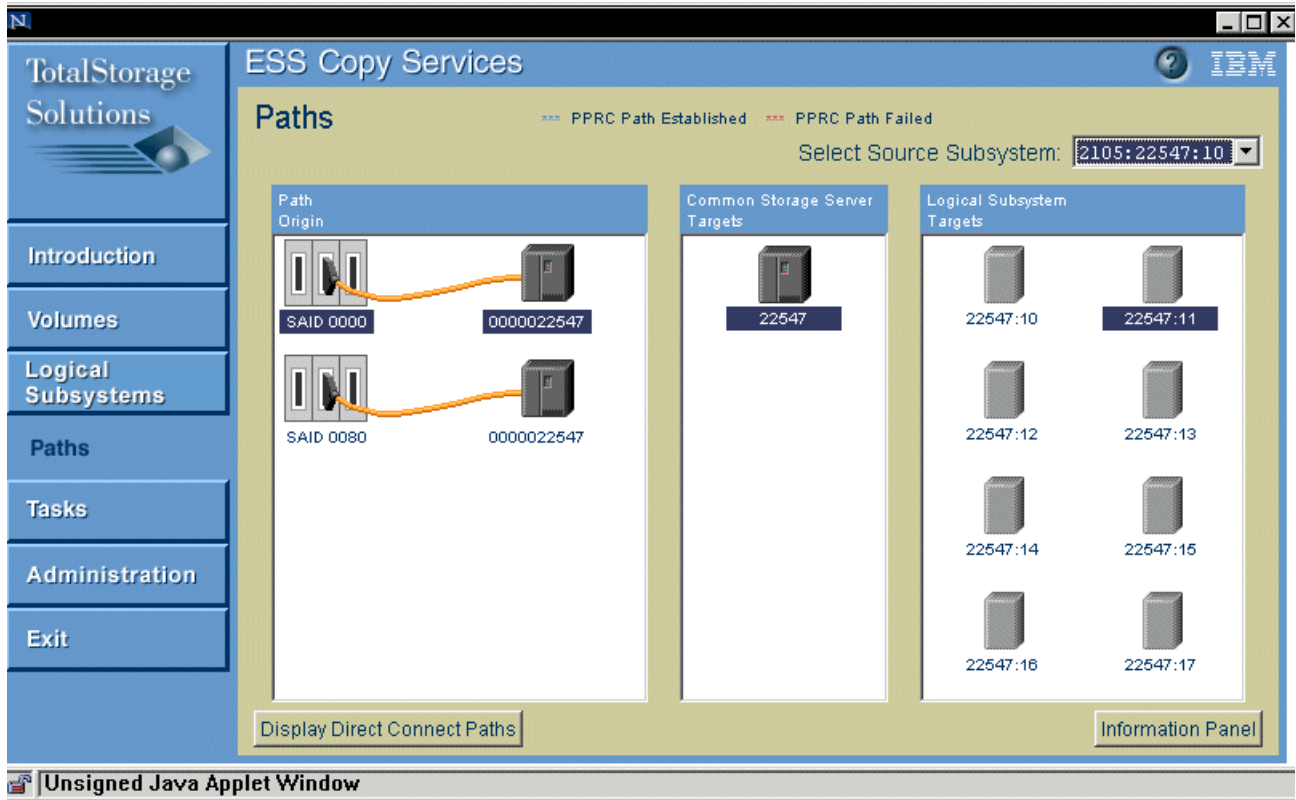
Figure 5-39   Establish PPRC path

Once the target and source of the PPRC path have been selected, click on one of the highlighted target LSSs to bring up the Task Wizard (Figure 5-40). Select the **Establish Path** option and click **Next**.
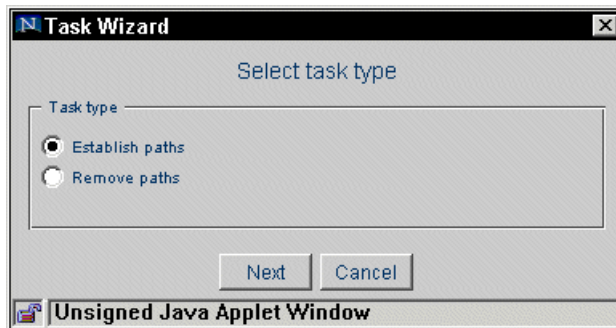
Figure 5-40   Task Wizard: Establish Path

If a switch is part of the path, the Select outgoing ports panel of the Task Wizard opens (Figure 5-41).
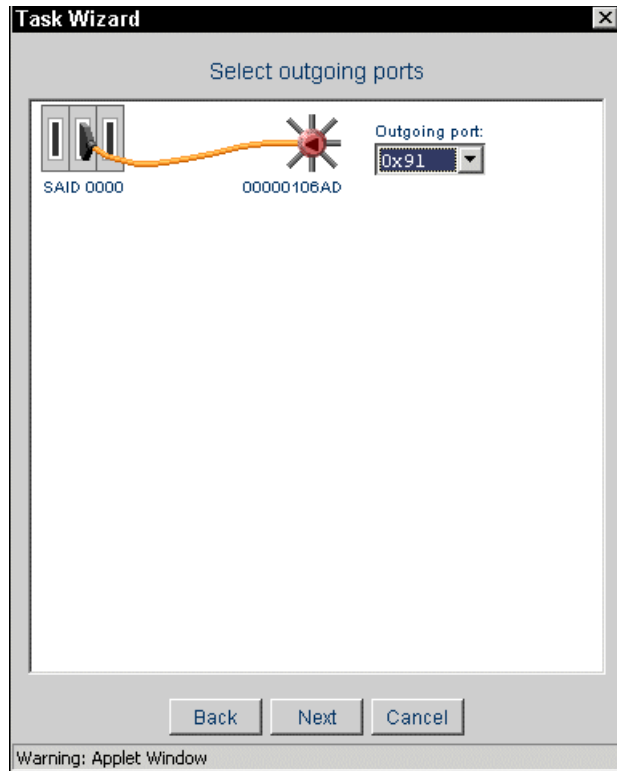
Figure 5-41   Task Wizard: Select outgoing ports panel

In the **Outgoing Ports list**, select the desired port on the switch. Click **Next**. Within the next panel, you can specify the path options (Figure 5-42).
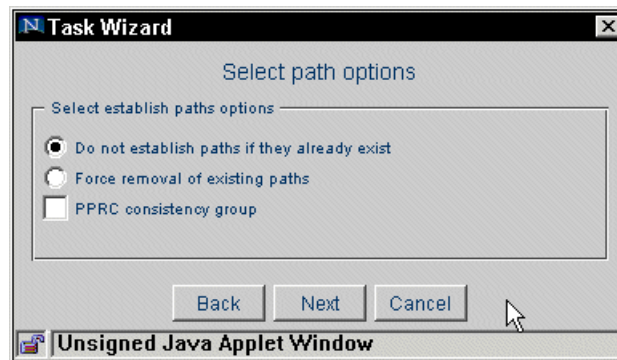
Figure 5-42   Establish PPRC options

► **Do not establish paths if they already exist:** If this option is checked and there is already a path defined from the source to the target, the operation of establishing paths will not be executed.

► **Force removal of existing paths:** You must select this option if the selected path goes through a switch. The selection causes any logical paths between a host system and the selected ESCON adapter to be removed before establishing logical paths.

> **Attention:** The **Force removal of existing paths** option removes all logical paths from the host to the port. This is potentially disruptive to host activity, since there may be applications using the port. It is recommended that you vary off all logical paths from the host to the port by means of a host command. Only as a last resort, or when the operator is absolutely positive there would be no negative impact, should the **Force removal of existing paths** option be used.

In either case, you can also select **PPRC consistency group** to create a consistency group. For information on consistency groups, see 4.5, "PPRC data consistency" on page 72.

Click **Next** and decide whether to **Save**, **Run**, or **Cancel** the task. Click **Save** if you want to save the task in the task repository. If you save the task, it can be executed form the **Tasks** panel any time.

> **Tips:**
>
> ► If you save the task, it can be executed from the Tasks panel at any time (see 5.7.3, "Running a task and viewing an error about a failed task" on page 126).
>
> ► After you create and save the task to establish paths, you need to create the corresponding task to remove the path. You can modify the task within the Tasks panel so that it removes the path (see 5.7.4, "Modifying a task" on page 127). This not only saves time, but it is also useful for disaster recovery.

When you run the task, the PPRC path is established and three blue asterisks appear right below the adapter in the Path Origin column of the Paths panel (see Figure 5-43).
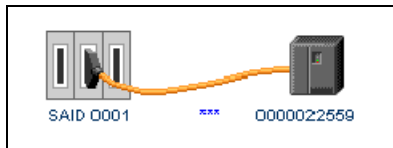
Figure 5-43   Path successfully established

5.11.2 Removing paths

If you do not need a path any more, you can remove it. Use the Paths panel to remove a path. From the Select Source Subsystem list, select the source logical subsystem. The blue asterisks below the path-connection symbol in the Path Origin column mean that a path is already established. Select the desired path in this column. The Paths panel displays the common storage servers and their designators (Figure 5-44 on page 146).
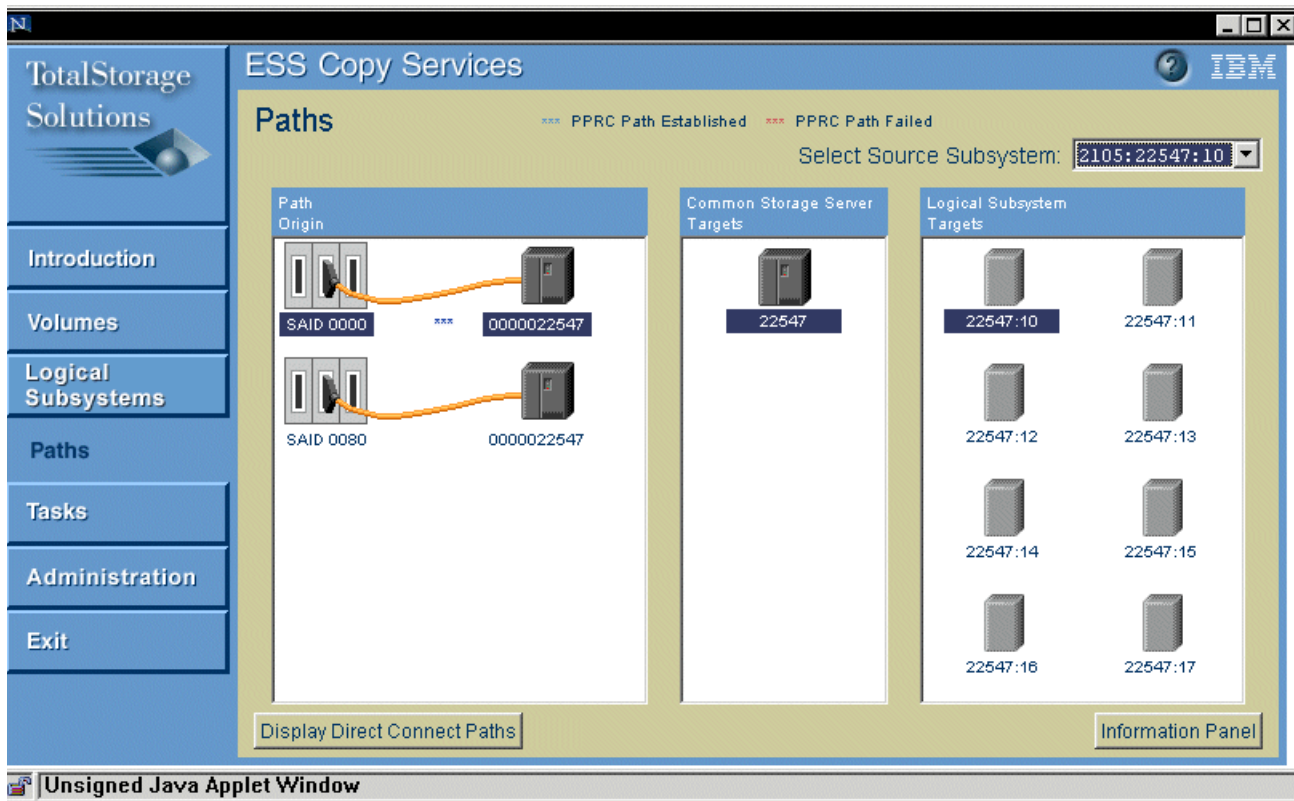
Figure 5-44   Paths panel: Remove paths

Select the desired ESS in the **Common Storage Servers Targets** column and the Paths panel displays a list of logical subsystems. Select the logical subsystem that is the destination path and right-click on the desired logical subsystem (Figure 5-44). The Task Wizard opens (Figure 5-45).
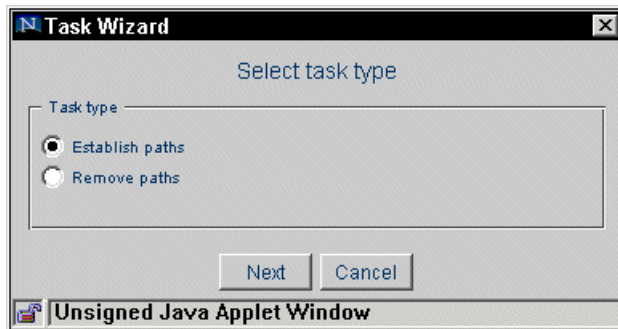
Figure 5-45   Task Wizard: Remove paths

Select the **Remove paths** option and click **Next**. In the next panel, you can specify the task options (Figure 5-46 on page 147).
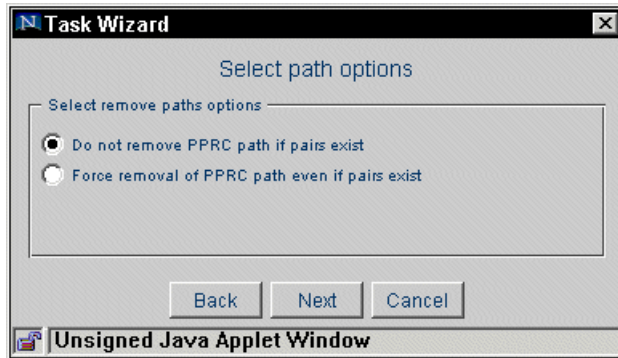
Figure 5-46   Task Wizard: Remove tasks options

► **Do not remove PPPC path if pairs exists:** If you choose this option, the path will not be removed if there is a PPRC pair using this path.

► **Force removal of PPRC path even if pairs exists:** This option enables you to force the removal of existing PPRC paths even if there are PPRC pairs established that could use this path. This option is helpful for the fail-over/fail-back mode. For further information, refer to 10.4, "Types of PPRC Failover and Failback procedures" on page 405.

From the next panel, you can either **Save**, **Run**, or **Cancel** the task. As mentioned before, once a task is saved, you can run it from the Tasks panel at any time.

Once you have run the task, all PPRC paths from the source to the target LSSs will be removed and the blue asterisks below the Path Connection symbol in the Path Origin column will disappear.

> **Tip:** To remove a path, you can either use the Paths panel, the Tasks panel (by modifying an Establish Paths task or by running a Remove Paths task) or by using the **Logical Subsystems** panel. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

5.11.3 Establishing a synchronous PPRC pair

Use the Volumes panel to establish PPRC pairs. On the left side, you select the source LSS; and on the right side, the target LSS. This is done using the drop-down menus at the top of the Volumes menu.

The source and target logical subsystem is specified the following way:

Device type (4 digits):ESS Serial number (5 digits):LSS number (2 digits)

In our example shown in Figure 5-47 on page 148 we have selected 2105:22547:10 as source and 2105:22549:11 as target of our PPRC pair.
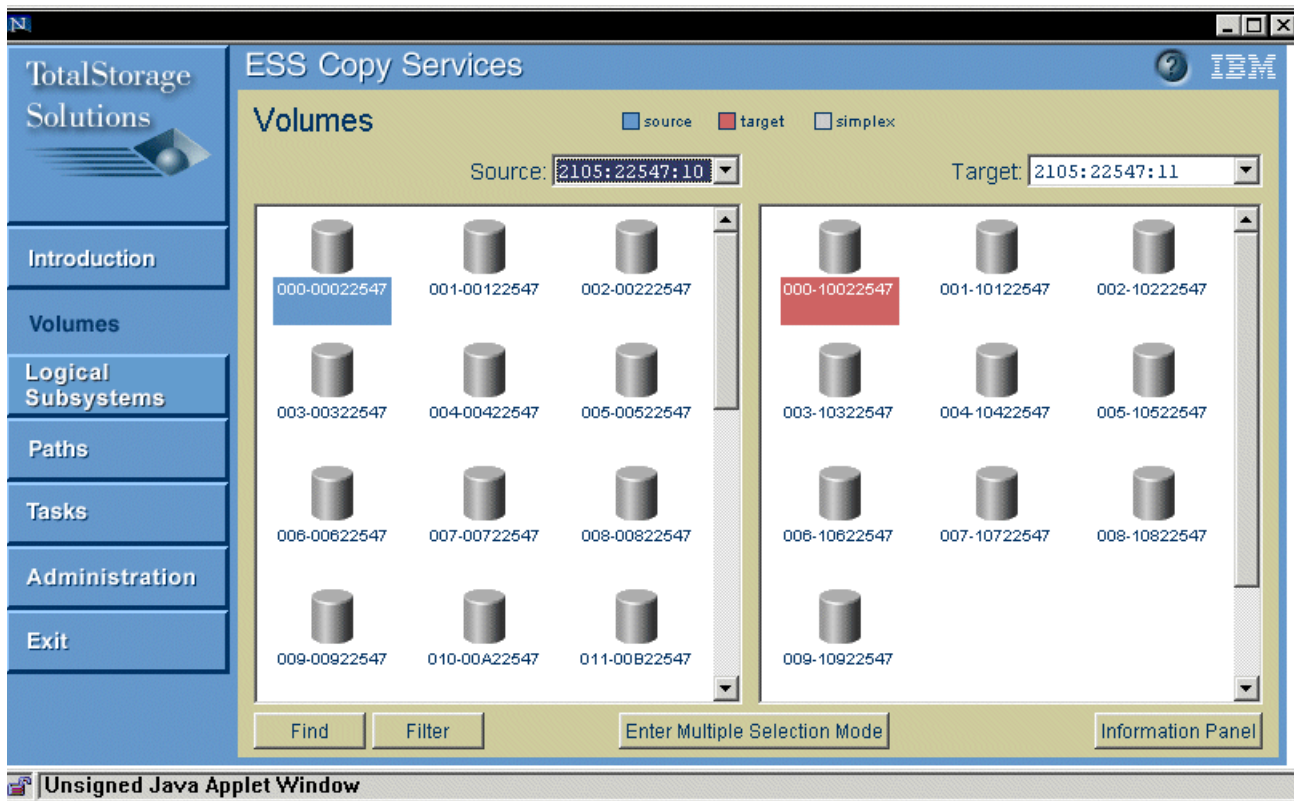
Figure 5-47   Volume panel: establish PPRC pair

You always need to have two components to establish a PPRC pair: a primary (source) and a secondary (target). Select the source volume, and, with a right-click, the target. If you have selected the wrong source or target volume, just click on the correct source volume again to clear the selection (Figure 5-47).

Once you have selected the primary and the secondary, you do a second right-click on the secondary (target) to bring up the Task Wizard (Figure 5-48). Select the **Establish PPRC copy pair** option and click **Next**.
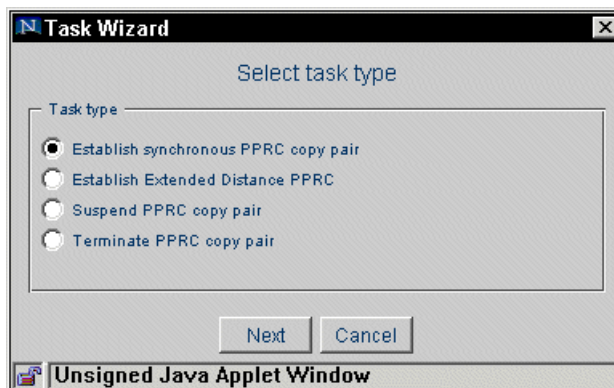
Figure 5-48   Task Wizard: Select task type

Within the next panel, you can specify the **Copy** options of the PPRC pair (Figure 5-49 on page 149). Click **Next** when you have finished the selection.
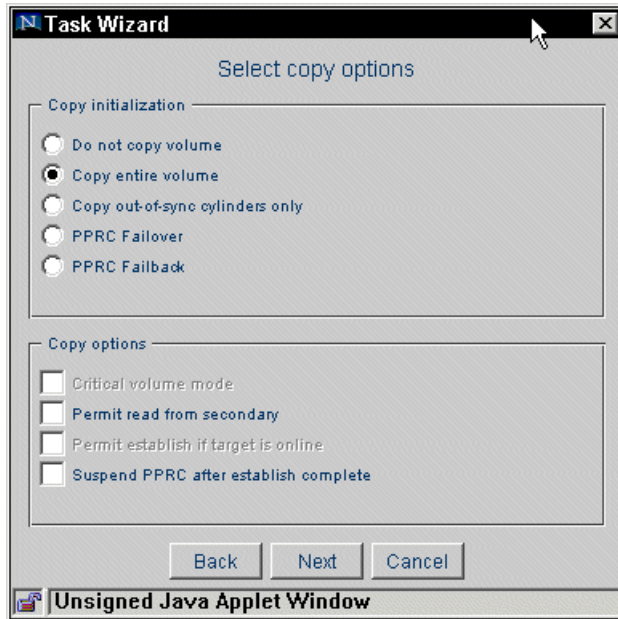
Figure 5-49   PPRC copy options

Options for PPRC synchronous

► Copy initialization:

– **Do not copy volume:** If this option is checked, the PPRC pair relationship is established without copying any data from the source to the target. This option is used when the source and target contain exactly the same data and are consistent while in the simplex state. Choosing this option causes the volumes to turn duplex immediately on the assumption that both source and target contain identical data.

– **Copy entire volume:** If this option is checked, all data is copied from the source to the target volume. This option has to be used the first time a PPRC relationship is going to be established and is needed to guarantee that source and target contain the same data.

– **Copy out-of-sync cylinders only:** This option copies only the data that was updated on the source volume since a PPRC copy pair was suspended. The option is used to resynchronize a PPRC synchronous or a PPRC-XD pair.

► PPRC modes:

– **PPRC Failover:** Select this option to reverse the direction of a PPRC pair. Use this option to make your current secondary volume at the recovery site become a primary volume (your current primary volume at the production side being the secondary volume of the PPRC relationship). For further information, refer to Chapter 10, "Disaster recovery" on page 399.

– **PPRC Failback:** Select this option to resynchronize the PPRC volumes at the recovery site with their peer volumes at the production site. For further information, refer to Chapter 10, "Disaster recovery" on page 399.

► Copy options:

– **Permit read from secondary:** Selecting this option allows host servers to read from the PPRC secondary volume. The PPRC pair must be in a full-duplex state in order for the host server to read the volume. This option is helpful for operating systems that use a Logical Volume Manager that reads data structures from the header of a secondary volume to update internal databases, such as the ODM under AIX, or /etc/lvmtab under

HP-UX. For further information, refer to Chapter 8, "Open systems specifics" on page 327.

– **Suspend PPRC after establish complete:** If you select the **Copy entire volume** or **Copy out-of-sync cylinders only** options described above, this means that after the data is completely copied to the target volume, the target volume goes into the suspended state. Selecting this option saves you the effort of creating a separate task to suspend the pair. This is useful in a PPRC-XD catch-up operation over extended distances.

From the next panel, you can either **Save**, **Run**, or **Cancel** the copy task, shown in Figure 5-50. Click **Save** to run the task and run it later. To run the task immediately, click **Run**.
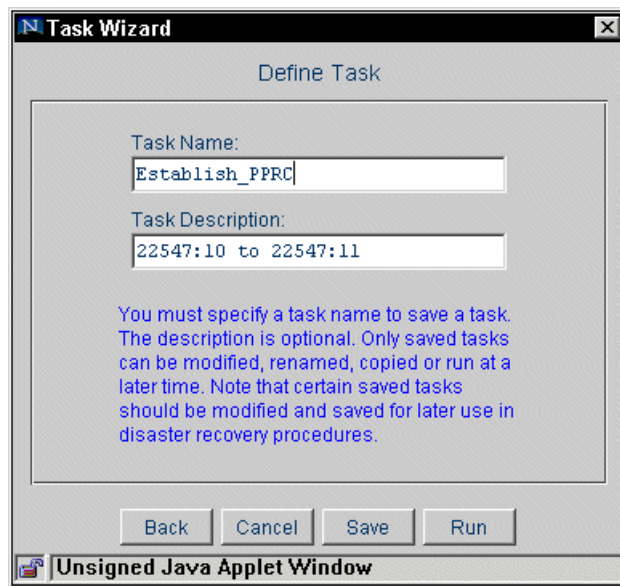
Figure 5-50   Task Wizard: Establish PPRC path

Tip: After you create and save the task to establish synchronous PPRC copy pairs, we recommend that you create the tasks that suspend the pairs and terminate the pairs. You can save time in creating tasks if you use the task you just saved as a template to create the suspend task. Then use the suspend task as a template to create the terminate task. In the **Task** panel, select the task you created and click on the **Modify** button to change the options of the task. Then, give a new name and save as a new task.

Once a synchronous PPRC pairs is established, the display of the primary and secondary volumes from the Volumes panel changes, indicating that the initial copy is in progress (Figure 5-51 on page 151).
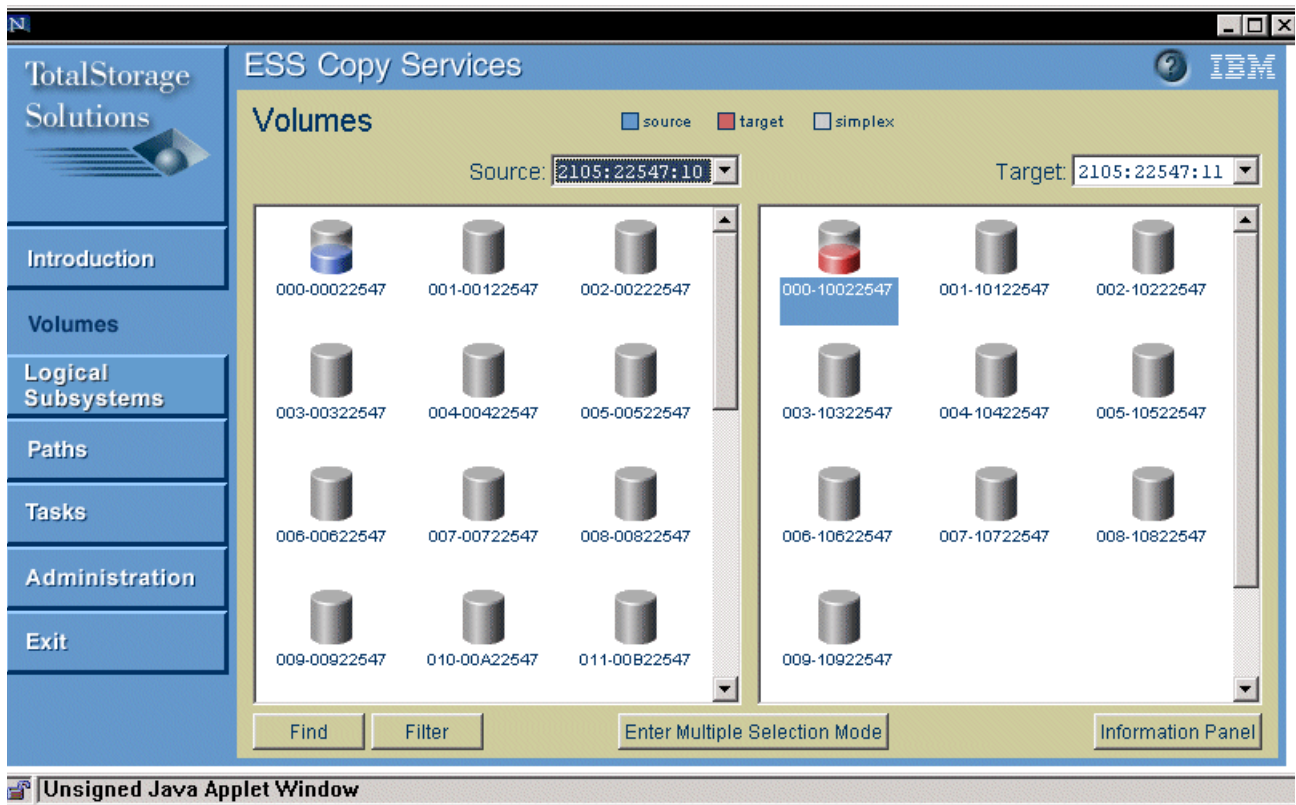
Figure 5-51   PPRC relationship in progress

Once the volumes are in sync, the icons of the volumes icons will change to solid (Figure 5-52).
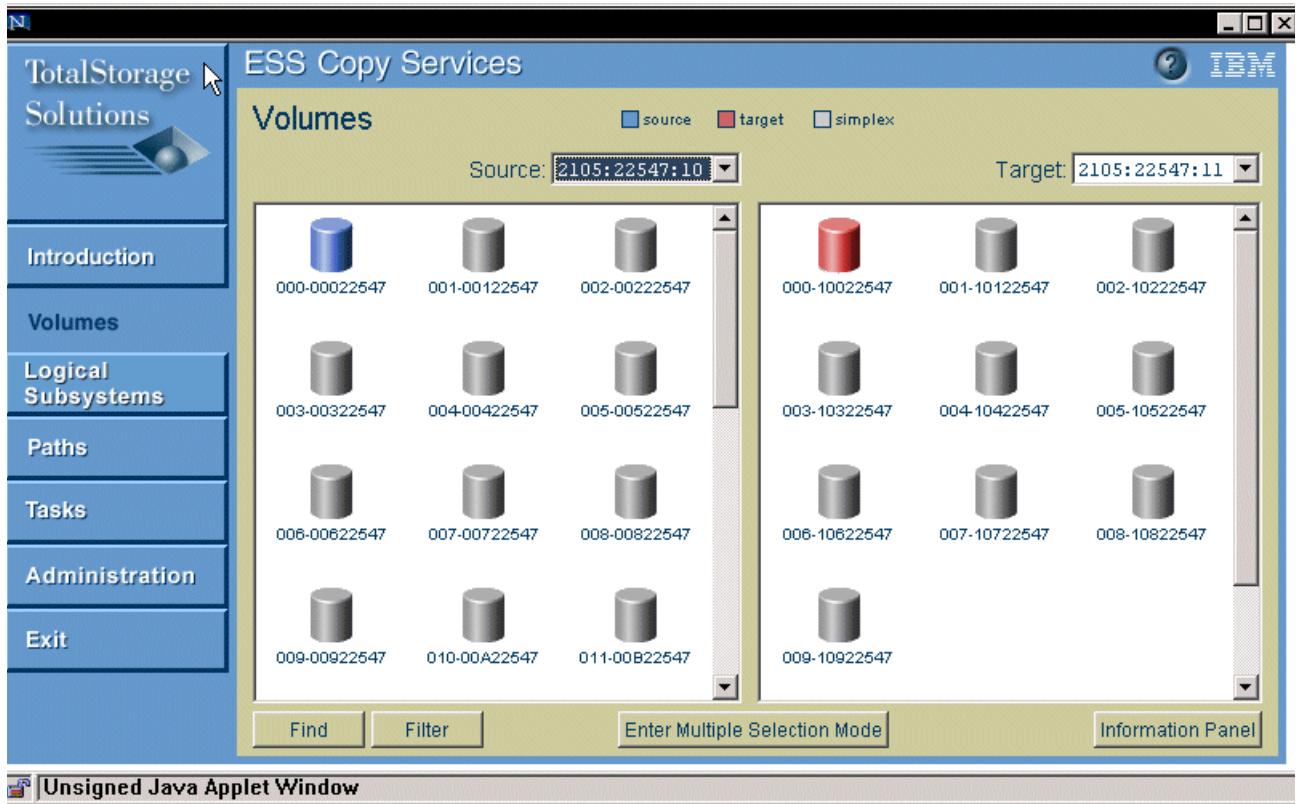
Figure 5-52   PPRC pair in sync

Select a volume and click the **Information Panel** button to retrieve more information about the status. If the source of a PPRC pair is selected, the number of out-of-sync cylinders that are still left to copy are displayed. Those are the tracks that need to be copied from the source to the target to achieve full copy mode (Figure 5-53 on page 153).
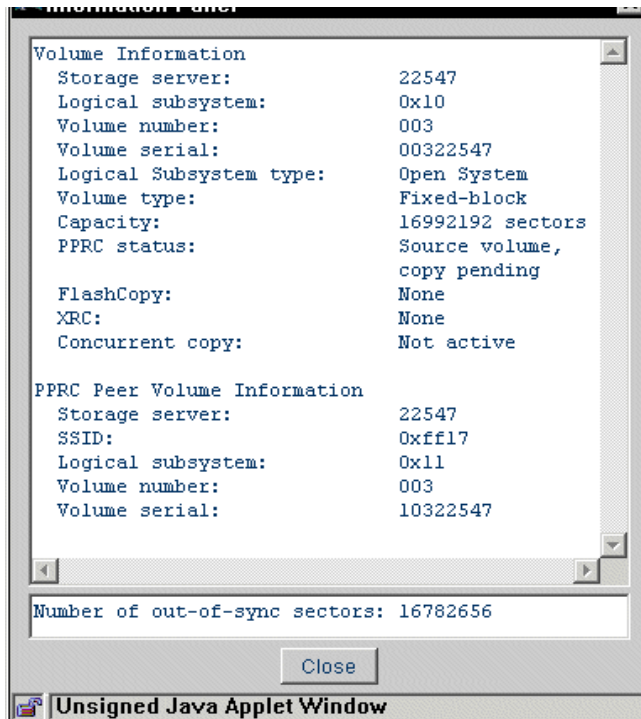
Figure 5-53   Information panel of PPRC source volume

> **Tip:** Alternatively, you can use the Logical Subsystem panel to select **Establish Synchronous PPRC Pairs** for an entire LSS. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

5.11.4 Establishing a PPRC Extended Distance (PPRC-XD) copy pair

PPRC Extended Distance provides the following advantages over the standard synchronous PPRC mode:

► Minimal response-time delays for application hosts that are writing to a PPRC primary device

► Copy distances that can be increased by thousands of miles with channel extenders

► Data migration without the response-time penalty of the synchronous PPRC copy mode

► No FlashCopy of the primary volume required to eliminate the synchronous-write penalty

To establish the PPRC Extended Distance copy pair, use the Volumes panel. This procedure is almost identical to the previous one, which you can use to create a synchronous PPRC copy pair, except that you select the **Establish Extended Distance PPRC** option in the Select task type panel of the Task Wizard (Figure 5-48 on page 148).

Also, the **Do not copy volume** option is not available for PPRC Extended Distance.

In addition, the **Suspend PPRC after established complete** option is not available for PPRC Extended Distance. The reason this requirement is not imposed is that the duplex state is not enforced by PPRC Extended Distance.

5.11.5 Converting a PPRC-XD copy pair to synchronous PPRC

There are two common situations when you would convert a Peer-to-Peer Remote Copy (PPRC) copy pair from PPRC Extended Distance (PPRC-XD) mode to synchronous PPRC:

► Situation 1: You have used PPRC-XD to complete the bulk transfer of data in the creation of many copy pairs, and you now want to convert some or all of those pairs to synchronous PPRC mode.

► Situation 2: You have PPRC-XD copy pairs for which you want to make FlashCopy backups on the remote site. You convert the pairs temporarily to synchronous PPRC mode in order to obtain a point-in-time consistent copy, as discussed in 5.11.6, "Creating a backup copy of volumes via PPRC Extended Distance" on page 156.

For an in-depth discussion of the use of PPRC Extended Distance, see the IBM Redbook, *IBM TotalStorage Enterprise Storage Server: PPRC Extended Distance*, SG24-6568.

You can convert a PPRC Extended Distance copy pair to a synchronous pair using the **Volume** panel. Select the LSS within which you want to convert the PPRC Extended Distance copy pair to a synchronous pair. Click on the primary PPRC volume. In the Target column of the Volumes panel, the secondary Volume of the PPRC pair will be displayed (Figure 5-54).
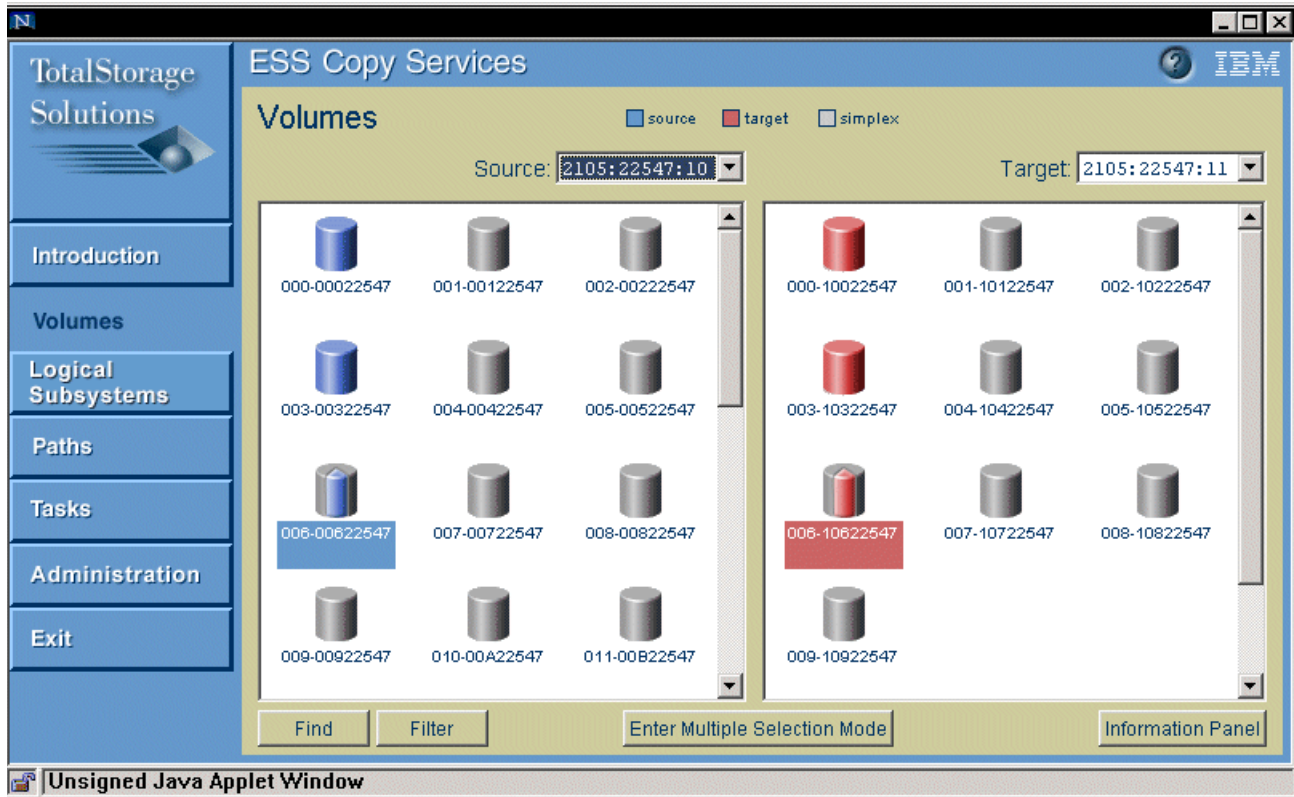
Figure 5-54   Volumes panel: Converting PPRC - XD to synchronous PPRC

Once you have selected the source volume, right-click the target volume, then right-click it again to open the Task Wizard (Figure 5-55).
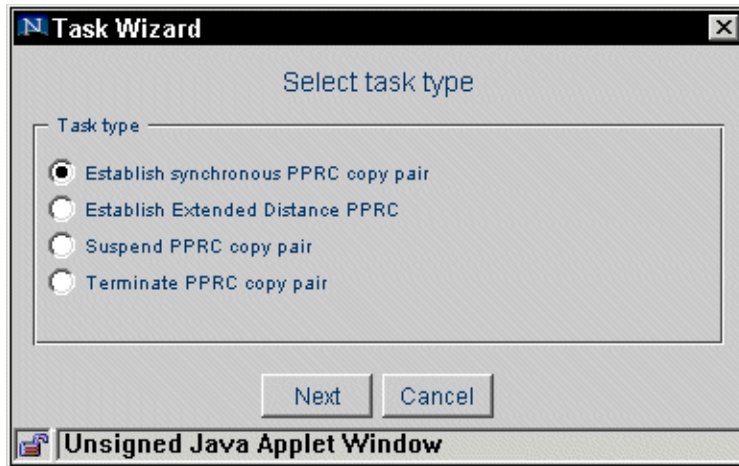
Figure 5-55   Task Wizard: Convert PPRC-XD to synchronous PPRC

Select the **Establish synchronous PPRC copy pair** option and click **Next**. On the next panel, you can select the **Copy out-of-sync cylinders only** option (Figure 5-56).
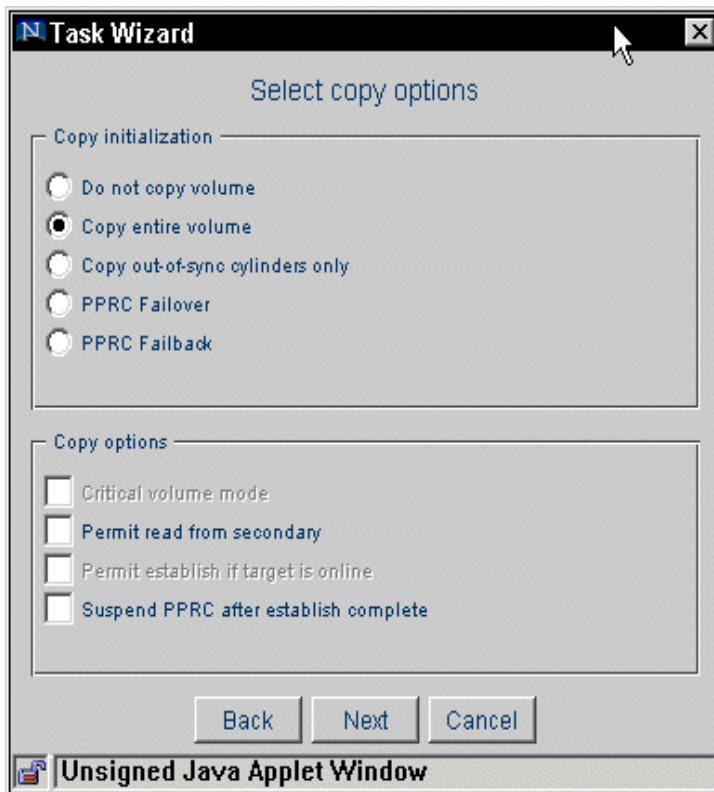
Figure 5-56   Task Wizard: Establish PPRC pair copy options

Click **Next**. From the next panel, you can either **Save**, **Run**, or **Cancel** the task.

Once you run the task, all out-of-sync cylinders will be copied from the primary to the secondary volume.

> **Tip:** Alternatively, you can use the Logical Subsystem Panel to select convert PPRC Extended Distance copy pairs to synchronous PPRC. Running the tasks at the LSS level will perform the task on several volumes. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

5.11.6 Creating a backup copy of volumes via PPRC Extended Distance

When you have used the PPRC Extended Distance (PPRC-XD) mode to transfer data, and the distance between the local ESS and the remote ESS is close enough for you to establish a synchronous PPRC copy pair that can achieve the full duplex state, you can temporarily convert from PPRC Extended Distance mode to synchronous PPRC mode when you want to make a consistent copy of the volumes.

This process would be useful as a back-up copy and for data migration purposes, but would not be particularly useful for disaster recovery if the source data set is subject to rapid change. For further information, refer to 4.3, "PPRC extended distance (PPRC-XD)" on page 59.

Perform the following steps to obtain a backup point-in-time consistent copy of the PPRC-XD volumes that are at the remote site:

1. Change the copy mode for the PPRC Extended Distance copy pair to synchronous PPRC. Select the **Suspend after establish** option in the Task Wizard in creating the synchronous PPRC. This avoids the step of creating a separate task just to suspend the pairs. See 5.11.5, "Converting a PPRC-XD copy pair to synchronous PPRC" on page 154.

2. Alternatively, if you did not select the **Suspend after establish** option, suspend the synchronous PPRC copy pair after it reaches the full duplex state (consistency). See 5.11.7, "Suspending a PPRC pair" on page 156.

3. Perform a FlashCopy on the secondary volumes. See 5.10.1, "Establishing a FlashCopy pair" on page 130.

4. Resume PPRC Extended Distance mode for the copy pair. If you have saved the task that created the original PPRC-XD relationship (see 5.11.4, "Establishing a PPRC Extended Distance (PPRC-XD) copy pair" on page 153"), you can invoke the saved task from the Tasks panel. For details, see 5.7.4, "Modifying a task" on page 127.

5.11.7 Suspending a PPRC pair

When you suspend a copy pair, PPRC stops transferring data to the target volume. Because the primary ESS keeps track of all changed cylinders on the source volume, you can resume PPRC operations at a later time. To resume operations, run a task that resynchronizes the PPRC copy pair. See 5.11.3, "Establishing a synchronous PPRC pair" on page 147.

ESS Copy Services now provides an alternative to creating a separate task that suspends a PPRC pair. When you establish the pair, the Task Wizard provides the option **Suspend PPRC after Establish**. For details, see 5.11.3, "Establishing a synchronous PPRC pair" on page 147".

To suspend a PPRC copy pair, you use the Volumes panel. This procedure is almost identical to the one described in 5.11.3, "Establishing a synchronous PPRC pair" on page 147, except that you select the **Suspend PPRC copy pair** option in the Select task type panel of the Task Wizard (Figure 5-57 on page 157).
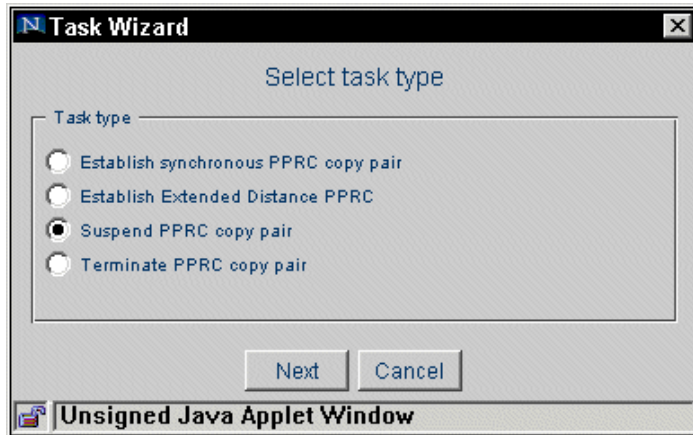
Figure 5-57   Task Wizard: Suspend PPRC copy pair

If you click **Next** within the next panel, you can decide whether to schedule the task with the source volume or to schedule the task to the target volume (Figure 5-58).
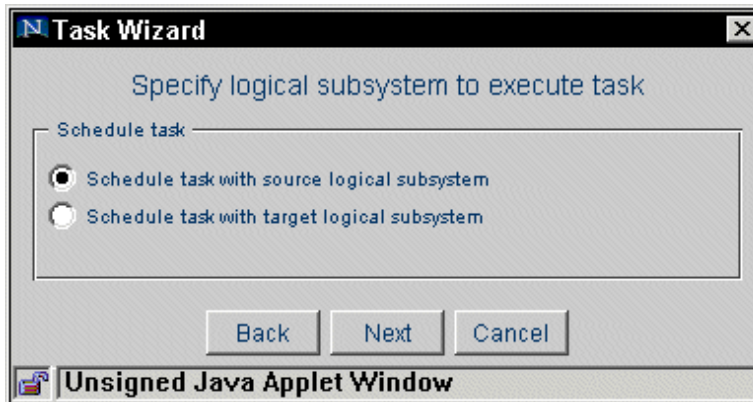
Figure 5-58   Task Wizard: Specify logical subsystem to execute task panel

The task runs on either the source or target ESS. For example, if the source LSS is not available, you can schedule a suspension of a PPRC pair with the target LSS. This is useful in a disaster situation when the primary site has gone down.

Note:

► If you schedule the task with the source logical subsystem, the source will become suspended. As a result, the target will become suspended immediately after.

► If you schedule the task with the target logical subsystem. The target will become suspended. But the source might not become suspended. You may use this option only in a disaster situation when the source LSS is not available anymore.

5.11.8 Terminating a PPRC copy pair

When you terminate a PPRC copy pair, the PPRC relationship between the source and target volume ends. The source and target volumes are removed from the PPRC configuration. You can either use this function to terminate a synchronous PPRC pair or an Extended Distance PPRC pair.

You can terminate a PPRC copy pair in the following ways:

► By modifying a task used to establish or suspend a PPRC copy pair that has been saved (refer to 5.7.4, "Modifying a task" on page 127 for further information)

► By using the Volumes or the LSS panel

This procedure is almost identical to 5.11.3, "Establishing a synchronous PPRC pair" on page 147, except that you select the **Terminate PPRC copy pair** option in the Select task type panel of the Task Wizard (Figure 5-59).
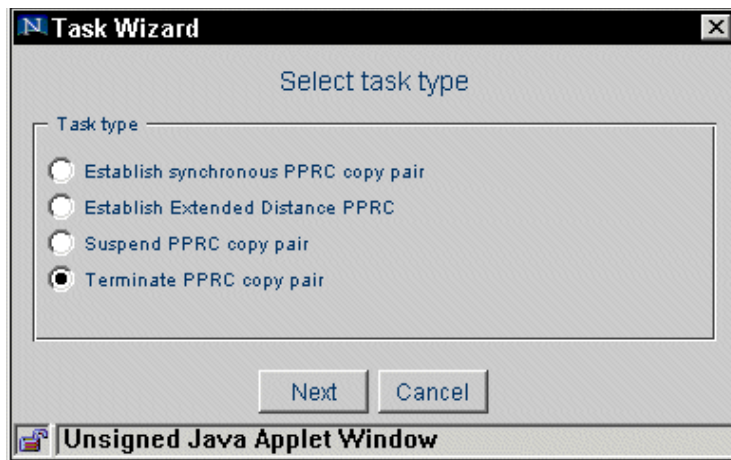
Figure 5-59   Task Wizard: Terminate PPRC copy pair

You will have two options when creating the terminate PPRC pairs: You can decide whether to **Schedule task with the source logical subsystem** or to **Schedule task with the target logical subsystem**.

Note:

► If you schedule a **terminate PPRC pair** task with the *source* logical subsystem and you run the task successfully, the source and the target volume become in simplex state.

► If you schedule a **terminate PPRC pair** task with the *target* logical subsystem and the task has been completed successfully, the source volume is in suspended state, and the target volume is in simplex state. This option is useful in a disaster situation when the primary site has gone down.

5.11.9 Resynchronizing PPRC copy pairs

When you resynchronize a copy pair, all changed cylinders on the source volume are copied to the target volume. PPRC synchronous operations are then resumed for the pair. Use the method described in 5.11.5, "Converting a PPRC-XD copy pair to synchronous PPRC" on page 154 to resynchronize a PPRC pair in suspend state.

Use the Volumes panel to resynchronize a PPRC copy pair. On the left side, select the source LSS, and on the right side, the target LSS. This is done using the drop-down menu at the top of the Volumes menu. Select the source volume and, with a right-click, the target. Once you have selected the source and the target, you do a second right-click on the target to bring up the Task Wizard (Figure 5-60 on page 159).
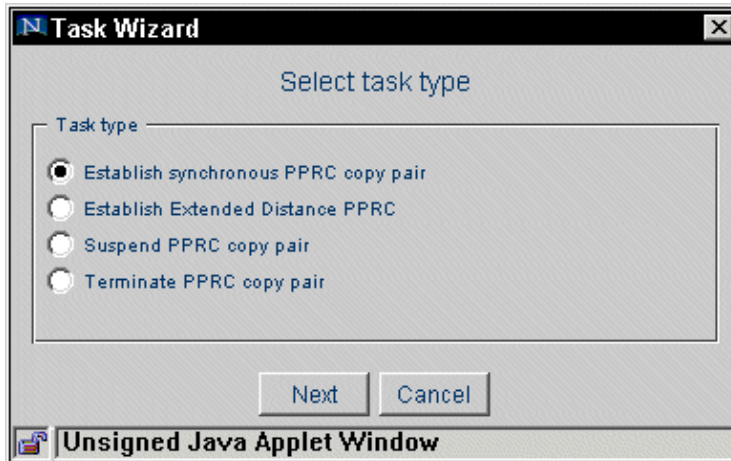
Figure 5-60   Task Wizard: Establish synchronous PPRC copy pair

Select the **Establish synchronous PPRC copy pair** option and click **Next**. In the next panel, you can select the copy options (Figure 5-61).
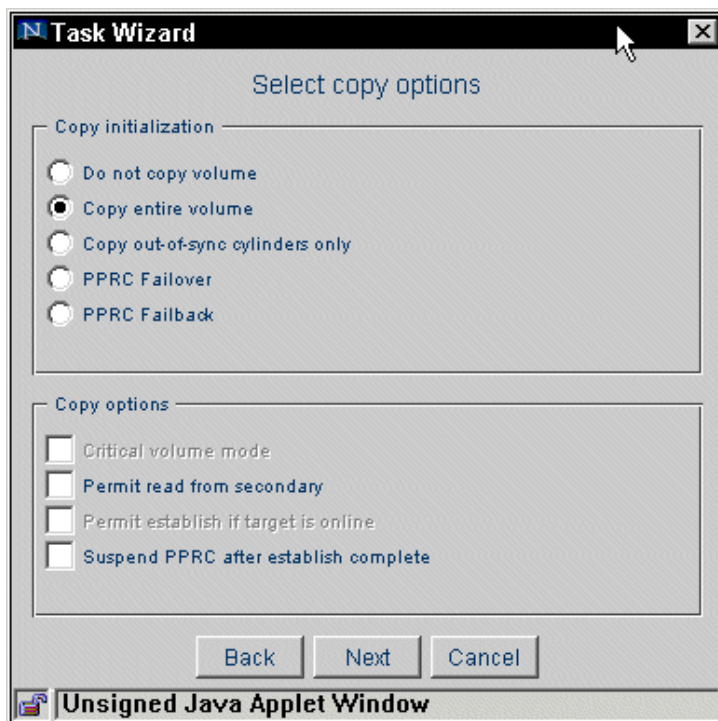
Figure 5-61   Task Wizard: Establish synchronous PPRC copy pair options

Select the **Copy out-of-sync cylinders only** initialization option and decide whether to use the following two options:

► **Permit read from secondary:** Selecting this option allows host servers to read from the PPRC secondary volume. The PPRC pair must be in a full-duplex state in order for the host server to read the volume. Select this option if you have an operating system that uses a Logical Volume Manager that reads data structures from the header of a secondary volume to update internal databases such as the ODM under AIX, or /etc/lvmtab under HP-UX. For further information, refer to Chapter 8, "Open systems specifics" on page 327.

► **Suspend PPRC after establish complete:** If you select the full-copy option described above, then selecting this option means that after the data is completely copied to the target volume, the target volume goes into the suspended state. Selecting this option if you want to suspend the pair right after the primary and secondary volume are in sync.

Click **Next**. From the next panel, you can either **Save**, **Run**, or **Cancel** the task. Once you run the task, all out-of-sync cylinders are copied from the primary to the secondary volume.

> **Tip:** Alternatively, you can use the Logical Subsystem Panel to select **Resynchronize PPRC Copy Pairs**. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

5.11.10 Creating a PPRC consistency group

Use the Paths panel to create a PPRC consistency group, which is a set of PPRC volume pairs that have the same source and target LSS. When an error occurs that affects any of these volumes, the ESS causes the volume where the error is detected to enter a suspended state. If the volume is participating in a consistency group, it enters a queue full state. The primary host temporarily queues all the updates to the primary volumes. ESS Copy Services creates the consistency group when you run a task to establish a path between the primary and secondary LSSs. For further information about PPRC consistency groups, refer to 4.5, "PPRC data consistency" on page 72.

To create a PPRC consistency group, click **Paths** in the navigation frame of an ESS Copy Services panel. The Paths panel opens (Figure 5-18 on page 123). From the Select Source Subsystem list, select the logical subsystem (LSS) from which you want to create the paths. Click the path in the Path Origin column. The ESS automatically displays the physical targets that you can access through this connection in the Common Storage Server Targets column. Click a physical target to select it. The ESS displays the LSS targets that you can access from this physical target in the Logical Subsystem Targets column. Select the target storage server that you want to select. Right-click the same target to open the Select task type panel of the Task Wizard.
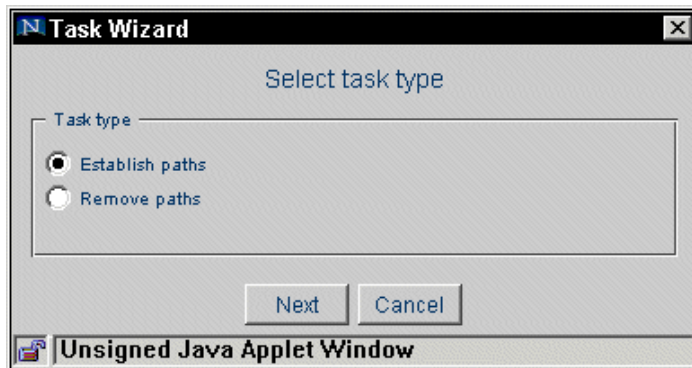
Figure 5-62   Task Wizard: Establish paths

Click **Next** and select the path options in the next panel (Figure 5-63 on page 161).
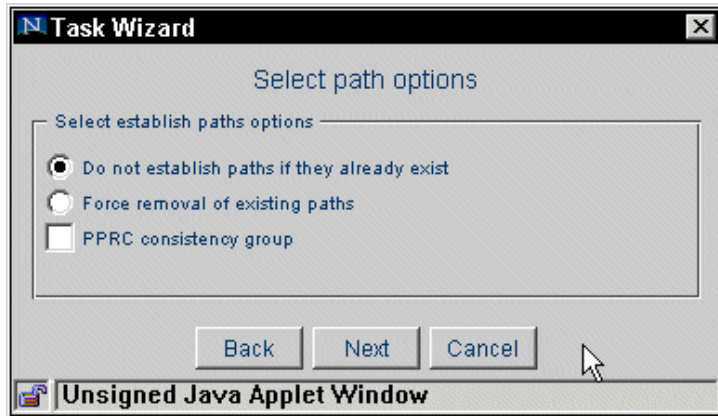
Figure 5-63   Task Wizard: Select path options panel

Select **PPRC consistency group** and click **Next**. The Define task panel of the Task Wizard opens (Figure 5-64).
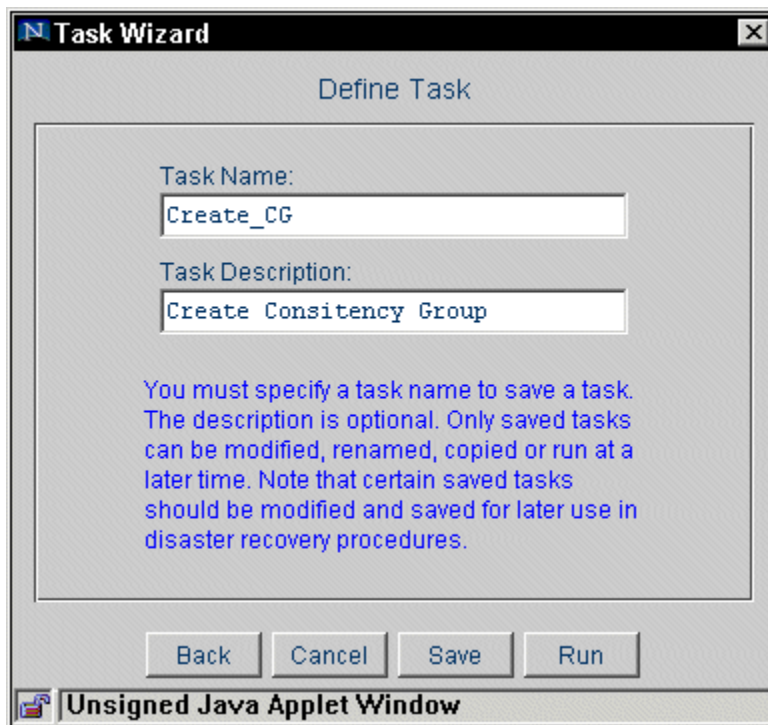
Figure 5-64   Task Wizard: Define Task panel

You can either click **Save** to save the task and run it later, or click **Run** to run the task immediately.

> **Attention:** When you create paths for PPRC consistency groups, it is important to save the task. After you run the Freeze command and corrected the problem that caused the suspension, you need to re-establish the paths. You must re-establish the paths before you run the task that resynchronizes the PPRC volumes.

> **Note:** You can also create a consistency group including all the volumes into one LSS by clicking on the PPRC Consistency Group check box in the LSS properties panel ("PPRC consistency group" on page 120).

5.11.11 Freezing a PPRC consistency group

Freezing a PPRC consistency group stops all operations from being propagated to the secondary volumes. The freeze operation is best implemented by creating and saving tasks for all the associated LSS pairs and then grouping the tasks. When you need to perform a freeze operation, the group can be run from the Task panel or from the command-line interface. See 5.7.1, "Grouping and ungrouping tasks" on page 125 for information on grouping tasks.

For the freeze operation to work and create a consistent image of the data at the remote site, the Establish Paths task must indicate that the LSS to LSS relationship is part of a consistency group (see 5.11.10, "Creating a PPRC consistency group" on page 160).

The task is invoked when a freeze trigger event (for example, an e-mail notification, an SNMP alert, or a host-detected failure condition) is detected. Following the freeze operation, a consistency-group-created operation should be run to stop any long busy states that the volumes are in. See 4.5, "PPRC data consistency" on page 72 for more information about the freeze and consistency-group-created operations.

To create a task that freezes a consistency group, click on the **Logical Subsystems** button in the navigation frame of an ESS Copy Service panel. The Logical Subsystems panel opens, and within the Select list in the upper right corner of the panel, you select the source ESS. A panel showing all the logical subsystems (LSSs) is displayed, with the appropriate names and serial numbers indicated below each LSS icon (Figure 5-65 on page 163). Click on the source LSS and click the select list again. Select the target ESS. The LSSs of this ESS are displayed and you right-click the target LSS. The LSS ID of the target LSS turns red.
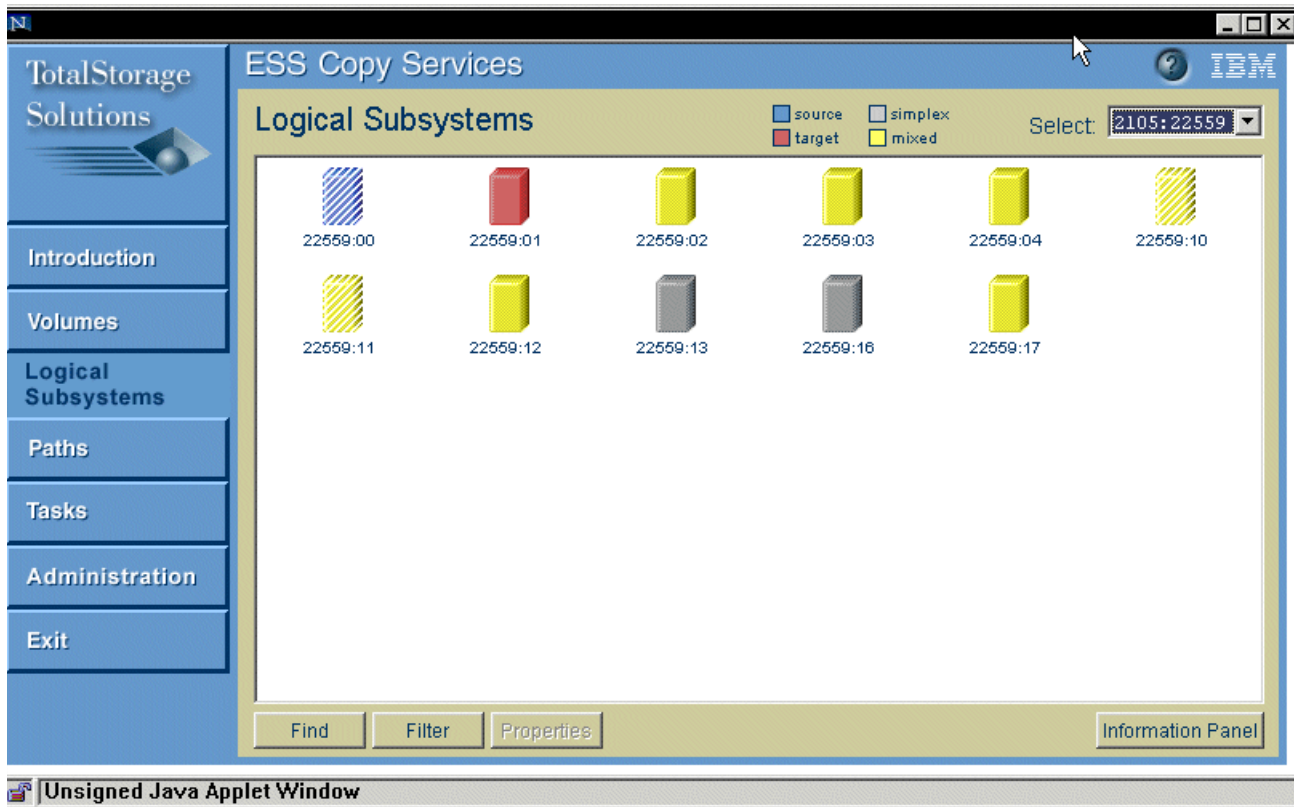
Figure 5-65   Logical Subsystems panel

Right-click the target LSS again to open the Task Wizard. Select **Freeze PPRC consistency group** (Figure 5-66).
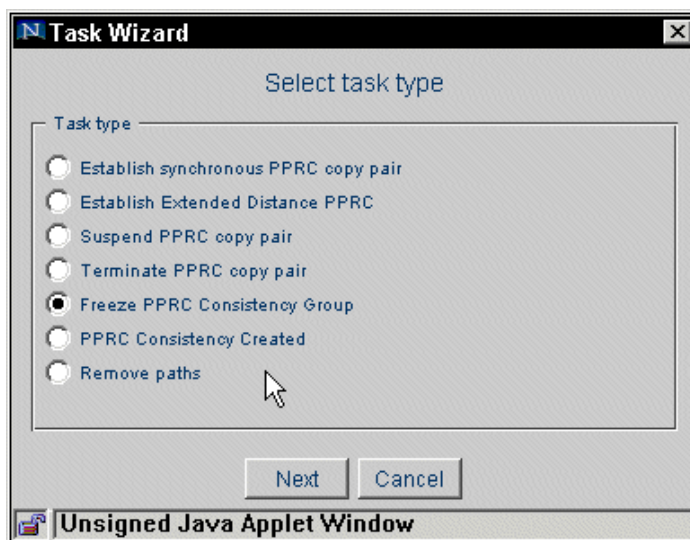
Figure 5-66   Task Wizard: Select task type

Within the next panel, you can decide whether to **Run**, **Save**, or **Cancel** the task. If you want to save the task, type the task name and the task description in the appropriate fields. Click **Run** to execute the task immediately.

> **Attention:** After creating and saving the task to freeze a PPRC consistency group, you should create the corresponding consistency-group-created task. We recommend that you modify the task that you just created so that it runs a consistency-group-created operation, and then name and save the task. For further information about modifying a task, refer to 5.7.4, "Modifying a task" on page 127.

5.11.12 Thawing a PPRC consistency group

The thawing of a PPRC consistency group (consistency-group-created operation) follows the successful freeze for all LSSs in the consistency group. It allows normal I/O operations from the host to the ESS to resume. For further information, refer to 4.5, "PPRC data consistency" on page 72).

You can either use the Task panel to modify a **Freeze PPRC Consistency Group** task to create an appropriate task to thaw the consistency group (refer to 5.7.4, "Modifying a task" on page 127 for further information), or you can use the Logical Subsystems panel to create a task to thaw a PPRC consistency group.

This procedure of using the Logical Subsystems panel is almost identical to the previous procedure, which you can use to freeze a PPRC consistency group, except that you select the **PPRC consistency created** option in the Select task type panel of the Task Wizard (Figure 5-66 on page 163).

5.11.13 PPRC Failover and Failback

PPRC Failover and Failback operations can be used in a disaster recovery process. The tasks are presented in detail in Chapter 10, "Disaster recovery" on page 399. They will be combined to restart the applications using the volumes at a recovery site, and establish a PPRC from the recovery site to the primary site.

We will discuss here the behavior of the two tasks, without giving too many details about the status of the volumes. To get more information, you can refer to 6.12.13, "Performing a PPRC Failover" on page 266 and 6.12.14, "Performing a PPRC Failback" on page 269.

PPRC Failover

This task is part of the process used to restart a production environment using the volumes from the recovery site. See Chapter 10, "Disaster recovery" on page 399 for details. In a disaster recovery process, the fail-over task is followed by a fail-back task once a path from the recovery site to the production site is established.

Use the Volumes panel to establish the PPRC pairs. The process is identical to the one used when establishing a PPRC relationship.

When creating the PPRC Failover task you must be very careful when choosing the source volume and the target volumes for the task:

► The secondary volume of the established PPRC pair should be selected as the primary (source) volume for the new PPRC Failover task. Click to select the volume.

► The Primary of the established PPRC pair should be selected as the secondary (target) volume for the new PPRC Failover task. Right-click to select the volume.

The failover is invoked by selecting the **Failover** option in the Task Wizard panel (Figure 5-67 on page 165).
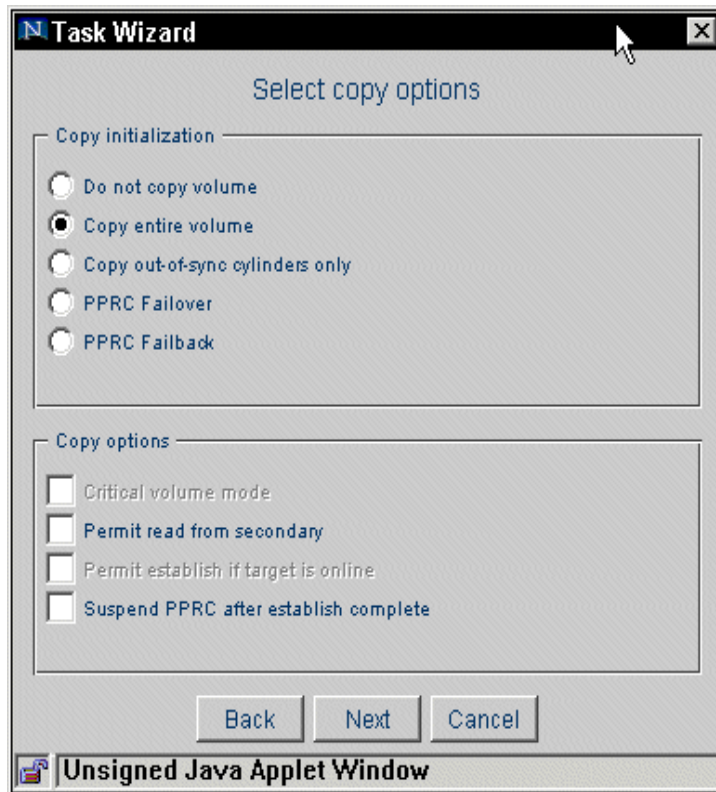
Figure 5-67   Task Wizard: Establish PPRC pair copy options

Once the PPRC Failover task is established on a synchronous PPRC pair, the volumes states change:

► The initial primary volume at the production site remains unchanged. Its state is "Source volume of a PPRC relationship". The target volume of this relationship is the volume at the recovery site. The status of the volume at production site is "Full duplex" of the "Synchronous PPRC" relationship. If you perform any write operation on this volume, its status will change from "Full duplex" to "Suspended".

► The initial secondary volume at the recovery site changed from "Target" volume to the new state of "Source volume of PPRC relationship". The target of this relationship is the initial primary volume at the local site. The status of the volume at recovery site is "Suspended" of "Synchronous PPRC" relationship.

After the PPRC Failover task is established, the volume at the recovery site is available to be mounted and used by a host. You can use this volume to restart your application from the recovery site.

PPRC Failback

This task is part of the process used to restart a production environment using the volumes from the recovery site (see Chapter 10, "Disaster recovery" on page 399, for details). The fail-back task is generally performed after a fail-over task.

After establishing a PPRC **Failover** task, the volume at the recovery site is available to be mounted and used by a host. You can use this volume to restart your application from the recovery site (see previous section). Establishing a PPRC with the **Failback** option will resynchronize the PPRC volumes at the recovery site with their peer volumes at the production site.

Important: Before you run the PPRC Failback task, it is necessary to establish a PPRC path from the recovery site to the production site between the concerned LSSs.

Use the Volumes panel to establish the PPRC pairs. The process is identical to the one used when establishing a PPRC relationship.

When creating the PPRC Failback task, you must be very careful when choosing the source volume and the target volumes for the task.

► The primary (source) volume of your new PPRC Failback task must be the volume at the recovery site. This volume was initially the secondary volume of the PPRC relationship. After the PPRC Failover task established, this volume state changed to become a primary (source) volume in a "Suspended" state. Click to select the volume.

► The secondary (target) volume of your new PPRC Failback task must be the volume at the production site. This volume is the primary volume of the initial PPRC relationship. After the PPRC Failover task established, this volume state remained unchanged. If you perform any write operation to this volume after the PPRC Failover task was established, then its status changes to "suspended". Right-click to select the volume.

The failback is invoked by selecting the **Failback** option in the Task Wizard window (Figure 5-68).
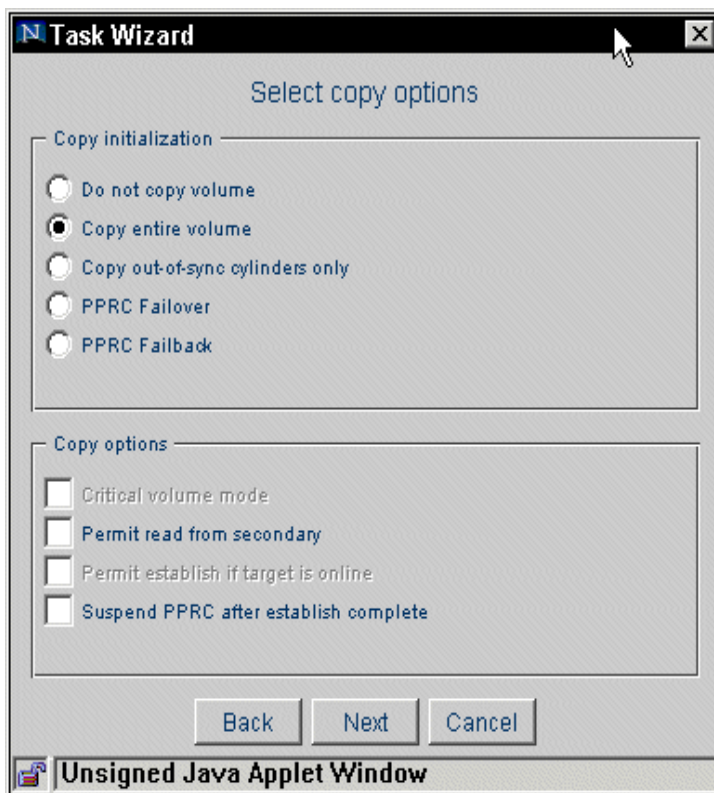
Figure 5-68   Task Wizard: Establish PPRC pair copy options

Once the PPRC Failback task is established, the PPRC relationship is resynchronized between the recovery site and the production site.

► The volume at the production site is now the secondary (target) volume of the PPRC relationship.

- ► The volume at the recovery site is now the primary (source) volume of the PPRC.
- ► The status of the two volumes will be first "Copy pending" than "Full duplex" when the copy of the out-of-sync sectors is finished.

6

ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0

There are three different methods of managing ESS Copy Services in the open systems environment:

► A Web-based Interface

► A Java-based Command Line Interface (CLI)

► And, starting with LIC level 2.3.0. the ESS API has been enhanced to support Copy Services configuration and use for PPRC and FlashCopy

Very important to note is the fact that the Web-based user interface panels and icons have had changes with LIC 2.2.0. first and later with LIC 2.3.0. These changes were not only to improve the look-and-feel for the users, but also related to the new functions that became available: LIC 2.2.0. brought in PPRC Version 2 and FlashCopy Version 2, and LIC 2.3.0. brought in PPRC over Fibre Channel links.

In this chapter, we explain how to use and set up the ESS Copy Services Web User Interface (WUI) when the ESS is at the minimum LIC level 2.3.0. Therefore, this chapter describes how to use the Web user interface when you are using PPRC Version 2 and FlashCopy Version 2, as well as when Fibre Channel links are used.

The Web user interface for ESSs having LIC levels prior to 2.2.0. is discussed in Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103.

The usage of the Command Line Interface is described in Chapter 7, "ESS Command Line Interface" on page 301.

The usage of the ESS API is described in Appendix A, "ESS Application Programming Interface (API)" on page 471.

6.1 Overview and requirements

ESS Copy Services runs within the IBM TotalStorage Enterprise Storage Server (ESS). Since LIC code 2.2.0, you can define two servers in the copy services domain to manage the copy services in two ways: the dual-active servers mode and the mixed mode (one active server and one passive server). The active servers are responsible for maintaining all Copy Services related information. One of the ESS clusters in the Copy Services domain has to be defined as the active server (ServerA). Optionally, you can define a second ESS cluster in the Copy Services domain to act as the second active server (ServerB).

On each ESS cluster that is intended to use Copy Services, there is a Copy Services client running that communicates with the Active Copy Services server.

Refer to Chapter 2, "Implementing ESS Copy Services" on page 13, to get more information about the Copy Services Domain, and the different modes (dual-active and mix-mode) for these domains.

Access to Copy Services is provided through a Web browser. Using a Web browser offers the possibility to easily control the ESS copy functionality over the network from any platform for which the browser is supported.

A Web server running in the ESS provides a Web interface that you can use to manage the ESS through a Web browser. The Web browser must be running on a workstation connected to the ESS through the IBM TotalStorage Enterprise Storage Server Network (ESSNet). The ESSNet is the ESS access facility that IBM installs when they install your ESS. The ESSNet consists of a switch (hub) and other networking components and the dedicated IBM TotalStorage Enterprise Storage Server Master Console (ESS Master Console).

In addition to using a Web browser on the ESSNet Master Console to connect to your ESS, you can also use a Web browser running on your own workstation, either by connecting your workstation directly to the ESSNet hub, or by connecting your workstation to your intranet and connecting your intranet to the ESSNet hub. If you use your own workstation, IBM recommends that it have at least 128 MB of memory.

The ESS Copy Services require one of the following Internet browsers:

► Netscape Communicator
► Microsoft Internet Explorer (MSIE)

For supported versions of Netscape Navigator and Internet Explorer, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.2 ESS Copy Services WUI

You can either use the ESS Master Console or a Web browser installed on a workstation connected through the ESSNet to access the ESS.

The ESS Master Console differs from other workstations in the way you access your ESS. The desktop of the ESS Master Console has a Netscape icon with the label ESS Specialist Launcher. That icon accesses a panel that provides a list of the IP addresses of your ESS clusters. Clicking on the address of the cluster that you wish to access will cause the resident ESS Launch panel in the selected cluster to open.

If you use a workstation other than the ESS Master Console to access your ESS, you must enter the desired cluster address in the Address field of your Web browser. This address,

which you click on in the *Specialist Launcher* of your ESS Master Console or that you type in the Address field of your Web browser, is the host name alias or the dotted decimal IP address of one of the clusters of your ESS. IBM configures these addresses in your ESS at installation time, based on the addresses you enter in the Communications Resources work sheet that you submit for your ESS. For details, see the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, GC26-7444.

The first page to open, the ESS Launch panel in Figure 6-1, is the same for either of the two clusters in the ESS, so access to either cluster allows you to configure the entire machine.

The ESS Launch panel contains two buttons related to Copy Services:

► The Copy Services button to connect to one active Copy Services server.

► The Tools button to connect to the Copy Services Web domain configuration and the copy services Help panel.
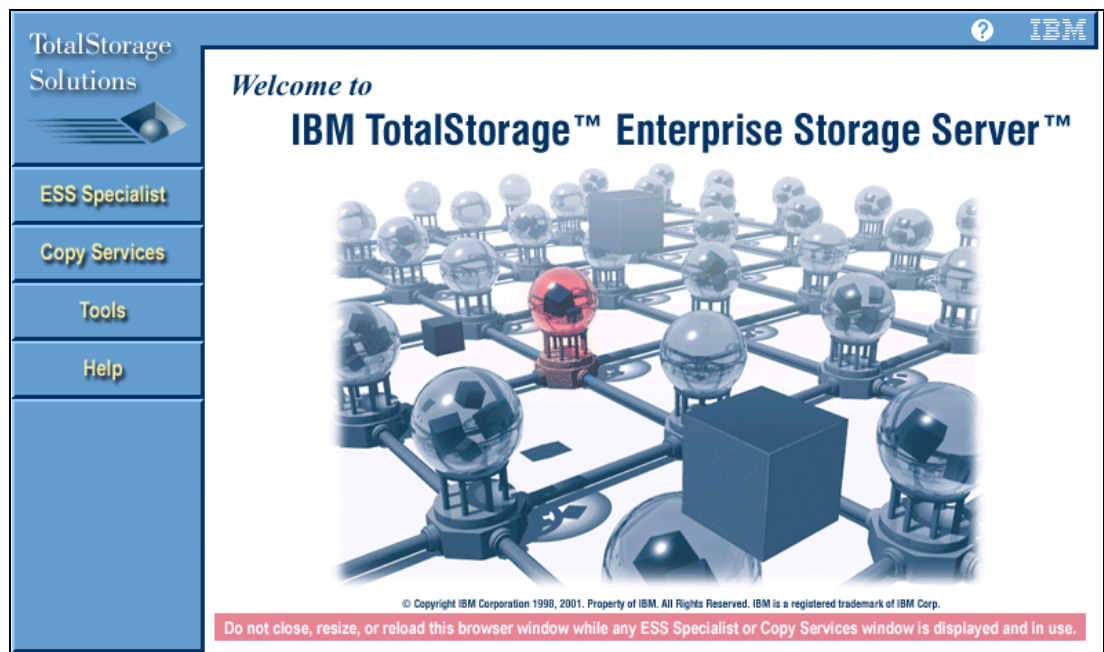
Figure 6-1   ESS welcome panel

6.3 Web Copy Services Domain: Tools and configuration

Since Version 2.2.0, the Copy Services allow the user to define the WEB Copy Services Domain using a browser. This configuration is performed using the different options accessed via the Tools panel. Since LIC 2.3.0, the Copy Services Domain supports up to eight ESSs.

Before going further, refer to Chapter 2, "Implementing ESS Copy Services" on page 13 to clearly understand all the concepts of the Copy Services Domain.

6.3.1 Copy Services Tools main menu

From the ESS Launch panel, click **Tools**. A new panel will appear with "Copy Services Main Service page" (Figure 6-2 on page 172).

Tools: Copy Services Main Service Page

- Determine IPs for Copy Services active servers
- Define Copy Services active servers
- Define Copy Services clients for the active server domain
- Recover from an 'Unable to Connect to the Copy Services Server' message
- Restart Copy Services on this cluster
- Perform a Domain-wide Reset
- Disable Copy Services

Determining IPs for Copy Services active servers

The active Copy Services server is the ESS cluster that is currently running the Copy Services server code.

Display the active Copy Services servers

Define Copy Services active servers

To define the Copy Services servers, enter an IP address for each of the active servers.

This action disables ESS Web Copy Services. When you perform this action, be aware that:

Figure 6-2   Copy Services Tools panel

From this page, you will be able to perform the following operations:

- ► Determine IPs for copyservices active servers
- ► Define copyservices active servers
- ► Define copyservices clients
- ► Recovering from an `Unable to Connect to the Copy Services Server` message
- ► Restart copyservices on this cluster
- ► Domain Wide Reset
- ► Disable copyservices
- ► Launch the Trouble Shooting and Disaster Recovery help

Click on the links to understand what each action will perform and to access the **Available actions** menu presented in Figure 6-3 on page 173.

6.3.2 Defining and restarting the Web Copy Services Domain

A Copy Services domain consists of two servers (ServerA and ServerB) and the clients that are associated with those servers. To set up your domain you must register the clusters and clients information using the Tools panel.

The copy services WEB domain is described in detail in Figure 6-3.

Once all the ESSs in the domain are connected into the same LAN, the actions you perform from the Tools panel are:

1. On *each ESS* in the copy services domain: Define the active servers for the Copy Services WEB domain.

2. On *the two clusters acting as ServerA or serverB*: Define the clients in the Copy Services WEB domain.

3. On *each ESS* clusters in the domain in the domain: Start the Copy services.
   An alternative to the restart process described above is the *Domain Wide Reset* action.

To perform these configuration and actions, use the buttons from the Tools panel (see Figure 6-3).

Action	Description
Restart	Restart Copy Services on this cluster.
Disable	Disable Copy Services on this cluster.
Domain-wide Reset	Perform a Domain-wide Reset.
Define Servers	Define Copy Services active servers.
Define Clients	Define Copy Services clients.
Cancel	Return to the Tools main page, without performing any action.

Figure 6-3   Copy Services action buttons from the Available Actions panel

Defining the active servers

Starting with LIC 2.2.0, dual-active servers for copy services in the same domain are supported. The two servers are designated as ServerA and ServerB. Depending on your domain configuration, the two servers will act in dual-active mode or in mix-mode. See Chapter 2, "Implementing ESS Copy Services" on page 13.

The definitions of the two servers is performed using the Tools panel on each ESS in the copy services domain.You need to perform this definition for all of the ESSs in the Copy Services Domain. Once the definition is submitted on one cluster (using the "Submit Configuration" button in Figure 6-4 on page 174), the configuration is automatically updated to the alternate cluster on the ESS.

You define the servers by entering the IP addresses of the two clusters in the **Define Copy Services Active servers** panel. This panel is accessed by clicking on the **Define servers** button. An IP address for serverA is mandatory; defining ServerB is optional. If you do not define ServerB, then you will not have dual-active mode for the copy services, since only one server will be active for the Copy Services.

Define CopyServices Active Servers

Enter IP Addresses Only

Active ServerA: 9.82.22.61
Active ServerB: 9.82.36.41

Submit Configuration   Reset Values

Figure 6-4   Define Copy Services Active Servers

Defining the clients

After you define the Active servers, the next step consists of defining the list of all the ESS in the Copy Services Domain. With LIC 2.3.0 installed, up to eight ESSs can be defined in a Domain. The two active servers (serverA and serverB) must appear in this list.

You will define the Copy Service Clients on the two clusters defined as active servers. Defining the clients on ServerB is not mandatory but highly recommended. If you do not define the clients on ServerB, you cannot use the **Domain Wide Reset** action.

Define the Copy Services clients by clicking the **Define Clients** button. Input the information requested for each of the client clusters (see Figure 6-5 on page 175).

Define CopyServices Clients

Enter Client Information For CopyServices Domain

Client1 Machine Serial Numbers:	18767
Client1 ClusterBay0 IP Address:	9.82.36.41
Client1 ClusterBay0 Host Name:	t18767c1.wsclab.washin
Client1 ClusterBay1 IP Address:	9.82.36.42
Client1 ClusterBay1 Host Name:	t18767c2.wsclab.washin
Client2 Machine Serial Numbers:	22331
Client2 ClusterBay0 IP Address:	9.82.22.61
Client2 ClusterBay0 Host Name:	t22331c1.wsclab.washin
Client2 ClusterBay1 IP Address:	9.82.22.62
Client2 ClusterBay1 Host Name:	t22331c2.wsclab.washin
Client3 Machine Serial Numbers:	undefined
Client3 ClusterBay0 IP Address:	undefined
Client3 ClusterBay0 Host Name:	undefined
Client3 ClusterBay1 IP Address:	undefined
Client3 ClusterBay1 Host Name:	undefined

page can be scrolled down for up to 8 Clients

Figure 6-5   Defining the clients

Once all the clients are defined, click on the **Submit Configuration** button, as in Figure 6-6, to update the active server with the configuration.

Figure 6-6   Define clients button

Start/Restart Copy Services

This operation is mandatory after you define the active servers and the client to be able to use Copy Services options.

To restart Copy Services, click the **Tools** button in the **ESS welcome** panel to open the **Tools** panel. On the **Tools** panel, click on the **Recover from an Unable to Connect to the Copy Services Server** message to link to the restart procedure. Be sure you understand the result of implementing a restart before proceeding. When you use this function, be aware that:

► You will lose any PPRC or FlashCopy tasks that have not completed.

► Established PPRC and FlashCopy relationships are maintained.

► You cannot submit any additional tasks through the CLI until ESS Copy Services has re initialized.

► You must exit the Copy Services panel and re-enter it to receive valid updates from the Copy Services server.

If you decide to restart ESS Copy Services, click the **Reset CopyServices** button. A warning message, `Resetting the ESS Copy Services`, appears. Click on **OK** to perform the reset.

> **Attention:** Use this function only if you are certain that no other recovery options are available.

Domain Wide Reset

You can use the Domain Wide Reset button to reset each cluster within the Copy Services Domain with the current active server.

> **Note:** Only ESSs that are running LIC level 2.2.0 or higher will be reset

To use the Domain Wide Reset function, you must have a common Web Administrator level user ID and password defined on all ESS clients in the ESS Copy Service Server domain.

Copy Services clients must be configured on the ESS clusters that are defined as Server A and Server B to support the Domain Wide Reset function.

To perform a Domain Wide Reset, click the **Tools** button in the **ESS welcome** panel to open the **Tools** panel. On the **Tools** panel, click on the **Perform a Domain Wide Reset** message to link to the Domain Wide Reset procedure. Be sure you understand the result of implementing a Domain Wide Reset before proceeding. When you use this function, be aware that:

► You will lose any PPRC or FlashCopy tasks that have not completed.

► Established PPRC and FlashCopy relationships are maintained.

► You cannot submit any additional tasks through the CLI until ESS Copy Services has re initialized.

► You must exit the Copy Services panel and re-enter it to receive valid updates from the Copy Services server.

► This function will not be successful on the cluster that does not have the Domain Wide Reset code installed.

► For the clusters that do not have the Domain Wide Reset code, the reset has to be performed on the each cluster separately.

If you decide to restart ESS Copy Services, click the **Domain Wide Reset** button. A warning message, `Resetting the ESS Copy Services`, appears. Click on **OK** to perform the reset.

> **Attention:** Use this function only if you are certain that no other recovery options are available.

6.4 Connecting to the Copy Services Web User Interface

When you click on the Copy Services button in the Welcome panel (Figure 6-1 on page 171), a new panel will prompt you to select which active server you want to connect to (Figure 6-7 on page 177).

By clicking **Start with serverA** or **Start with serverB**, you will open the Copy Services Web user interface in another panel. If you have not previously selected one of those two buttons, you will be prompted for the user name and the password before starting the Copy Services Web User Interface.

Enterprise Storage Server: Start CopyServices

Server Configuration:

IP Address	Description
9.82.22.61	Currently configured Active serverA.
9.82.36.41	Currently configured Active serverB.

NOTE: to refresh the above display, select the reload/refresh menu option from your browser's context menu for this browser frame (i.e. right-click in this frame and select *Refresh* or *Reload Frame* from the context menu).

Available Actions:

Action	Description
Start with serverA	Start CopyServices with serverA.
Start with serverB	Start CopyServices with serverB.

Figure 6-7   Start CopyServices panel

Attention: Once the Copy Services Web user Interface is started, do not close the Start Copy Services panel, or it will also close the Copy Services GUI panel.

The message panel shown in Figure 6-8 will be displayed while connecting to the Copy Services Server.

IBM Copy Services

Connecting to Copy Services Server, please wait...

Warning: Applet Window
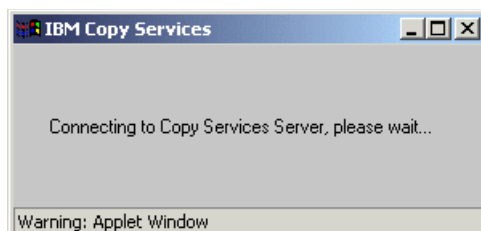
Figure 6-8   Connecting to the Copy Services server

Once the connection to the Copy Services server is successful, the main menu of the ESS Copy Services WUI will be displayed (Figure 6-9 on page 178). From here, you can access all Copy Services menus by selecting one of the buttons on the left side.

Figure 6-9   Copy Services GUI: Welcome panel

There are six buttons in the Navigation panel on the left. The actions and panels available via these buttons are discussed in the sections below:

► **Introduction button**: This button brings you the initial Copy Services panel (Figure 6-9).

► **Volumes button**: Refer to 6.5, "Volume panel" on page 179.

► **Logical Subsystems button**: Refer to 6.6, "Logical subsystems panel" on page 186.

► **Paths button**: Refer to 6.7, "The Paths panel" on page 191.

► **Tasks button**: Refer to 6.8, "Tasks panel" on page 195.

► **Administration button**: Refer to 6.9, "Administration panel" on page 200.

► **Exit button**: Refer to 6.10, "Exiting ESS Copy services" on page 202.

6.4.1 Failure to connect to the ESS Copy Services

If you receive the following message when you click Copy Services in the navigation frame of the ESS Launch page:

```
Failed to connect to Copy Services Server. Server may be down or not configured and cannot
access ESS Copy Services
```

the failure is most likely the result of at least one of the following three conditions:

► The Copy Services active servers are not defined. (You need at least one active server defined).

  – If you have successfully accessed ESS Copy Services before, this would not be the source of the problem.

  – If this is the first time that you are connecting to ESS Copy Services, see 6.3.2, "Defining and restarting the Web Copy Services Domain" on page 173.

► Check for the active server configuration and verify that the Copy Services are running on the servers. If the Copy Services server is not running, see "5.3.3, "Restarting ESS Copy Services" on page 108 for instructions.

► Your Web browser has a temporary operational or connection problem.

> **Note:** For all three conditions noted above, if you have the CLI active on one of your host systems, you can verify that there is a connection to the active Copy Services server with the `rsTestConnection` command. For details, see Chapter 7, "ESS Command Line Interface" on page 301.

For Web browser problems, take these corrective actions:

1. Close any open ESS Copy Services browser panel.

2. Access ESS Copy Services again by clicking **Copy Services**.

3. If step 2 is unsuccessful, close all browser windows (if you have ESS Specialist running, first complete any unfinished task definition), then restart your browser.

4. Again, access the ESS Launch page and click **Copy Services**.

5. If Copy Services still does not launch, consider restarting the workstation that is running your browser.

6.4.2 Restart the ESS Copy Services

If ESS Copy Services is not available, and you fail to connect to any of the active servers, the tools panel proposes two ways to restart ESS Copy Services.

► Reset CopyServices procedure
► Domain Wide Reset procedure

Be aware that you need to have administration authority to do this. Refer to "Start/Restart Copy Services" on page 175 and "Domain Wide Reset" on page 176.

6.5 Volume panel

Volumes are defined with the ESS Specialist in order to provide a fixed storage capacity to the connected host system. They are the base components of each data copy task. The ESS assigns each volume a unique eight digit identifier (ID). This identifier is used to address each volume within the ESS.

From the Volumes menu, you will be able to:

► Find a volume.

► Display volumes based on a filter.

► View information about a volume.

► Establish and withdraw FlashCopy pairs with different options.

► Establish, suspend, and terminate synchronous PPRC copy pairs with different options.

► Establish, suspend, and terminate PPRC Extended Distance copy pair.

► Convert a PPRC Extended Distance copy pair to synchronous PPRC.

► Establish multiple selection mode to select volume-pairs when creating relationships for FlashCopy and PPRC.

► Establish an Asynchronous Cascading PPRC relationship.

The window is divided in two panels (see Figure 6-10 on page 180). In each panel you will display all the volumes in the same LSS.
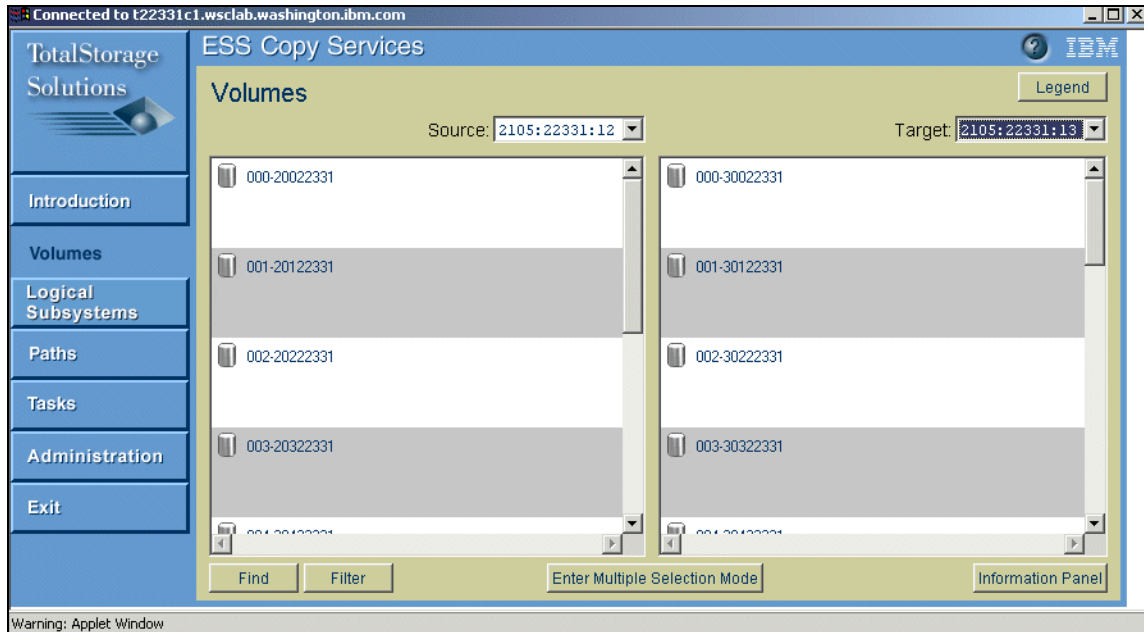
Figure 6-10 Volume panel with source and target selected

► Use the left panel to display and select the volumes you will use as source volumes for FlashCopy or as primary volumes for PPRC.

► You will use the right panel to display and select the volumes you will use as target volumes for FlashCopy or as secondary volumes for PPRC.

To display the volumes shown in Figure 6-10, do the following:

1. Select the source LSS using the scrolling list above the panel on the left

2. Select the target LSS using the scrolling list above the right panel

The source and target logical subsystem are specified in the following way: Device type (4 digits):ESS Serial number (5 digits):LSS number (2 digits). For example, 2105:22331:16 designates the LSS number 16 on the ESS, which has the serial number 22331.

The Command line user interface provides a very useful command that will allow you to display all the volumes assigned to a host in the panels. This command is `rsPrimeServer`.

1. Issue the `rsPrimeServer` command from any host. A list of volumes is sent by the host to the Copy Services Server. Refer to Chapter 7, "ESS Command Line Interface" on page 301.

2. In the Volume panel, select the host name in the source and target list to display all the volumes associated with this host (see Figure 6-11 on page 181).
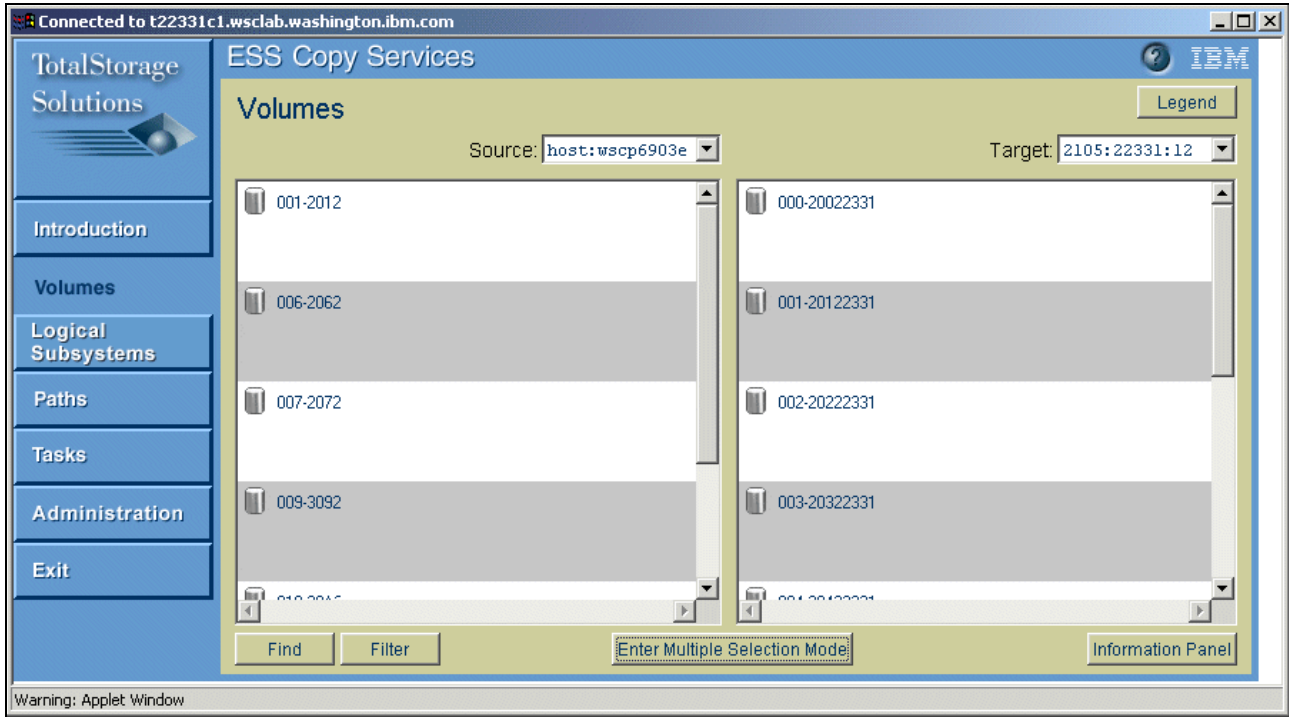
Figure 6-11   Example of results of the rsPrimeServer command

> **Tip:** Additionally, using the CLI command `esscli set volume`, you can add a label on the the volumes to help you identify volume functions. This label will appear below the volume serial number.

6.5.1 Volume panel icons

The Volumes menu shows all volumes defined within one LSS. To the right of each volume icon, you will find its unique serial number. To the right of the serial number, you will find one or more icons. Those icons will describe how the volume is used in any copy services relationship. No icon means that the volume is not in use by any copy relationship.

In the volume panel, you get a short description about the meaning of each icons when you click on the **Legend** button located at the upper right corner of the panel.

Table 6-1 presents the meaning of the different icons and colors in the Volume panel.

Table 6-1   Meaning of the icons and colors in the Volume panel

Icon	Description
	Volume icon. This icon designates a volume created in the ESS. This icon is always displayed whether this volume is or is not in a copy relationship.
	Designated source volume. When you right-click on a volume to select it as a source, then the volume icon and its serial number are surrounded by a rectangle filled with blue color.
	Designated target volume. When you click on a volume to select it as a target, the volume icon and its serial number are surrounded by a rectangle filled with red color.

Icon	Description
	Designated source and target volume. When you click, followed by a right-click, on the same volume, then the volume icon and its serial number are surrounded by a rectangle filled with grey color.
	PPRC source volume. The volume is a source volume in an established PPRC relationship. And the volume is in duplex mode for this relationship.
	PPRC target volume. The volume is a target volume in an established PPRC relationship. The volume is in duplex mode for this relationship.
	PPRC source volume copy in progress. The volume is a source volume in an established PPRC relationship. The volume is in duplex pending mode for this relationship.
	PPRC target volume copy in progress. The volume is a target volume in an established PPRC relationship. The volume is in duplex pending mode for this relationship
	PPRC source suspended. The volume is a source volume in an established PPRC relationship. The volume is in suspended mode for this relationship.
	PPRC target suspended. The volume is a target volume in an established PPRC relationship. The volume is in suspended mode for this relationship.
	PPRC extended distance source. The volume is a source volume in an established PPRC-XD relationship.
	PPRC extended distance target. The volume is a target volume in an established PPRC-XD relationship.
	FlashCopy source. The volume is a source volume in an established FlashCopy relationship.
	FlashCopy target. The volume is a target volume in an established FlashCopy relationship.
	Change recording enabled. This icon displays on the source and target volume when you establish a FlashCopy with the Change recording enabled option. The pair can be used to perform incremental FlashCopy or reverse restore FlashCopy.
	Dataset FlashCopy (for S/390 only).

Icon	Description
	Volume Copy. The volume is the target of a FlashCopy relationship. This Icon remains even if the FlashCopy relationship is withdrawn. To remove this Icon manually, you need to withdraw the relationship with a "withdraw from target" command.
	Multiple FlashCopy relationship. The volume is a source volume for more than one FlashCopy relationship.
	The state of this volume cannot be determined.

6.5.2 Volume Information Panel button

You can get detailed information about a single volume by selecting the volume and clicking on the **Information Panel** button.

You can leave the Information Panel open as you work in the volume panel and select different volumes.

Figure 6-12 shows the information panel of a volume that is the source of two FlashCopy relationships.
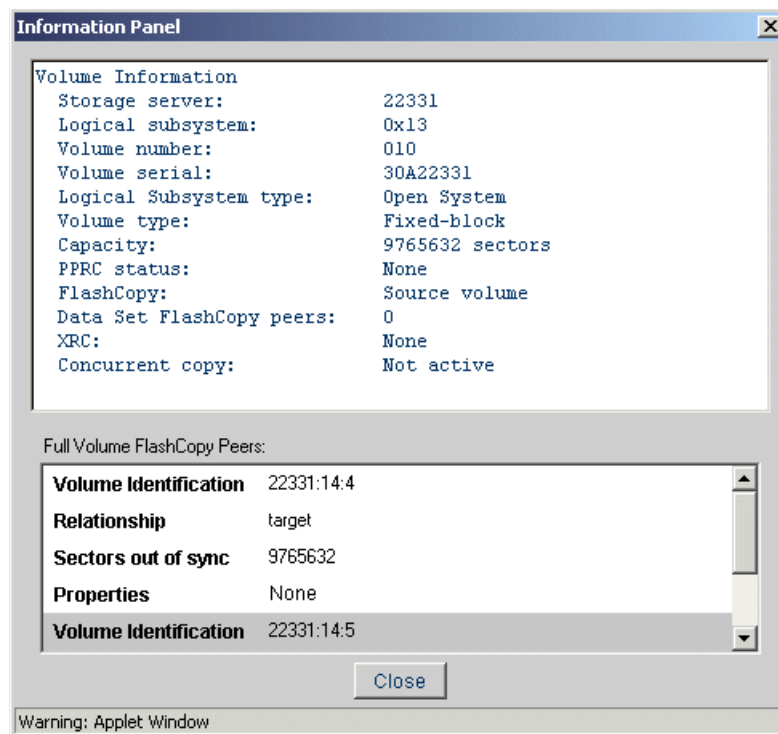
Figure 6-12   Volume Information Panel

The following information is displayed for a volume:

► The storage server (ESS) serial number.

► Logical subsystem (LSS) designations.

► The volume number.

- ► The volume serial number.

- ► The LSS type ("S/390" or "Open System").

- ► The volume type (fixed block for open system; 9337, 3390, and other type for S/390).

- ► The volume capacity in sectors (for fixed block) or in cylinders (CKD format).

- ► PPRC status: If the volume is in an active PPRC relationship, it indicates whether the volume is a primary (source) volume or a secondary (target) volume. The properties for the relationship, the volume status, and the number of out-of-synchronous sectors are displayed in the Full Volume PPRC peers (Figure 6-12 on page 183). If there are no active PPRC relationships active on a volume, a value of None is displayed.

- ► FlashCopy Status: If the volume is in any FlashCopy active relationship, it indicates whether the volume is a source or a target. Information about the corresponding volumes and out-of-synchronous sectors are displayed in the Full Volume FlashCopy Peer (Figure 6-12 on page 183). If there are no active FlashCopy relationships active on a volume, a value of None is displayed.

- ► Any active extended remote copy (XRC) operations (this function is available for S/390 only).

- ► Whether the Concurrent Copy function is active (this function is available for S/390 only).

6.5.3 Finding volumes

The **Find** button in the volumes panel enables you to search for a specific volume. The volume must be specified by its eight digit ID.
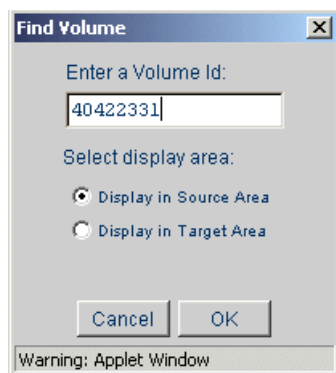
Figure 6-13   Find volume

If the volume ID is valid, the volume is displayed in the area you selected by clicking on **Display in source area** or **Display in target area**, as in Figure 6-13.

The volume is displayed on the top of the area and is automatically selected as a "source" volume (as if you selected the volume and performed a right-click on it).

6.5.4 Filtering volumes

You can filter the volumes displayed in the volume panel. To do this, click on the filter button, and select the filter volume option, as in Figure 6-14 on page 185.
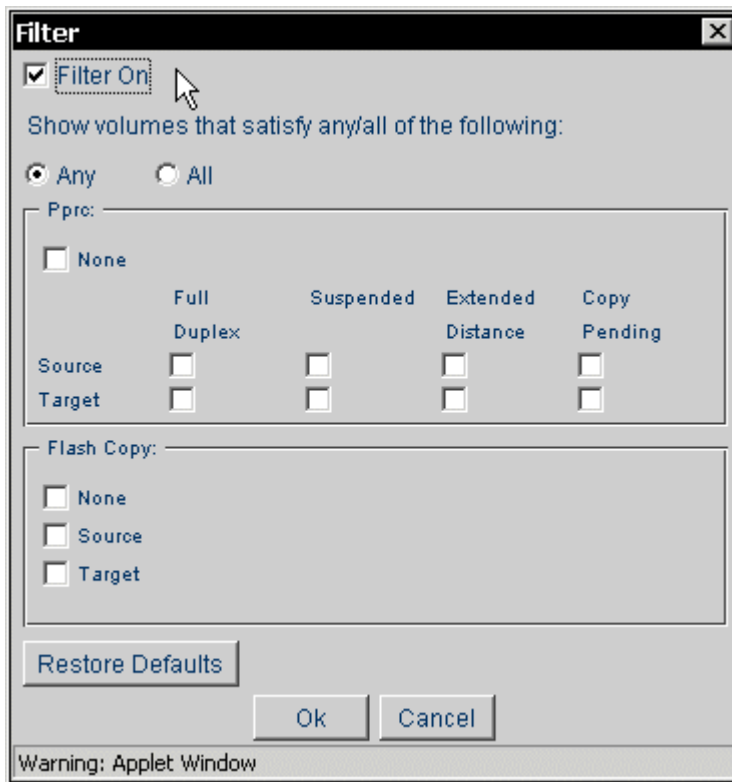
Figure 6-14   Filter volumes panel

6.5.5 Multiple Selection Mode

When you enter in the **Multiple Selection Mode**, you can select multiple volume pairs when you create tasks for both FlashCopy and PPRC. This mode is very useful for saving time when creating the tasks.

You might use this mode when selecting the volume for FlashCopy and PPRC tasks (for further information, refer to 6.11.1, "Establishing a FlashCopy pair" on page 204, 6.12.3, "Establishing a synchronous PPRC pair" on page 239, and 6.12.4, "Establish a PPRC Extended Distance copy pair" on page 244). Rather than defining one copy pair at a time, you can define all the copy pairs at one time by selecting multiple source and their targets in the same task definition (Figure 6-15 on page 186).

To enter the Multiple Selection Mode:

1. Click on the **Multiple Selection Mode** button at the bottom of the panel. This allows multiple pairs to be defined.

2. Select the desired source volume and right-click the desired target volume.

3. Repeat the selection of source and target volumes until you have defined all volume pairs.

4. Right-click one of the target volumes again and the Task Wizard opens. You can decide whether to run one of the following operations:

   – Establish synchronous PPRC copy pairs

   – Establish Extended Distance PPRC copy pairs

   – Suspend PPRC copy pairs

   – Terminate FlashCopy pairs

– Establish FlashCopy pairs

– Withdraw FlashCopy pairs

To exit the Multiple Selection Mode, click the **Exit Multiple Selection Mode** button at the bottom of the Volume Panel.

Figure 6-15   Multiple Selection Mode

Restriction: You can only perform the multiple selections on volumes within the same LSS (all sources must be in the same LSS and all targets must be in the same LSS).

Attention: If you click Exit Multiple Selection Mode before you perform any operations on them, you lose all your previous selections.

6.6 Logical subsystems panel

The **Logical Subsystems** panel displays the Enterprise Storage Servers and the logical subsystems (LSSs) within the storage network. The storage network includes all Enterprise Storage Servers that are configured to use the same Copy Services Server. Each of the logical subsystems is specified by the serial number of the ESS it belongs to and its 2-digit LSS number within the ESS.

With the **Logical Subsystems** panel, you will be able to:

► Establish PPRC copy pairs for LSSs.

► Convert PPRC Extended Distance copy pairs to synchronous pairs
  for all the volumes in a source LSS and target LSS.

► Filter a list of LSSs.

► Find an LSS.

► Freeze a PPRC consistency group.

► Modify the PPRC consistency group time-out value.

► Remove orphaned paths.

► Resynchronize PPRC copy pairs for LSSs.

► Run a consistency-group-created operation.

► Suspend PPRC copy pairs for LSSs.

► Terminate PPRC copy pairs for LSSs.

► View information about an LSS.

In Figure 6-16, you can see the **Logical Subsystems** panel for a selected ESS. The color indicates the state of the LSS, whether it contains volumes that are currently in a copy relationship (source, target, or mixed), or that they are not part of a copy pair at all.
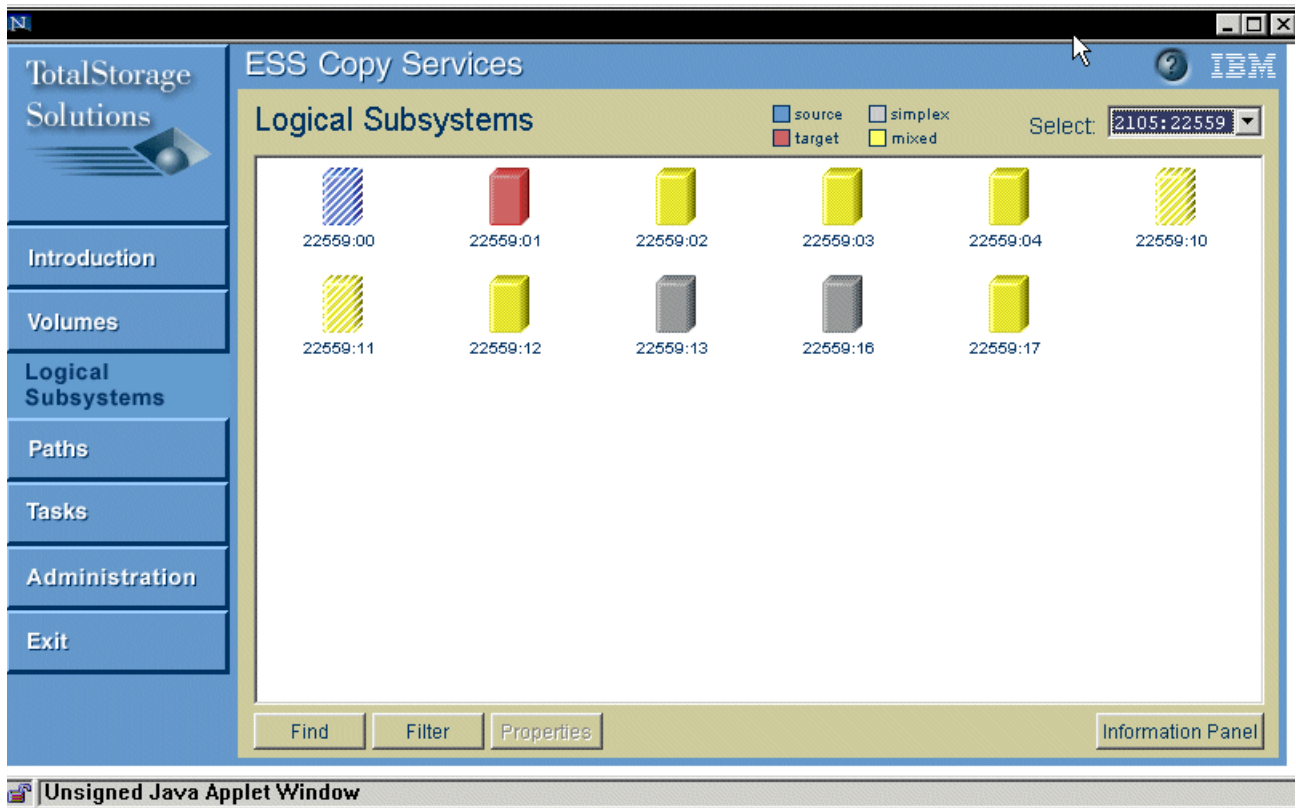
Figure 6-16   Logical Subsystem panel

Table 6-2 explains the meanings of the icons that appear on the Logical Subsystems panel. The Volumes panel displays similar icons for individual volumes (see Table 6-1 on page 181).

Table 6-2   Meaning of icons on Logical Subsystems panel

LSS Icon	Appearance	Meaning
	Gray solid ("Simplex state")	All the volumes in the LSS are in simplex state, meaning that they are not in a copy relationship.

LSS Icon	Appearance	Meaning
	Blue solid (PPRC sync. source)	All the volumes in the LSS are in PPRC pairs (synchronous or Extended Distance) in a non-suspended state.
	Red solid (PPRC sync. target)	
	Blue and white stripes (PPRC suspended source)	All the volumes in the LSS are in PPRC pairs (synchronous or Extended Distance) in a suspended state.
	Red and white stripes (PPRC suspended target)	
	Yellow solid	Not all volumes in the LSS are in the same type of Copy Services relationship (this apply for PPRC and FlashCopy relationships), and none of the volumes are in a suspended state.
	Yellow and white stripes (mixed types and states)	The volumes in the LSS are either the same type, but some are in a non-suspended state, or not all of the volumes are the same type, but some are in a suspended state.

6.6.1 LSS Information Panel button

You can get detailed information about a single LSS by selecting the LSS and clicking on the **Information Panel** button.

In our example in Figure 6-17 on page 189, we selected LSS10 and opened the information panel. LSS 10 was colored with yellow solid. In the information panel, we read that LSS10 contains six volumes. All six volumes are currently source volumes of FlashCopy relationships.
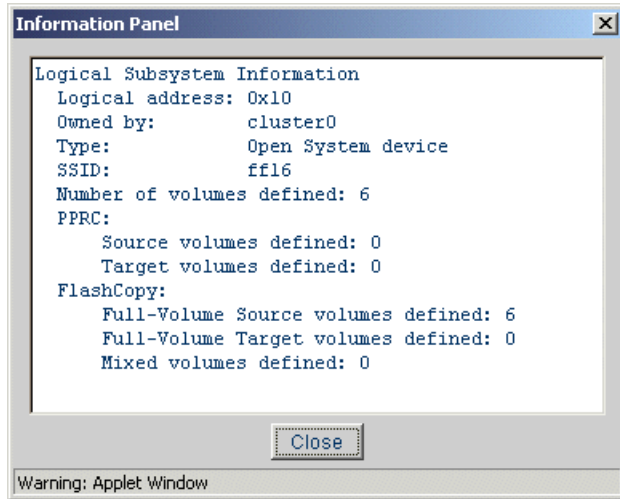
Figure 6-17   LSS Information Panel

6.6.2 Finding LSSs

You can search for a specific LSS by clicking on the **Find** button. The LSS must be referred to using the following syntax: ESS serial number:LSS number.

In our example in Figure 6-18, we want to find LSS 10 in the ESS with the serial number 22331.
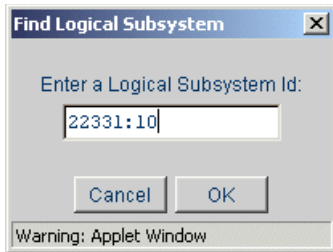
Figure 6-18   Find Storage Server LSS

6.6.3 Filtering LSSs

You can have many filtering options from the Logical Subsystems panel. Click on the **Filter** button, and select the filter volume option (see Figure 6-19 on page 190).
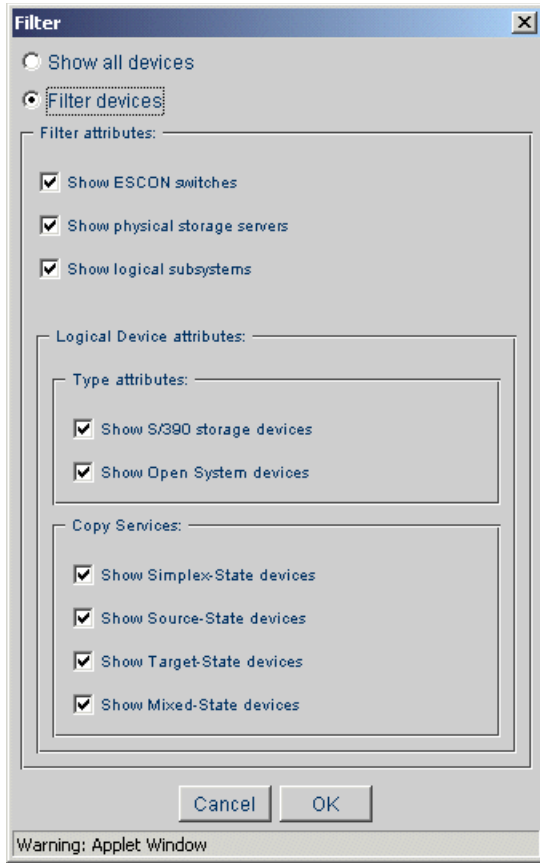
Figure 6-19   Filtering LSS

6.6.4 LSS properties

To view or modify the properties of an LSS, select the LSS and click on the **Properties** button, as in Figure 6-20.
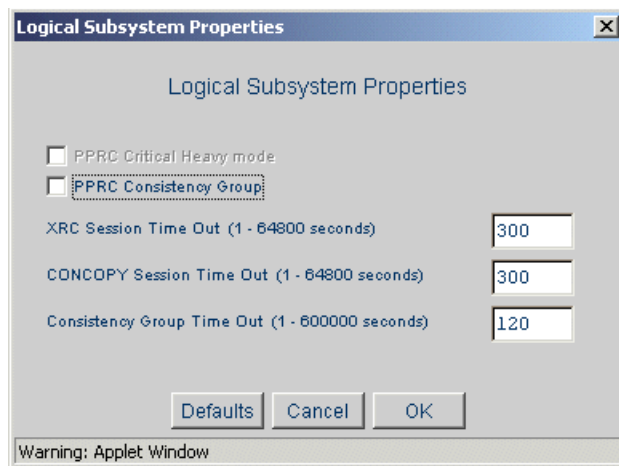
Figure 6-20   LSS properties

PPRC Critical Heavy mode

This option is not available for open-system LSS.

PPRC consistency group

To maintain consistency of data across volumes at the recovery site, volumes in a PPRC relationship can be collected into a *consistency group*. If you mark this check box, then all the volumes in that LSS will be considered as part of one Consistency group. Creating a consistency group of volumes provides the ability to temporarily queue subsequent write operations to all PPRC consistency group volumes on an single LSS pairing when an error occurs on one of the volumes in the group (primary or secondary), or when a total link failure is detected between the primary and secondary LSS pair. The other way to establish a consistency group is to select the PPRC Consistency group option when establishing a path between two LSSs. For further information about PPRC Consistency groups, refer to 4.5, "PPRC data consistency" on page 72.

XRC session timeout

This parameter does not apply to open systems.

CONCOPY session timeout

This parameter does not apply to open systems.

Consistency Group timeout

This parameter, used by FlashCopy V2 and PPRC V2, indicates the amount of time that an I/O is withheld from updating a primary volume of a consistency group in case of an error event. This timeout enables automation software to detect that an error has occurred and to issue commands to *freeze* all other members of the consistency group. When the consistency group is created, this parameter is set to a default of two minutes. For further information about PPRC consistency groups, see 4.5, "PPRC data consistency" on page 72. For further information about FlashCopy consistency groups, see 3.3.4, "FlashCopy Consistency Groups" on page 37.

6.7 The Paths panel

A path is used to send data between the primary and secondary volumes of PPRC pairs. The physical path consists of the ESCON connection between two Enterprise Storage Servers, while a logical path describes the connection of the PPRC source and targets.

Be aware that before you can establish PPRC copy pair relationships, you must use the Paths panel to establish logical paths between logical subsystems between the source and the target LSS. After you establish the paths, the Paths panel displays the current status of the paths for the selected LSS. There can be multiple logical paths established over a single physical path. For further information, refer to 4.7.3, "PPRC paths definition" on page 87.

From the Paths panel (see Figure 6-21 on page 192), you will be able to:

► Establish paths.
► Add paths.
► Remove a group of established paths.
► Remove one or more paths from a group of established paths.
► Create a PPRC consistency group.
► View information about paths.

To view configured paths over Fibre Channel links (ESS serial number 23953, in our example) on the Paths panel, follow the steps listed:

1. Start by clicking on **View Path Status**, as shown in Figure 6-21 on page 192.
2. From the drop-down box, select the source ESS.

3. Click on the source LSS.
4. From the drop-down box, select the FCP.

For paths successfully configured, three blue asterisks will show under the link.
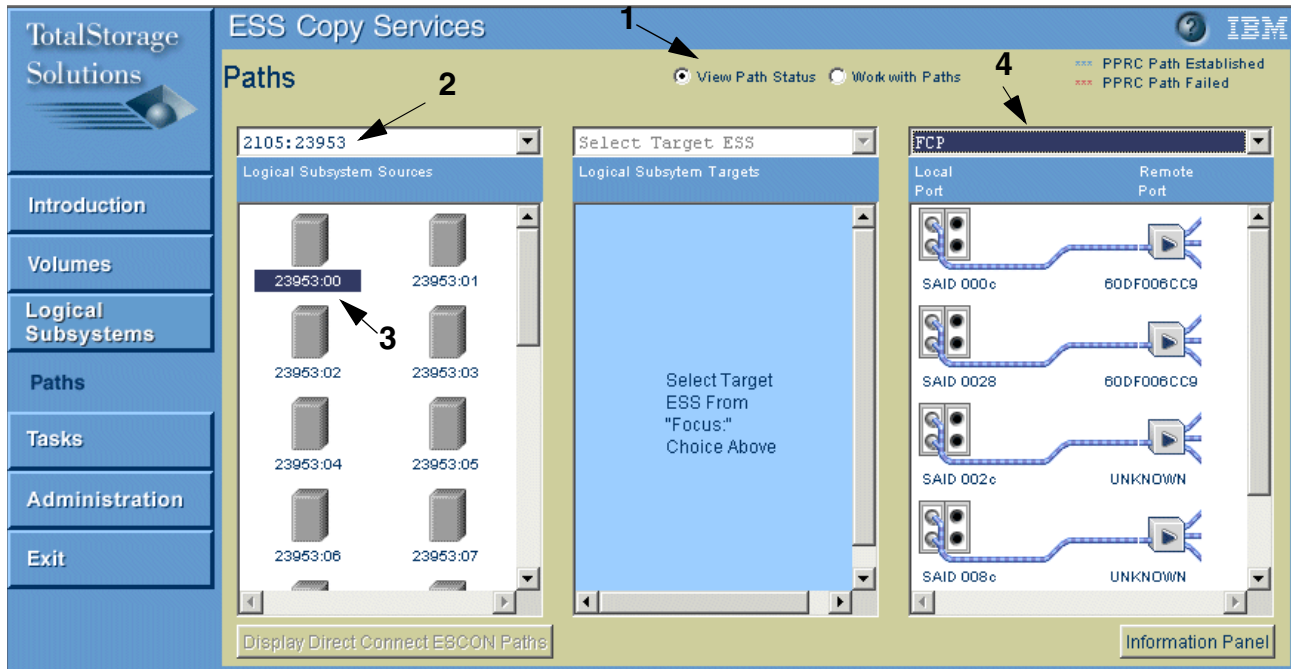
Figure 6-21   View FCP paths status

To view configured paths over the ESCON links, on the Paths panel, repeat steps one through three, but for step 4, from the drop down box, select the adapter type ESCON, as shown in Figure 6-22.
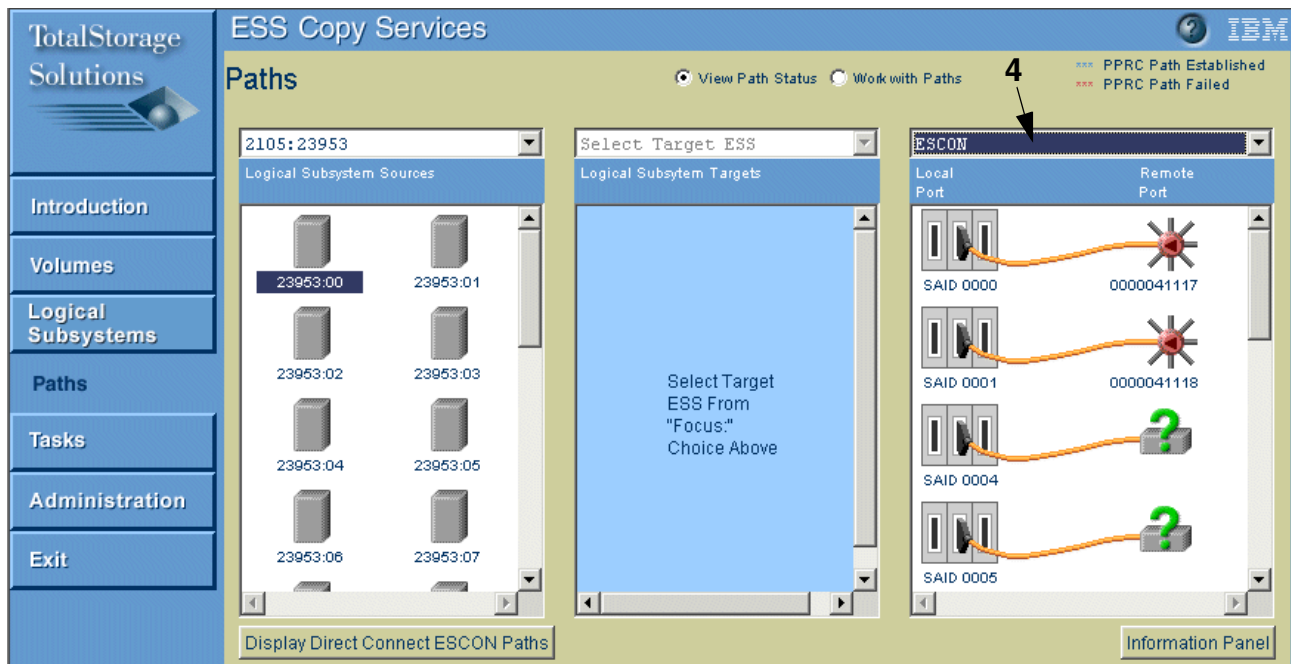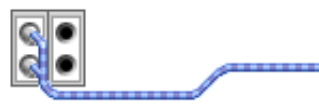
Figure 6-22   View ESCON paths status

6.7.1 Icons and symbols in the Paths panel

Table 6-3 gives you an overview of the connection symbols used in the "Local Port / Remote port" column on the right side of the panel.

Table 6-3   Meaning of the Paths panel icons

Connection icon	Description
	ESCON Host adapter port, with the system adapter identification number (SAID) provided under it
	ESCON director
	Host server
	Storage server
	Unknown device
 SAID 0084	Local FCP host adapter port, with the system adapter identification number (SAID) provided under it
	Remote FCP host adapter port, with serial number provided under it.

If there are logical paths defined on an ESCON or FCP adapter, you will find three blue asterisks right below the adapter in the Local/Remote Ports column on the Paths panel. Three red asterisks below the path-connection symbol mean that the last attempt to establish the path failed.

Figure 6-23 on page 194 shows an ESCON adapter SAID without, then with, the path successfully defined, and finally with a failure to establish the path.
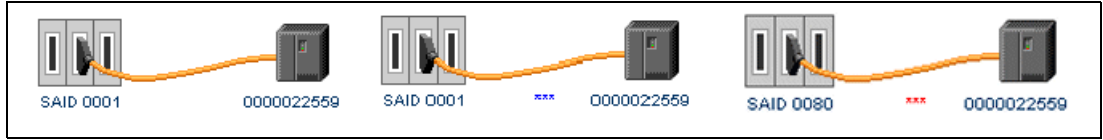
Figure 6-23   ESCON link: No path defined, path defined, and path definition failed

6.7.2 ESCON System Adapter ID (SAID)

The ESCON adapters are specified by their System Adapter ID (SAID). Figure 6-24 shows the SAID of all ESS ESCON adapters.

SAID numbers of the ESS ESCON adapters

Host adapter bays viewed from the frontside of the Enterprise Storage Server
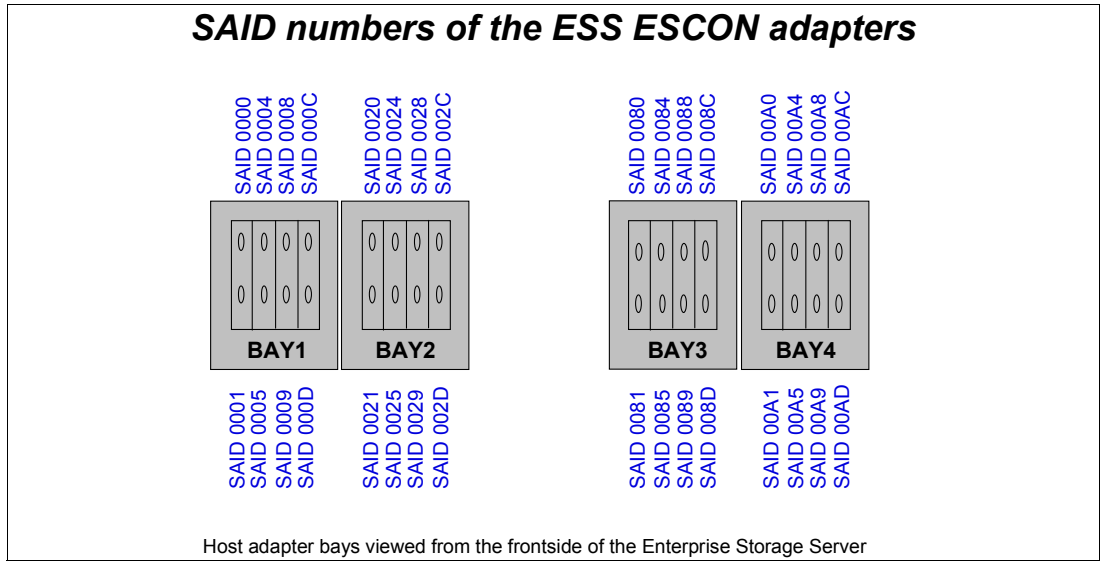
Figure 6-24   SAID numbers of the ESCON adapters

6.7.3 FCP System Adapter ID (SAID)

The FCP adapters are also specified by their System Adapter ID (SAID). Figure 6-25 shows the SAID of all the Fibre Channel ports.
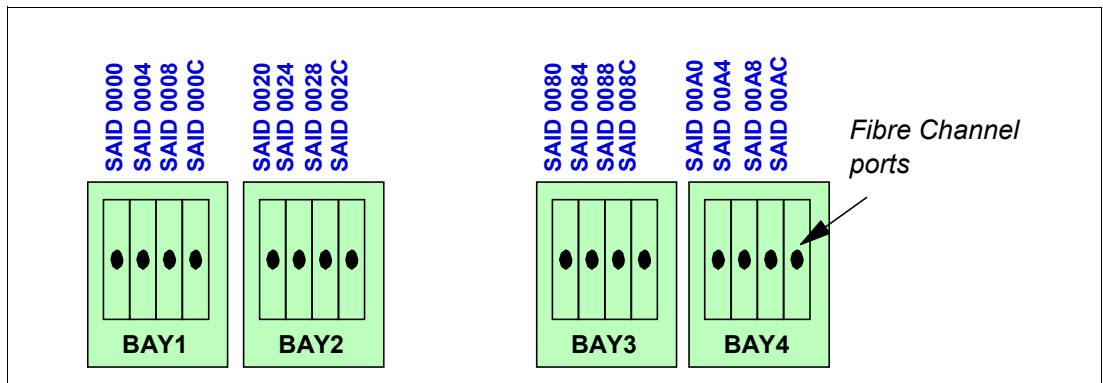
Figure 6-25   SAID numbers of the FCP adapters

The System Adapter IDs (SAID) are further discussed in Appendix D, "System Adapter ID (SAID)" on page 493.

6.7.4 Display Connection Paths

If you have two ESSs that are directly connected (that is, no switch between them), you can click **Display Direct Connect Paths** to display those paths.

> **Note:** You must perform a refresh function for the following conditions:
> ► If this is the very first time anyone is accessing the ESS Copy Services WUI.
> ► After an update of the ESS microcode.
> ► After a concurrent service action.
> ► If the cabling for the ESCON links is modified.

6.7.5 Path Information Panel button

Once an ESCON or FCP adapter is selected, you can get more information about the paths by clicking the **Information Panel** button at the bottom of the Paths panel as in Figure 6-26.

```
Information Panel                                    [X]

Source SAID: 0081
  Serial number:        18767

Remote SAID: 0080
  Serial number:        22331
  Manufacturer:         IBM
  Type:                 2105
  Model number:         800
  Connection type:      Storage Server
  Total PPRC Paths:     2
  PPRC Source Paths:
    0: destination:               22331
       SSID:                      0xff19
       status of path:            Established
       switch port destination:   0x0
       logical subsystem source:  0x16
       logical subsystem target:  0x13

                      Close

Warning: Applet Window
```

Figure 6-26   Path Information Panel

The example shown in Figure 6-28 on page 197 shows a path defined between source LSS 16 on ESS 18767 and target LSS 13 on ESS 22331 using the ESCON adapters SAID0081 on the source ESS and remote SAID0080 on the target ESS.

6.8 Tasks panel

With the ESS Copy Services, you have the possibility to save the setup of any data copy action within a Task. This could be any kind of FlashCopy, PPRC, or path operation.

With the Tasks menu, you will be able to:

► Group tasks.
► Modify a task.
► Remove a task group.
► Remove a saved task.
► Run a saved task.

- ► Ungroup tasks.
- ► View error information about a failed task.
- ► View information about a saved task.

Figure 6-27 shows the Tasks panel of the ESS Copy Services. For each task, the name, a description, and the last status of the execution is displayed.

You will be prompted to choose a task name when you create a task. Use the following guidelines:

- ► The name can only contain alphanumeric characters.
- ► The name cannot contain blanks.
- ► The name can contain the underscore (_) or hyphen (-) characters.
- ► The name cannot exceed 16 characters.

Note: In Appendix E, "Sample naming convention" on page 499, we present the task naming convention used in this chapter. You can refer to this appendix to find some hints and tips about naming conventions.
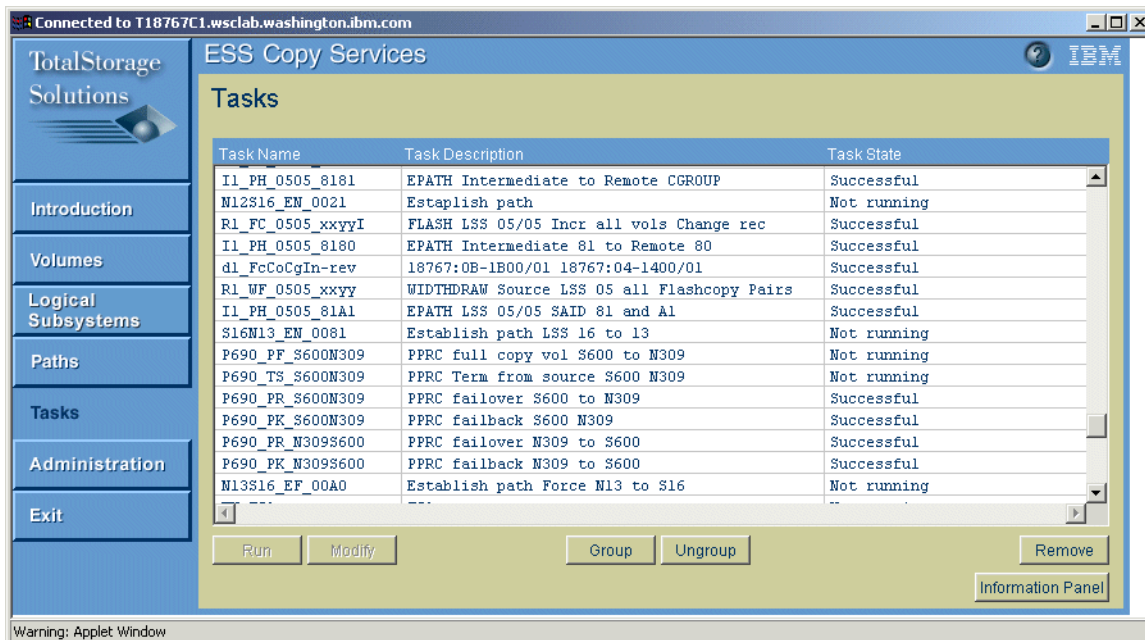
Figure 6-27   Tasks panel

6.8.1 Display task information

You can see information about a task by selecting the task and clicking on **Information Panel**.

In our example in Figure 6-28 on page 197, we show a task that will perform a FlashCopy with the **No copy** and **Freeze** options selected. The FlashCopy will be performed on multiple volumes (there are four source volumes and four target volumes). The task was created using the Multiple selection mode.
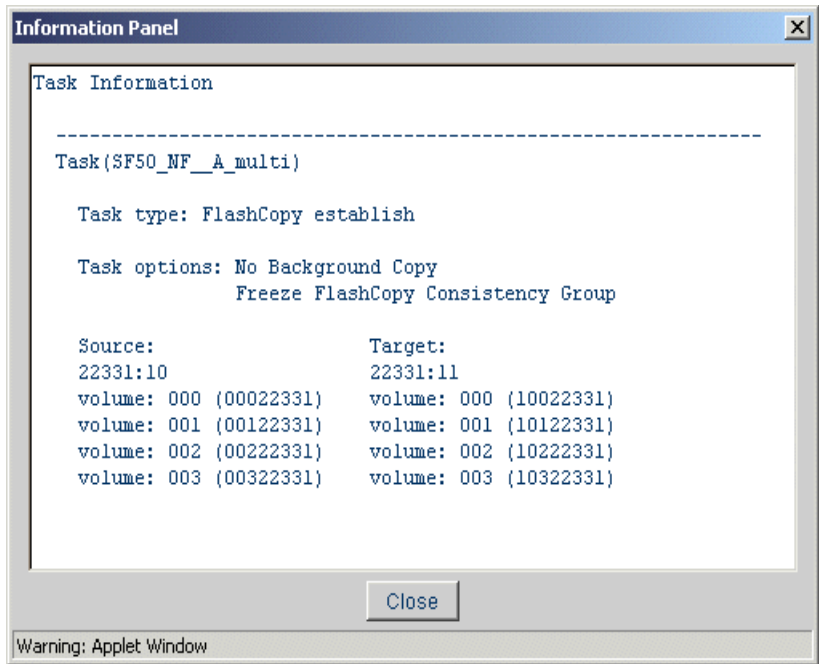
Figure 6-28   Task Information panel

6.8.2 Grouping and ungrouping tasks

To create a group task, click the single tasks you want to group together while holding the Shift key or the Control key (mark single tasks). Once you are finished, click the **Group** button and specify the group name. It is not possible to include a group into another task group. All tasks within a task group will be processed in parallel.
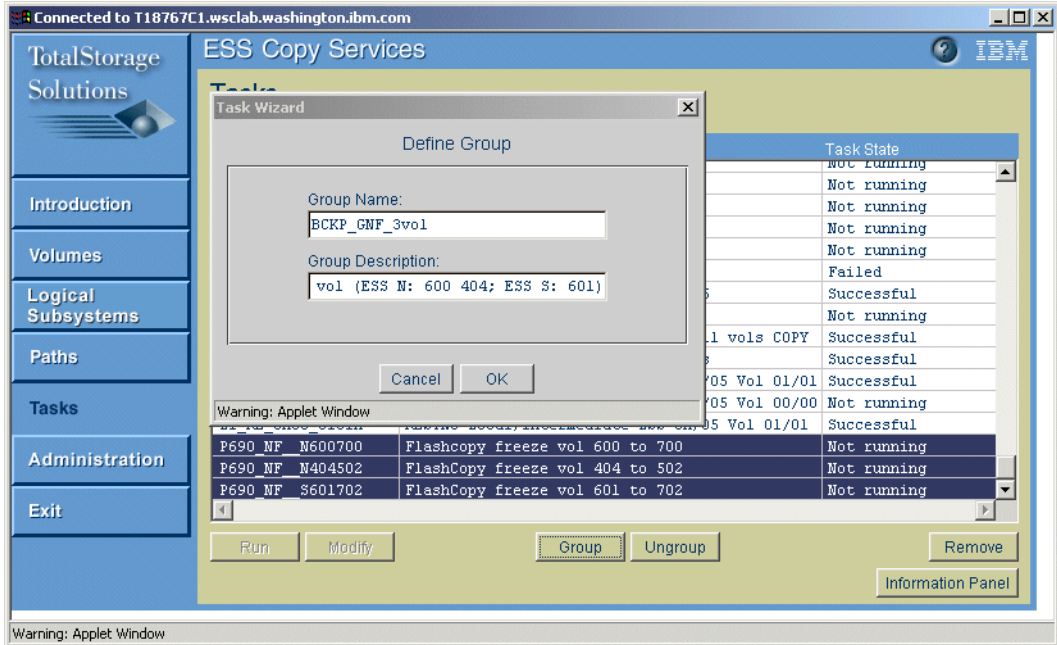
Figure 6-29   Tasks panel: Grouping tasks

An example for the usage of a task group would be multiple FlashCopy pairs from different logical subsystems or different ESSs that need to be issued all at the same time with the **Freeze** option in order to do a backup using the target volumes consistency group.

You can get detailed information about the setup of the task. Select the task and click the **Information Panel** button at the lower right.

In our example in Figure 6-30, we show the information panel for a grouped task. This group contains three single tasks. The three tasks are FlashCopy tasks with **No copy** and **Freeze** options. They will be performed in two different ESSs.

Grouping tasks will cause the individual tasks to collapse into the overall group task.

```
Information Panel                                                    ×

Task Information

   --------------------------------------------------------------
  Group Task(BCKP_GNF_3vol)

    Task(P690_NF__N600700)

      Task type: FlashCopy establish

      Task options: No Background Copy
                    Freeze FlashCopy Consistency Group

      Source:                 Target:
      22331:16                22331:17
      volume: 000 (60022331)  volume: 000 (70022331)

    Task(P690_NF__N404502)

      Task type: FlashCopy establish

      Task options: No Background Copy
                    Freeze FlashCopy Consistency Group

      Source:                 Target:
      22331:14                22331:15
      volume: 004 (40422331)  volume: 002 (50222331)

    Task(P690_NF__S601702)

      Task type: FlashCopy establish

      Task options: No Background Copy
                    Freeze FlashCopy Consistency Group

      Source:                 Target:
      18767:16                18767:17
      volume: 001 (60118767)  volume: 002 (70218767)

                        Close

Warning: Applet Window
```

Figure 6-30   Information panel: Task group

To ungroup a task, select the task and click the **Ungroup** button at the bottom of the Tasks panel. It is necessary to ungroup the overall group task in order to gain access and to execute the individual tasks.

6.8.3 Removing a task

To remove a task, select the task and click the **Remove** button at the bottom of the Tasks panel.

6.8.4 Running a task and viewing an error about a failed task

To run a saved task, select the task, and click the **Run** button at the bottom of the Tasks panel. The task will be executed immediately.

If the task fails, a failure report is available through the Information Panel (Figure 6-31). The failure report contains the error messages and the sense data. The messages are documented in the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7346, available on the ESS Web site.

You can also view the logs in the administration panel.

```
Information Panel                                        [x]

Task Information

  --------------------------------------------------------------
  Task(N13S16_EN_0000)

    Task type: Establish copy paths

    Task options: No Force Establish

    Source:      SAID    Through                    Target:
    22331:13     0x0000  DIRECT-CONNECTION          18767:16

    Failure Report:
     Failure:(ANTP0117I)(22331:13)(18767:16) sense=80000000-00ffff0f

                            [ Close ]
Warning: Applet Window
```
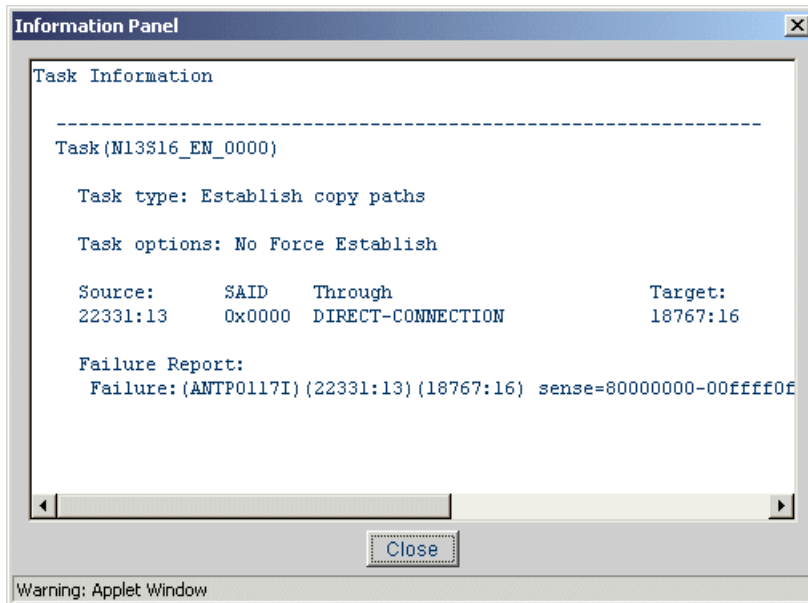
Figure 6-31   Information panel: Failing task report

6.8.5 Modifying a task

Use the Tasks panel to modify a previously saved task. You can change the task goals, task options, task name, and description.

You will be prompted if you want to run, to replace the selected task, or create a new task.

You can change a task from one that establishes a PPRC copy pair to one that suspends, re-synchronizes, or terminates a PPRC copy pair. For FlashCopy, you can change a task that establishes a FlashCopy pair to one that withdraws a FlashCopy pair, or you can change the options of one task.

> **Note:** Grouped tasks cannot be modified. The tasks must be ungrouped and then the tasks can be modified.

6.9 Administration panel

Use the Administration panel (see Figure 6-32) to manage the server logs, reports, and ESS Copy Services Command Line Interface (CLI) user IDs (user IDs) and passwords. You can also use this panel to refresh the volume and LSS information for an ESS.
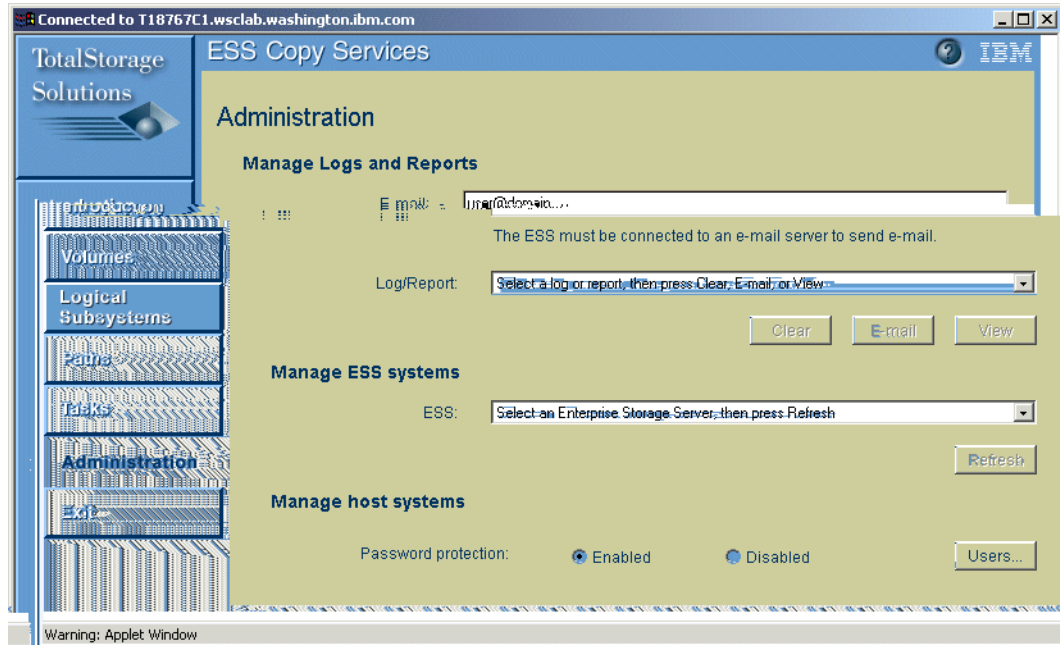
Figure 6-32   Copy Services Administration panel

You can perform the following actions from the Administration panel:

► Clear the server logs.

► Manage the CLI user ID and password for an open-systems host:
   – Define the user ID and password.
   – Remove the user ID and password.

► Refresh the path and volume information for an ESS.

► Send the ESS Network Configuration report to an e-mail address.

► Send the ESS Resource Configuration report to an e-mail address.

► Send the server logs to an e-mail address.

► Specify the e-mail address that receives server logs, ESS Network Configuration reports, and ESS Resource Configuration reports.

► View the ESS Network Configuration report.

► View the ESS Resource Configuration report.

► View the server logs.

6.9.1 Working with the copy services logs and reports

► The Copy Services server maintains the following logs and reports:

   – **Copy Services status log (copyservices.log):** This contains messages that are associated with user actions issued through ESS Copy Services. This log includes messages associated with such actions as saving tasks and executing tasks. The log

also includes information about the success or failure of the tasks. You can click the error message to get a description of the problem and its possible solutions.

> **Attention:** The status log is useful for problem determination. Therefore, you should only clear it when no problems exist or when you have completed data collection for all failures.

– **Copy Services Timing log (rsCStiming.log):** This contains information about the time it took for data-copy functions (PPRC or FlashCopy) to complete on a specific volume.

– **ESS Network Configuration report:** The ESS Network Configuration report contains important information about your network-defined ESS resources:

  • IP addresses for your primary and Backup Copy Services servers
  • IP addresses and host names for the ESS clusters
  • IP addresses for the ESS clients

– **ESS Resource Configuration report:** The ESS Resource Configuration report contains the following important information about your ESS resources:

  • The 2105 connection information (SAIDs, port types, and values)

  • PPRC path information (SAIDs, LSSs, remote SAIDs, ports, and status)

  • S/390 or zSeries volumes (LSSs, labels, Concurrent Copy status, FlashCopy status, PPRC status, or XRC status)

  • Open-systems volumes information (LSSs, volume serial numbers, user defined label, FlashCopy status, or PPRC Status)

An extract of the ESS Resource Configuration report is presented in Figure 6-33.
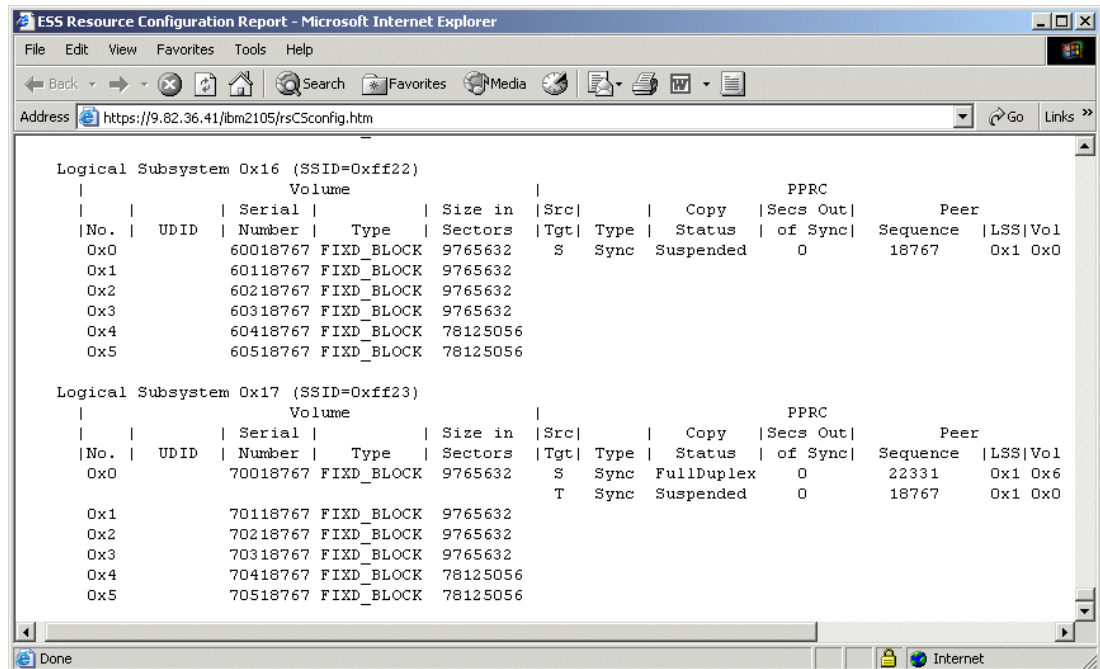
Figure 6-33   Copy Services Administration panel

To clear a log, select the log in the **Log/Report** list and click the **Clear** button. Notice that you cannot clear a report. You can view the log or report by selecting the appropriate log/report and click the **View** button.

To send a report or a log to an e-mail address, select the appropriate log/report in the **Log/Report** list and click the **E-Mail** button. Notice that you have to specify the e-mail address in the **E-Mail** field.

6.9.2 Refreshing path and volume information on an ESS

Use the refresh function of ESS Copy Services to get the most up-to-date information about paths and volumes for an ESS. To perform this action, select the ESS for which you want this information in the ESS list and click the **Refresh** button.

A status message opens on the top of the Administration panel noting that:

```
A refresh request was sent to <the selected ESS>. It may take a few minutes before the
server can refresh the data.
```

6.9.3 Managing the CLI user ID and password for an open systems host

Use the Administration panel to manage the user IDs (user IDs) and passwords for the Copy Services Command Line Interface (CLI). Open-systems hosts that issue CLI commands to a Copy Services server must provide a valid user ID and password as part of the command string. If the **Password Protection** option is enabled and the hosts do not provide the user ID and password, the command fails.

> **Attention:** in a dual-active copy services server configuration, when you define the user ID and password on one active server A, server A does not send them to the other Copy Services server (server B) as it does for Copy Services tasks. Therefore, you must define the user ID and password at the Copy Services server B in a separate step.

To enable the password protection for access to an ESS through the CLI click **Enabled** in the Password Protection field. To add users, click the **Users** button and the Authorized Users dialog box will appear. Click the **Add** button and enter the user name and password. If you have entered the user name and the password, click the **Add** button and the new user will be authorized. To remove a user, select the user in the Authorized Users dialog box and click the **Remove** button. After you have performed all actions, click **Done** and the Authorized Users dialog box will be closed.

6.10 Exiting ESS Copy services

There are several legitimate ways to quit the ESS Copy Services Web User Interface:

- ► You can click the **X** at the top right corner of the ESS Copy Services panel.
- ► You can click **Exit** in the navigation frame of any of the ESS Copy Services panels.
- ► You can shut down the whole browser from another browser window.

In addition, if you simply want to return to the ESS Launch page without closing ESS Copy Services, such as to launch ESS Specialist, you can keep the ESS Copy Services browser window open and use standard navigation methods such as Alt-Tab to switch to the ESS Launch page.

The difference between the first two options above is nothing more than a warning message. If you click **Exit** in the navigation frame of any of the ESS Copy Services panels, the warning

message opens on top of the browser window and tells you that the tasks that you have submitted will continue running after you close the window. Click **OK**. The ESS Copy Services browser window closes and the ESS Launch page comes into focus. In either case, the next time you click **Copy Services** on the ESS Launch panel, the ESS Copy Services browser window opens without requiring you to log in. The browser loads a cached version of the applet.

> **Note:** Loading the cached applet is a problem in two situations:
> - ► If you load a new version of the ESS microcode during that time (while you have not yet quit your browser), the cached applet might not conform to the new code on the ESS.
> - ► If unauthorized users have access to your workstation when you leave it, they can access the ESS through your cached user ID.

6.11 Performing FlashCopy operations with the WUI

In this section, we explain how to set up FlashCopy relationships between volumes in an ESS using the Web User Interface (WUI).

Be aware of these requirements of the FlashCopy functionality:

- ► The source and target volume must be in the same ESS. Within the same ESS, they can be in a different LSS.

> **Note:** This is available only if the FlashCopy feature V2 is installed. If the FlashCopy V1 is installed, then the source and target volumes must be in the same LSS.

- ► The target volume must be the same size as the source volume or larger.
- ► A target volume can be in only one FlashCopy relationship at a time.
- ► A Source volume can be in multiple relationship at a time. Among the multiple FlashCopy relationships using the same volume as source volume, only one can be an "incremental" FlashCopy relationship.

> **Note:** This is available only if the FlashCopy feature V2 is installed. If the FlashCopy V1 is installed, then the source can only be in one FlashCopy relationship at a time.

> **Note:** Prior to establishing a FlashCopy, we recommend that you do the following:
> 1. Unmount the target volume from all host systems, since the FlashCopy process is a destructive operation to the target and will overwrite the data on the target volume.
> 2. Additionally, stop all active reading/writing operations to the source.
> 3. Also, quiesce your application and flush the data to the source volume.

There are two different ways of establishing a FlashCopy pair:

- ► From the Volumes panel
- ► From the Tasks panel (once a task for a FlashCopy is created and saved)

6.11.1 Establishing a FlashCopy pair

Use the volume panel to establish a FlashCopy pair. In the source area, select the LSS within which you will select the source volume. To select the LSS, you will use the source drop-down list. When the source LSS is selected, select the LSS within which you will select the target volume in the target area.

You always need to have two components to establish a FlashCopy pair: a source and a target. Click on the source volume to select it, then right-click on the target volume to select the volume. If you have selected the wrong source or target volume, just click on the correct source volume again.

The source and target logical subsystem is specified the following way: Device type (4 digits):ESS Serial number (5 digits):LSS number (2 digits).

In our example shown in Figure 6-34, we have selected 2105:22331:10 as the source LSS and 2105:22331:11 as the target LSS. The volume 000 in LSS 10 is selected as a source volume and volume 000 in LSS 11 is selected as target volume.
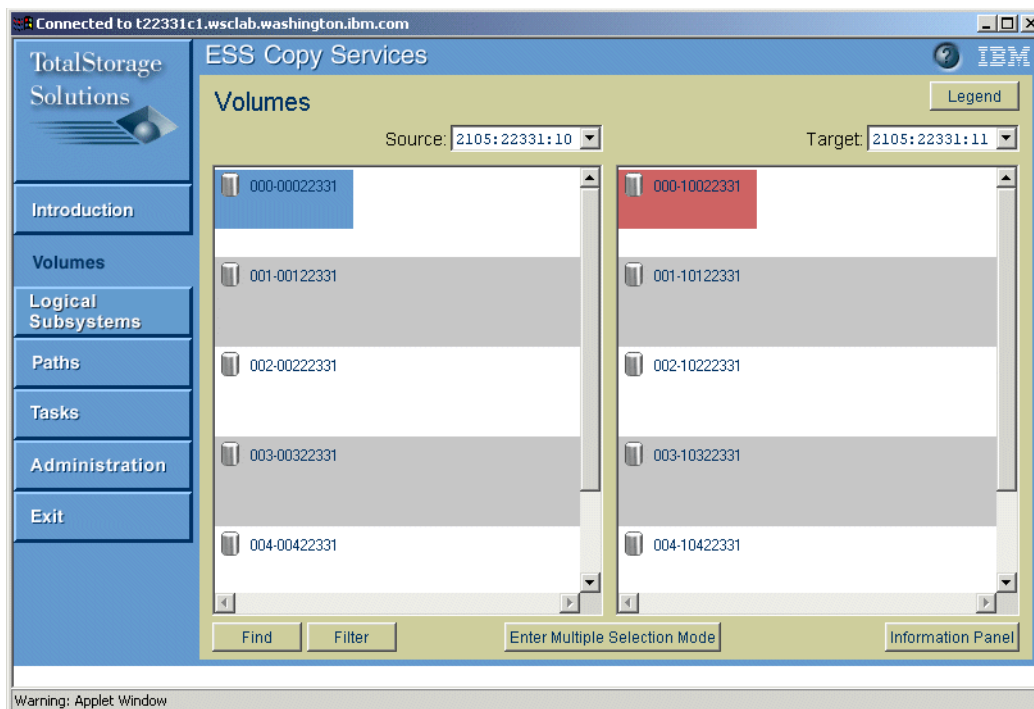
Figure 6-34   Volumes panel with source and target volumes selected

Another right-click on the target volume will bring up the Task Wizard panel. In the Select Task panel, select the **Establish FlashCopy** option. Click on **Next**.
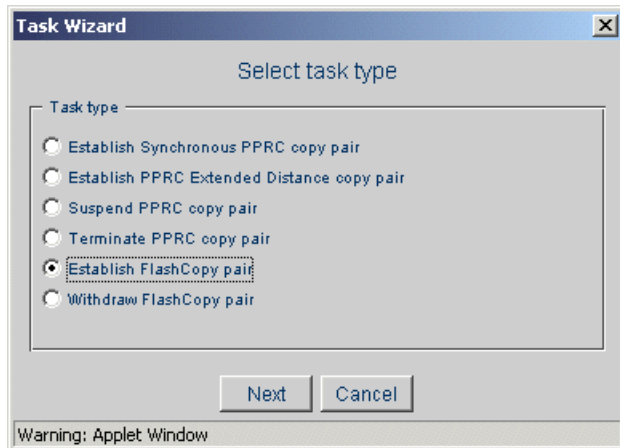
Figure 6-35 Task Wizard panel

Select the options for the FlashCopy task (see Figure 6-36 and the paragraphs below for details about the options).
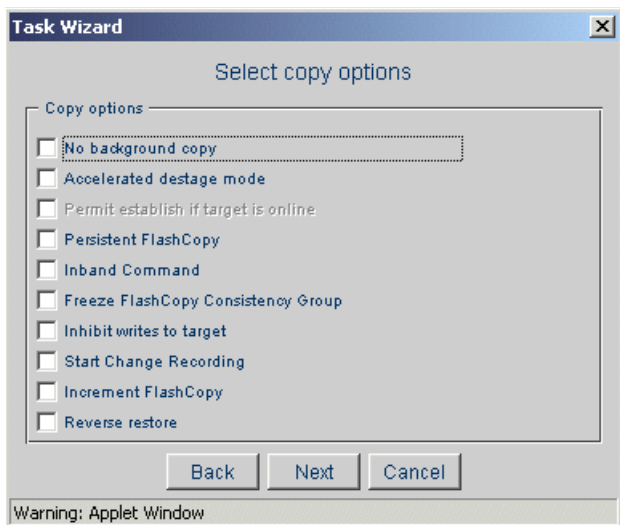
Figure 6-36 Define task panel

Selecting the FlashCopy options

3.4, "FlashCopy tasks and options" on page 39 gives detailed descriptions of the different options for FlashCopy. Refer to this section to fully understand the different options. We present only a summary here of the different options.

No background copy option (nocopy)

If this option is checked:

► Only the tracks modified on the source volume are copied to the target volume.

► This relationship between source and target volume remains forever and has to be broken manually.

► By default, this option is not selected and all data is copied from the source to the target volume of a FlashCopy pair (background copy). Once all data is copied, the relationship ends automatically (unless the Persistent FlashCopy option is selected).

► If two volumes are in a FlashCopy "nocopy" relationship, you can perform a FlashCopy Start Background Copy to create a physical background copy.

Accelerated destage mode

Select this option to cause a FlashCopy source volume track, which is being modified and is in a FlashCopy relationship, to be destaged from cache sooner than it would be if normal cache algorithms were applied. This feature minimizes (for the source volume) the number of modified tracks that are resident in cache.

Note that:

► With FlashCopy Version 1, this option might (under specific circumstances) speed up the completion of the background copy process.

► With FlashCopy Version 2, this feature has no effect, since the cache algorithms have changed.

Permit establish if target is online

Select this option if you want to establish a FlashCopy relationship even if the target is online to the S/390 or zSeries host.

This option does not apply to open-systems volumes.

Persistent FlashCopy

Select this option if you want the FlashCopy relationship to remain even after the FlashCopy operation completes. This prevents another FlashCopy task from writing on your target volume before you have withdrawn the FlashCopy relationship. The FlashCopy relationship between source and target volumes remains indefinitely and must be broken with a Withdraw FlashCopy task.

Inband command

Select this option to allow a FlashCopy relationship between two volumes to be established at a remote ESS using inband commands (the command will be sent over an established PPRC Path). For details, refer to 3.4.3, "FlashCopy V2 establishment options" on page 41 and the paragraph 6.11.11, "Inband FlashCopy" on page 229 later in this chapter.

Freeze FlashCopy consistency groups

Select this option to freeze write activity to each LSS containing FlashCopy source volumes that are in a consistency group. A consistency group is a group of volumes participating in FlashCopy relationships that need to be kept in a consistent state to ensure data integrity.

The freeze option provides a mechanism to ensure data consistency across multiple FlashCopy volumes in an LSS, or across LSSs, or across Model 2105 ESSs. The freeze option causes volumes to remain in a Queue Full condition (volumes are not available) until the condition is reset or the time-out value expires (default is two minutes).

Inhibit writes to target

Select this option to ensure that writes are inhibited on the target volume until an incremental FlashCopy operation is complete. This ensures data consistency on the target volume. If you select the Inhibit writes to target option, the change recording option will not be effective on the target volume.

Start change recording

Select this option to create the bitmaps that record the changed tracks on both volumes participating in a FlashCopy relationship. For details, refer to 3.4.3, "FlashCopy V2

establishment options" on page 41 and paragraph 6.11.6, "Incremental FlashCopy" on page 216.

Increment FlashCopy
Select this option to create a new point-in-time data copy without copying all of the tracks from the source volume to the target volume. This feature provides an alternative to copying an entire volume for each point-in-time data copy. For details, refer to 3.4.3, "FlashCopy V2 establishment options" on page 41 and paragraph 6.11.6, "Incremental FlashCopy" on page 216.

Reverse restore
Select this option to reverse the FlashCopy relationship and copy the modified tracks from the target volume to the source volume. The background copy process must complete before you can reverse the order of the FlashCopy relationship from its original source and target relationship. For details, refer to 3.4, "FlashCopy tasks and options" on page 39 and the paragraph 6.11.6, "Incremental FlashCopy" on page 216 later.

Running or saving the task
Once you select the FlashCopy options, click on **Next**. From the next panel, as in Figure 6-37, you will decide whether to **Save**, **Run**, or **Cancel** the task.
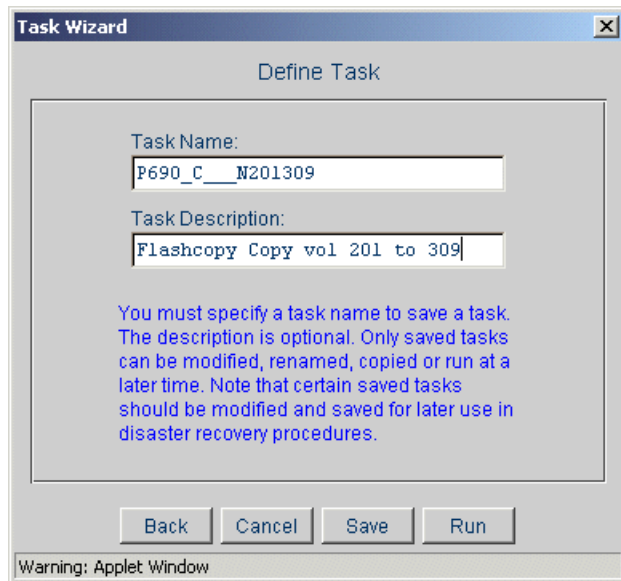
Figure 6-37   Define task panel

If you save the task, it can be executed from the task panel at any time.

When the FlashCopy is issued, a bitmap is created for the datacopy from the source to the target volumes. The time to build this bitmap is only a few seconds (depending on the size of the volumes). Once the bitmap is created, the FlashCopy relationship is established and the two volumes are immediately available.

Once the FlashCopy relationship is established, the display of the volumes in the volume panel changes (see Figure 6-38 on page 208):

► A blue lightning bolt appears next to the FlashCopy source volume.

► A red lightning bolt appears next to the FlashCopy target volume.

► A "I-am-a-target" icon appears right to the red lightning bolt next to the target volume. This icon is a small gray square. The icon remains when the FlashCopy relationship ends after, for example, a background copy is finished. To remove this icon, you must withdraw the FlashCopy relationship with a "withdraw to target" task (see 3.4.4, "FlashCopy V2 withdraw options" on page 42).

► The lightning bolts disappear when the copy operation is complete. If the FlashCopy relationship is established with a **No copy** option or with the **Persistent FlashCopy** option, the icons remain until a withdraw operation is performed.

► You can get information about the progress of the FlashCopy background copy by clicking on the information panel.

Figure 6-38   FlashCopy lightning bolts icons

Tip: After you have created and saved the task to establish the FlashCopy with the **No background copy** option, you need to create the corresponding task to withdraw the pair. We recommend that you modify the task that you just created so that it withdraws the pair, then name and save the task. This not only saves time, but it is also useful for disaster recovery. For further information, refer to 6.8.5, "Modifying a task" on page 199.

Establishing FlashCopy relationship using the Logical Subsystem panel

To establish FlashCopy relationships between all the volumes in two LSSs, you can create a task at the LSS level in the Logical Subsystem panel (LSS panel).

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 6-2 on page 187).

Click on the source LSS and the LSS ID turns blue. Right-click the target LSS and the LSS ID turns red. Right-click again on the target LSS and the Task Wizard opens. Select the options as you would do for volumes pairs.

6.11.2 Getting information about a FlashCopy pair

By selecting one volume in the volume panel and clicking the Information button, you get information about the state of the FlashCopy pairs the volume is involved in.

In the example in Figure 6-39 on page 209, we selected a volume that is the source for two FlashCopy relationships. The target volumes IDs are in the Full Volume FlashCopy Peers panel.
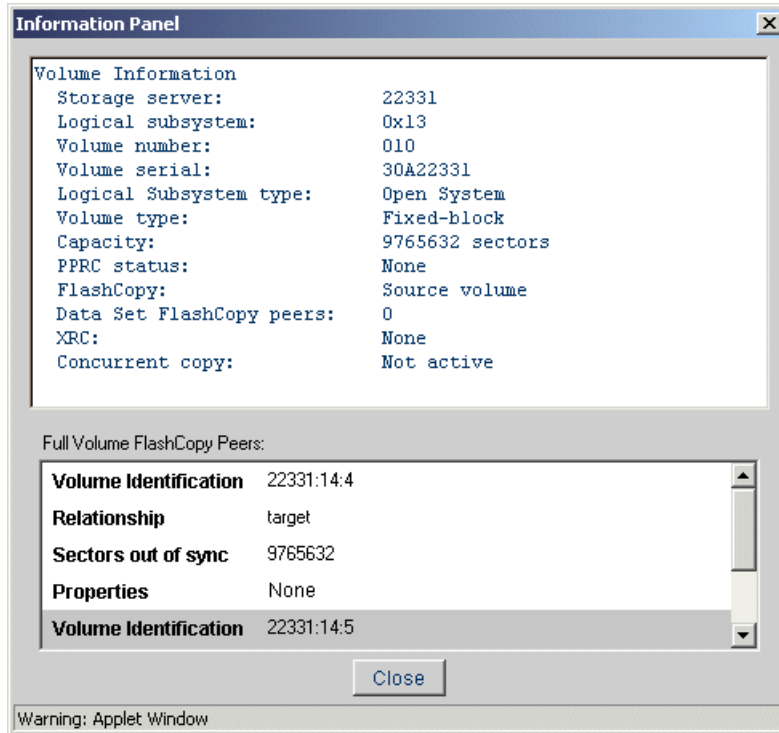
Figure 6-39   Volume Information Panel

6.11.3 Withdrawing a FlashCopy pair

In the following cases, you have to *withdraw* a FlashCopy pair:

► If a FlashCopy pair is not needed anymore, but has not yet finished the background copy.

► If a FlashCopy pair that was created with the **Do not perform background copy** option is not needed anymore.

► If a FlashCopy pair that was created with the **Persistent FlashCopy** option is not needed anymore.

The command to **Withdraw FlashCopy pair** results in the termination of the FlashCopy operation. All the data on the target can be corrupted and unusable if it is inadvertently written on by another application.

You can use the Volumes panel or the Logical Subsystem panel to withdraw a FlashCopy pair.

Open the Volumes panel and select the source and target LSS within which you want to perform the withdraw FlashCopy in the Source and Target panels. You always need to have two components to withdraw a FlashCopy pair, a source and a target. Select the source volume and right-click on the target. If you have selected the wrong source or target volume, just click on the correct source volume again. See Figure 6-40 on page 210 for an example.
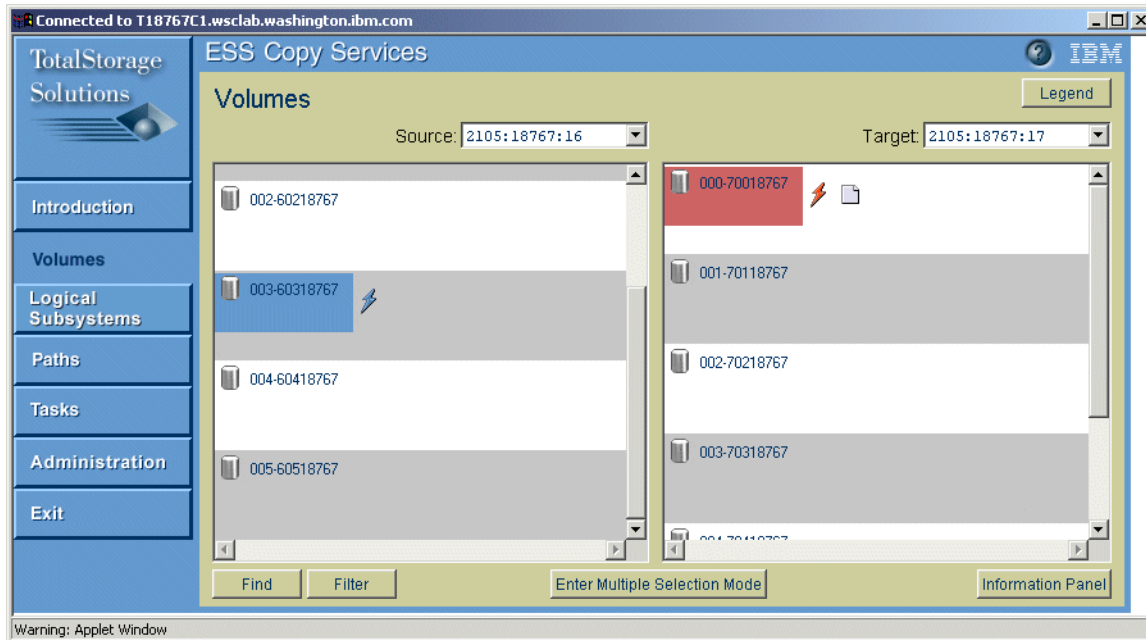
Figure 6-40   Volumes panel: Select volumes

Once the source and target are selected, start the Task Wizard (Figure 6-41) by right-clicking on the target volume again.
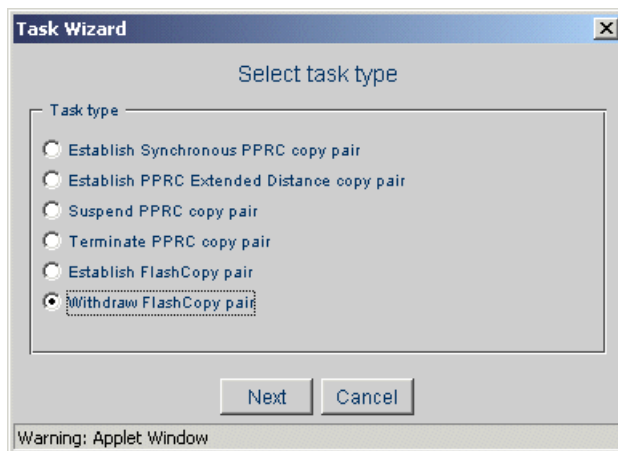
Figure 6-41   Task Wizard: Select task type

Select the **Withdraw FlashCopy pair** option and click on **Next** to select if you want to issue this command using an established PPRC path (inband command). See 6.11.11, "Inband FlashCopy" on page 229 for more details.
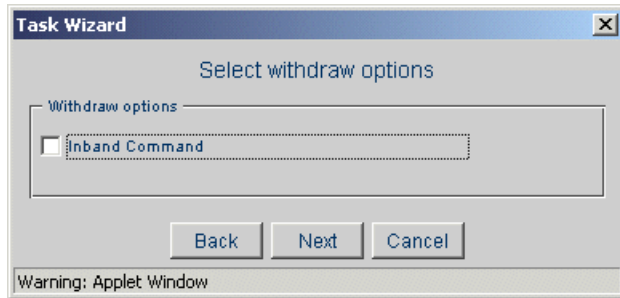
Figure 6-42   Task Wizard: Select withdraw options

Click on **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (see Figure 6-43).
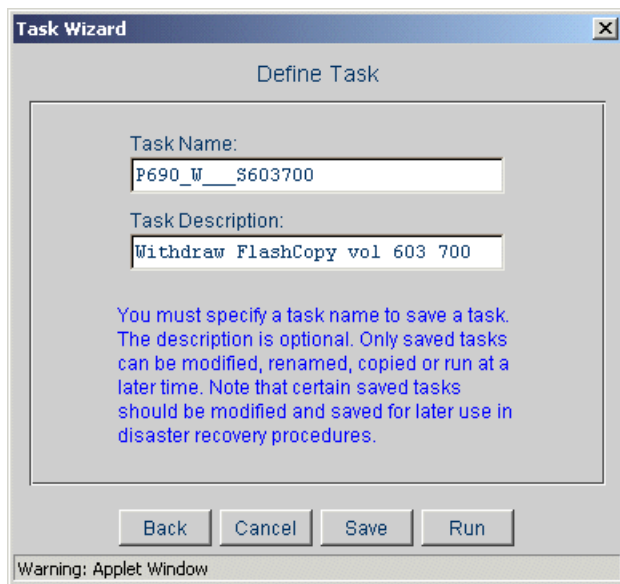
Figure 6-43   Task Wizard: Define Task

Withdrawing FlashCopy pairs using the Logical Subsystem panel

To withdraw all FlashCopy pairs between two logical subsystems, you can create a task at the LSS level in the Logical Subsystem panel.

> **Note:** All the FlashCopy relationships that have source volumes in the selected source LSS and their associated targets in the selected target LSS will be terminated.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 6-2 on page 187).

Click on the source LSS and the LSS ID turns blue. Right-click on the target LSS and the LSS ID turns red. Right-click again on the target LSS and the Task Wizard opens. Select the options as you would do for volumes pairs.

Withdraw FlashCopy to the target

Use the option **Withdraw FlashCopy to the target** against the target volume when you want to remove its FlashCopy icon (represented by a white square at the lower left of the volume

icon; see Table 6-1 on page 181) after the background copy has been successfully completed. Typically, you would remove the indicator because you have decided that it is no longer necessary to mark the volume as a copy of another.

> **Attention:** If you perform a Withdraw to target action to a target volume of an established relationship, it will withdraw the relationship and remove the "I-am_a_target" icon.

To perform this action, click on the target volume and right-click on the target volume. The volume icon and the volume ID turns gray (see Figure 6-44).
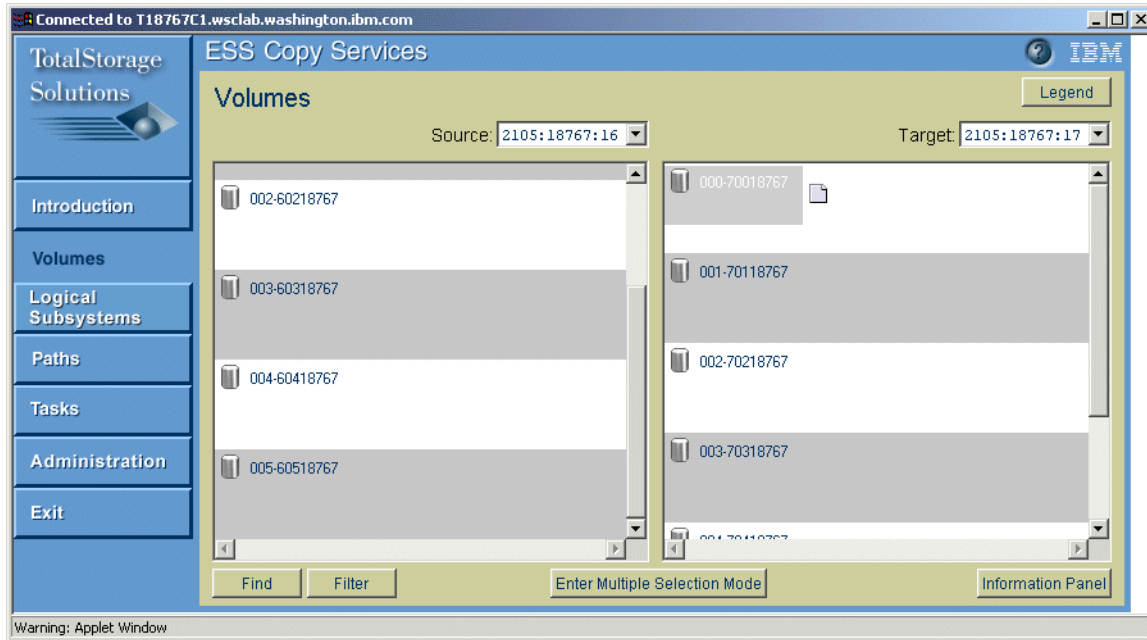
Figure 6-44   Volumes panel: Selecting a volume as source and target for FlashCopy withdraw

Right-click on the target volume again to open the Task Wizard (Figure 6-45).
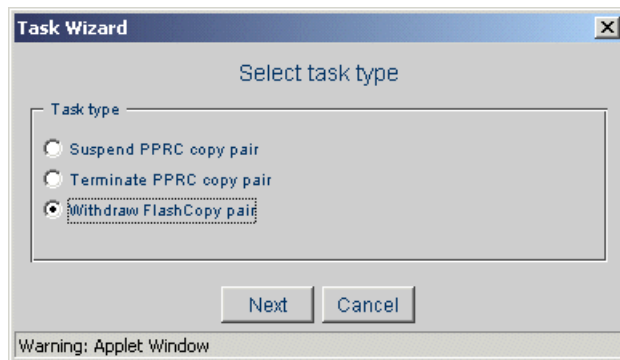
Figure 6-45   Task Wizard: Select task type

Select the **Withdraw FlashCopy pair** option and then click **Next** to select the withdraw options (Figure 6-46 on page 213).

Figure 6-46   Task Wizard: Select withdraw options

Click on **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-47).

Figure 6-47   Task Wizard: Define task

> **Note:** If you perform the **Withdraw FlashCopy to the target**, make sure you select a target volume of a FlashCopy relationship in the volumes panel. You cannot perform this action by selecting a source volume.

Withdrawing FlashCopy to the target using the Logical Subsystem panel

To withdraw all FlashCopy relationships to the target for all target volumes in the same logical subsystem, you can create a task at the LSS level in the Logical Subsystem panel.

> **Note:** All FlashCopy relationships that have a target volume in the LSS that you select will be terminated.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystems panel do have different meanings (for further information, refer to Table 6-2 on page 187).

Click on the LSS and the LSS ID turns blue. Right-click on the same LSS and the LSS ID turns gray. Right-click again on the LSS and the Task Wizard opens. Select the options as you would do for volumes pairs.

6.11.4 Multiple FlashCopy using a single source volume

Multiple FlashCopy using the same volume is possible with LIC 2.2.0 or higher and FlashCopy V2 feature installed. FlashCopy V2 enables you to have multiple FlashCopy relationships using the same source volume.

To do this, create different FlashCopy tasks using the same source volume, and then establish each FlashCopy relationship by running each individual task.

When a volume is the source of more than one established FlashCopy relationship, a specific Multiple Relationship FlashCopy Icon will appear next to the FlashCopy Source icon, as shown in the Figure 6-48.
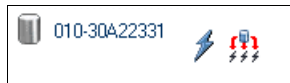
Figure 6-48   FlashCopy source and multiple relationship FlashCopy icons

Restriction: The maximum FlashCopy relationships a volume can be the source for is 12.

Restriction: Among the Multiple FlashCopy relationships the volume is the source for, only one can be a FlashCopy relationship with the **Start Change Recording** option (this option is a prerequisite prior to establishing an incremental or a reverse restore FlashCopy relationship).

6.11.5 FlashCopy Start Background Copy

If you have performed a FlashCopy pair with the **No background copy** option and you need to create a permanent physical copy of the data for backup or disaster recovery purposes, you can perform a **FlashCopy Start Background Copy**. All data will be copied from the source to the target. After the copy is complete, the FlashCopy relationship is terminated, unless the FlashCopy relationship is persistent, in which case, the relationship remains.

FlashCopy Start Background Copy using the Volumes panel
Open the **Volumes** panel and select in the source LSS and the target LSS in the source and target panels.

To perform a **FlashCopy Start Background Copy**, click on the source volume of the FlashCopy pair. Then right-click on the source volume (the target ID of the volume turns gray, as presented in Figure 6-49 on page 215).

Note: If you perform **FlashCopy Start Background Copy**, make sure you select a source volume of a FlashCopy relationship in the volumes panel. You cannot perform this action by selecting a target volume.
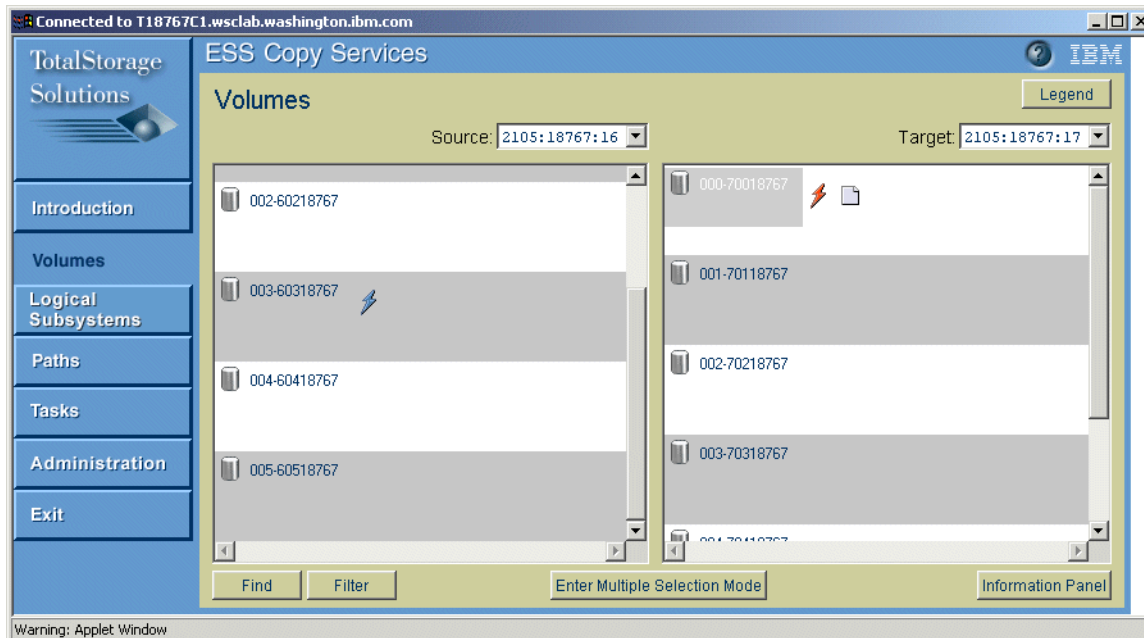
Figure 6-49   Volumes panel: The target volume selected as source and target

Then, right-click on the target volume again to open the Task Wizard (Figure 6-50).
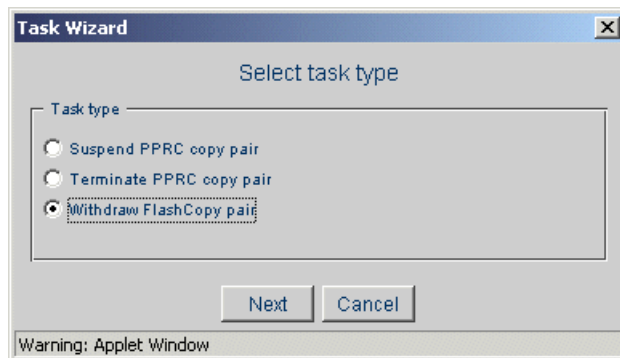
Figure 6-50   Task Wizard: Select task type

Select the **Withdraw FlashCopy pair** option and click **Next**. In the next panel, you select the **FlashCopy Start Background Copy** option (Figure 6-51).
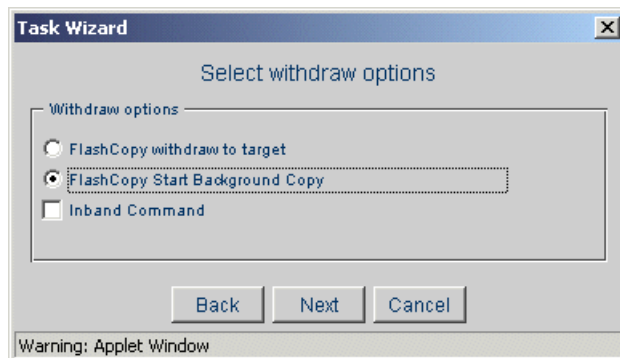
Figure 6-51   Task Wizard: Select withdraw options

Click the **Next** button and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-52).
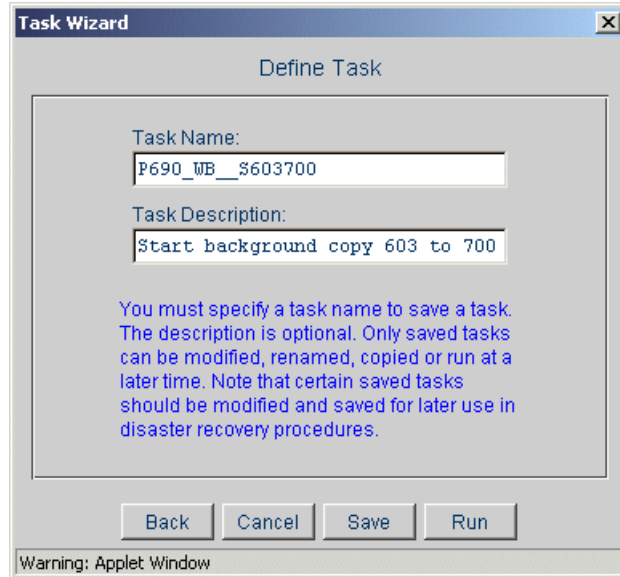
Figure 6-52   Task Wizard: Define Task

Once you run the task, the background copy task copies the data from the source to the target. After the copy is complete, the FlashCopy relationship is terminated and the lightning bolts disappear on the source and the target volumes, unless the FlashCopy relationship is persistent, in which case, the relationship remains.

FlashCopy Start Background Copy using the LSS panel

To start a background copy of all FlashCopy pairs that have a source volume on a logical subsystem, you can use the Logical Subsystem panel.

Note: When you run the task at the LSS level, be aware that the task will run on all the source volumes in that LSS.

Open the Logical Subsystems panel and select the ESS in the Select list at the upper right corner. A panel showing all the LSSs is displayed, with the appropriate name and serial number indicated below each LSS icon. The different icons on the Logical Subsystem Panel do have different meanings (for further information, refer to Table 6-2 on page 187).

Click on the desired LSS. The LSS ID turns blue. Right-click on the LSS and the LSS ID turns gray. Right-click on the LSS again and the Task Wizard opens. Select the options as you would do for volumes pairs.

6.11.6 Incremental FlashCopy

Incremental FlashCopy is possible with LIC 2.2.0 or higher and the FlashCopy V2 feature installed. With the incremental FlashCopy function, it is no longer necessary to copy an entire volume for each point-in-time copy. Instead, only tracks that have changed on the source volume since the last incremental FlashCopy are copied to the target volume. To enable incremental FlashCopy, the ESS uses a feature called *change recording* to monitor writes and record changes that were made to volumes participating in FlashCopy relationship since the initial (or last) incremental FlashCopy was performed.

A source volume can participate in multiple (up to 12) FlashCopy relationships. However, when you establish a FlashCopy pair with the **Start Change Recording** option on a volume pair, the option only applies to that specific pair, regardless if the same source volume is participating in multiple FlashCopy relationships. Incremental FlashCopy operations (with the start change recording featured enabled) are performed on a one-to-one basis, that is, incremental changes are copied from one source volume to one target volume.

The following describes how you can enable an incremental FlashCopy:

► The initial setup requires that you enable the **Start Change Recording** option when you establish the initial FlashCopy volume pair. This process copies the entire source volume to the target volume to keep the volumes in sync and activates the change recording feature. See "Initial task: Establishing a FlashCopy pair to start change recording" on page 217 for instructions.

► Each time you want only incremental updates made to the target volume, you must establish the FlashCopy operation with the **Increment FlashCopy** option. Note that the **Start Change Recording** option should also be selected. The ESS copies only data that has changed since the last incremental copy and avoids recopying data that has not changed. See "Establishing an Incremental FlashCopy" on page 219 for instructions.

Note:

► If you perform an incremental FlashCopy on a FlashCopy pair without the required change recording activated, the operation fails. You must withdraw the FlashCopy relationship, then re-establish the FlashCopy relationship with the Start change recording enabled.

► The Persistent FlashCopy option is enabled for incremental FlashCopy operations. Because the FlashCopy relationship between source and target volumes remains indefinitely, you can withdraw the relationship with a withdraw FlashCopy pair task.

Initial task: Establishing a FlashCopy pair to start change recording

Open the Volumes panel and select the source LSS in the left panel, then select the target LSS in the right panel. Select the source volume and, with a right-click, select the target volume. With another right-click on the target volume, you will open the Task Wizard. Select the **Establish FlashCopy pair** option (Figure 6-53 on page 218), and then click on **Next**.
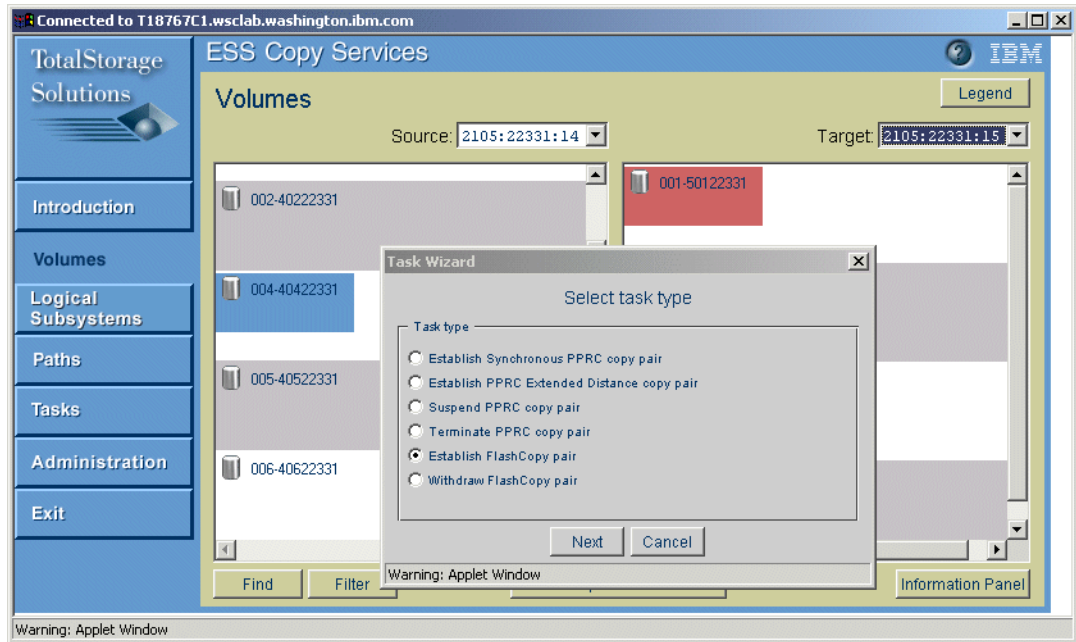
Figure 6-53   Volumes panel: Select volumes

In the next panel, you have to select the **Start Change Recording** option (Figure 6-54). This option will enable you to run an incremental FlashCopy task or a reverse restore FlashCopy task later using these two volumes. When you select the **Start Change Recording** option, the **Persistent FlashCopy** option is automatically selected (Figure 6-54). You can also select other options if needed.
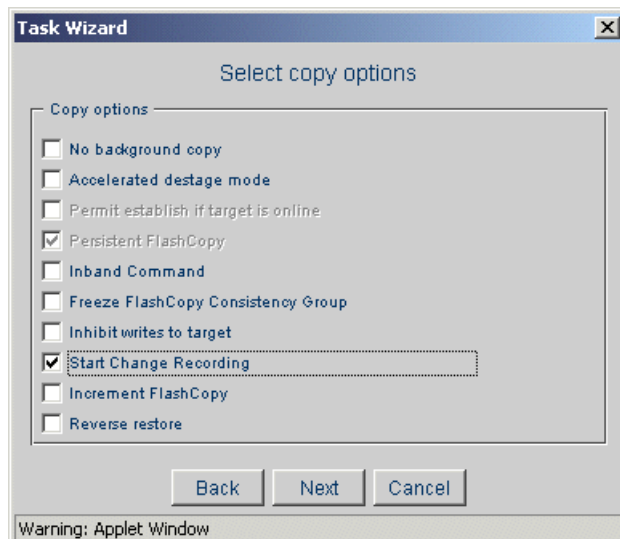
Figure 6-54   Task Wizard: Select task type

Once the options are selected, click on **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-55 on page 219).
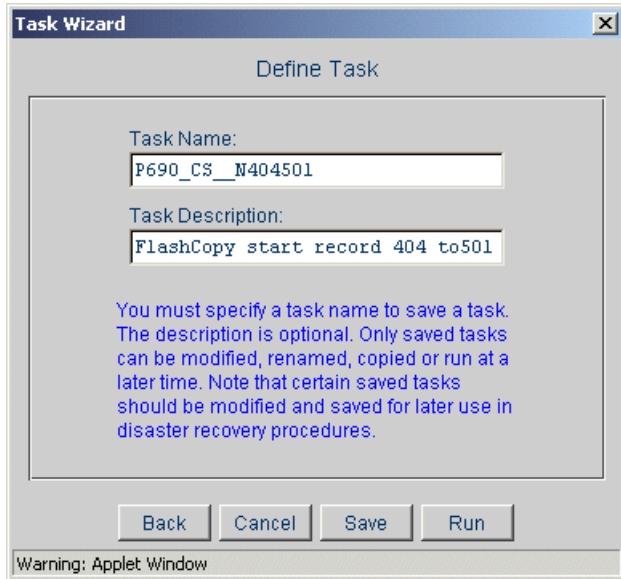
Figure 6-55   Task panel: Define task

Once the FlashCopy is established, a new icon will appear next to the FlashCopy lightning bolts in the source and target volume. This new icon is a yellow and red lightning bolt above a green triangle (see Figure 6-56 for an example). This Icon means that the **Start Change Recording** option is enabled on the source and target volumes and that the volumes are eligible for an incremental FlashCopy.
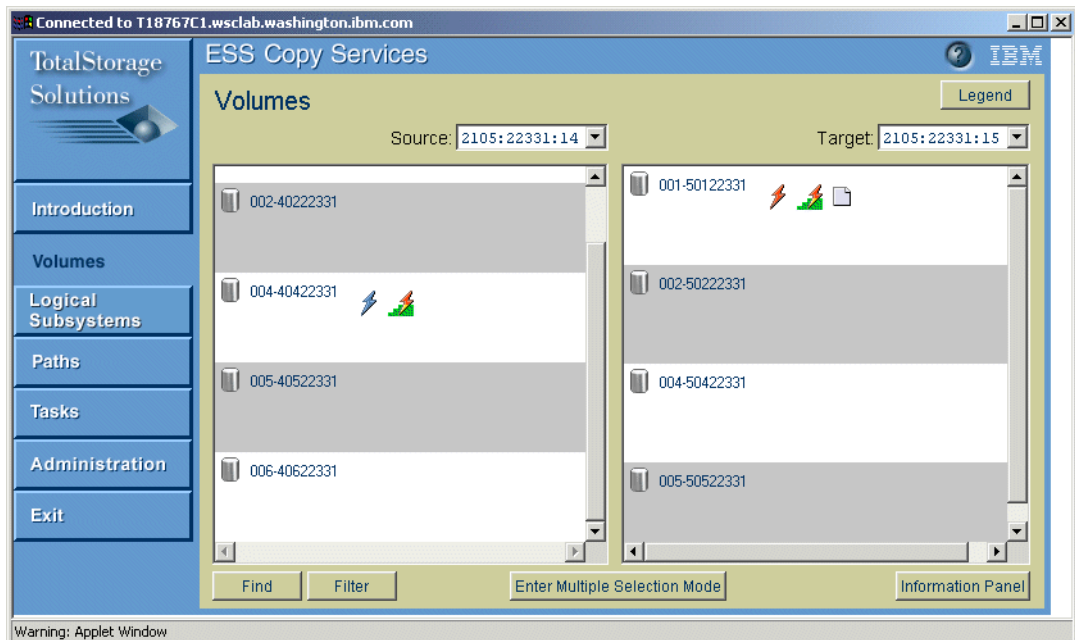
Figure 6-56   Volumes panel: Result of a Start Recording FlashCopy

Establishing an Incremental FlashCopy

You can perform an Incremental FlashCopy using two volumes in a FlashCopy relationship with **Start Change Recording** enabled. To verify that the Start Change Recording is enabled on a FlashCopy pair, you can select the source volume and click on **Information Panel** to check for the FlashCopy properties (Figure 6-57 on page 220).
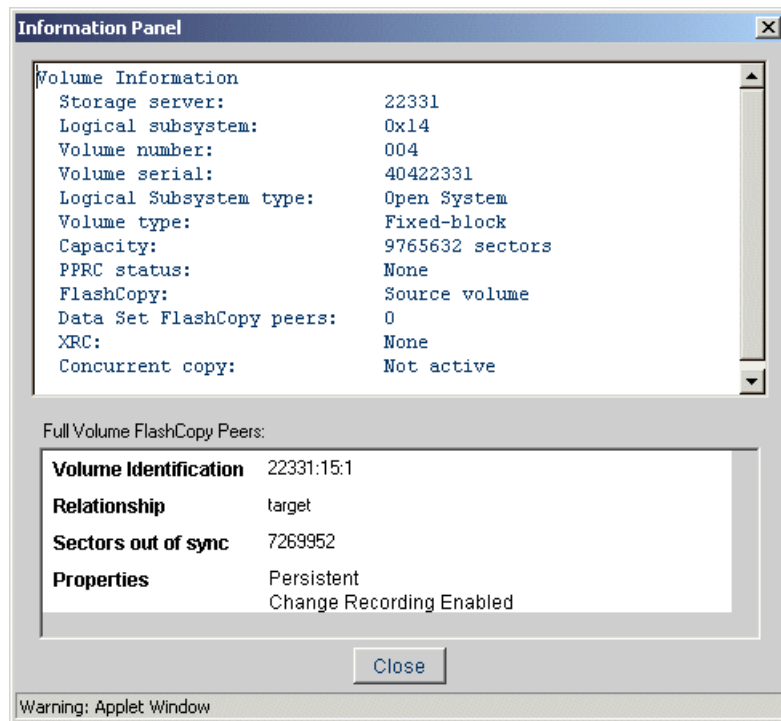
Figure 6-57   Information panel: Source volume and start recording

Open the Volumes panel and select the source LSS in the left panel, then select the target LSS in the right panel. Select the source volume and, with a right-click, select the target volume. With another right-click on the target volume, you will open the Task Wizard. Select **Establish FlashCopy pair** option (Figure 6-58), and then click on **Next**.
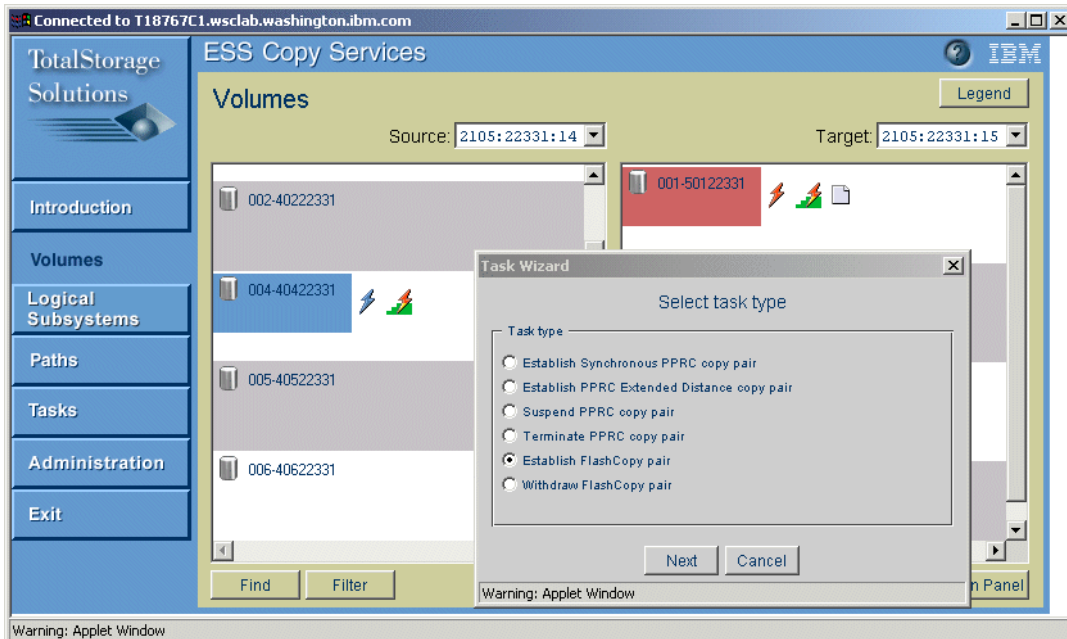
Figure 6-58   Volumes panel: Selecting volume for an incremental FlashCopy

Select the **Incremental FlashCopy** and **Start Change Recording** options. The **Persistent FlashCopy** option is automatically selected when you click on the **Start Change Recording** option (Figure 6-59).

> **Important:** When you perform an incremental FlashCopy operation, change recording is not re-enabled by default. You must select the **Start Change Recording** option each time you establish an incremental FlashCopy operation to ensure that change recording is maintained on the volume pair. If you do not select the **Start Change Recording** option, you will not be able to run another incremental FlashCopy or a reverse restore FlashCopy after you run the task.
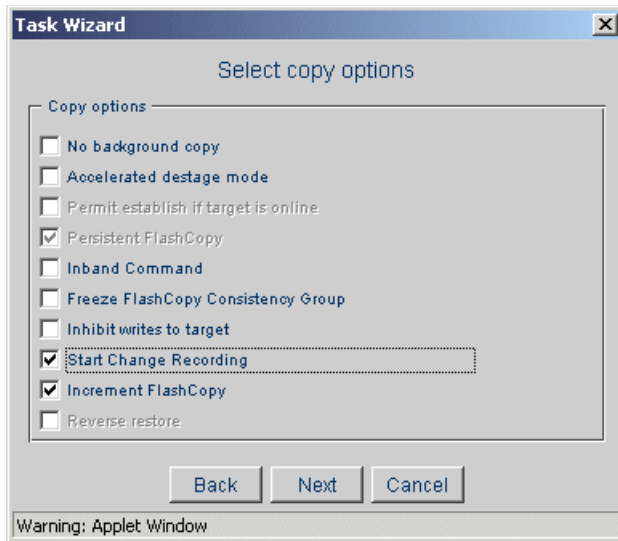
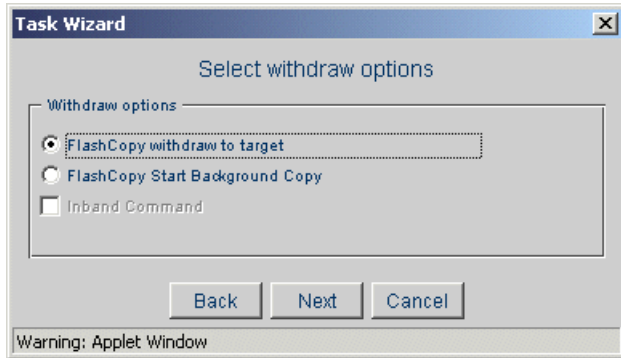*Figure 6-59   Task Wizard: Select copy options*

Click on **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-60).
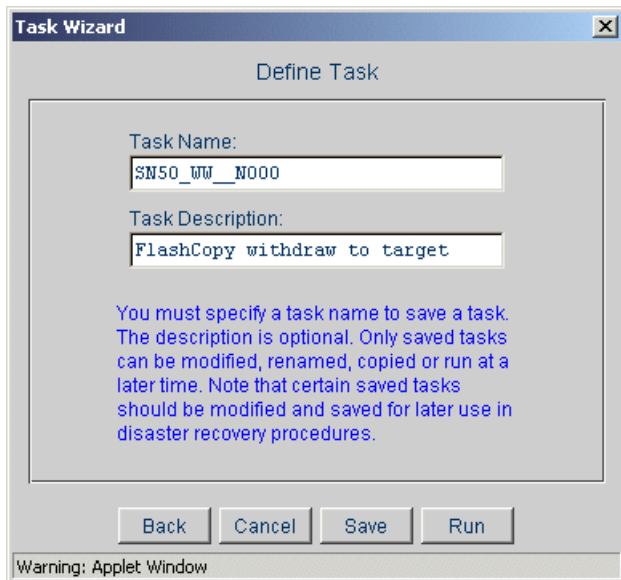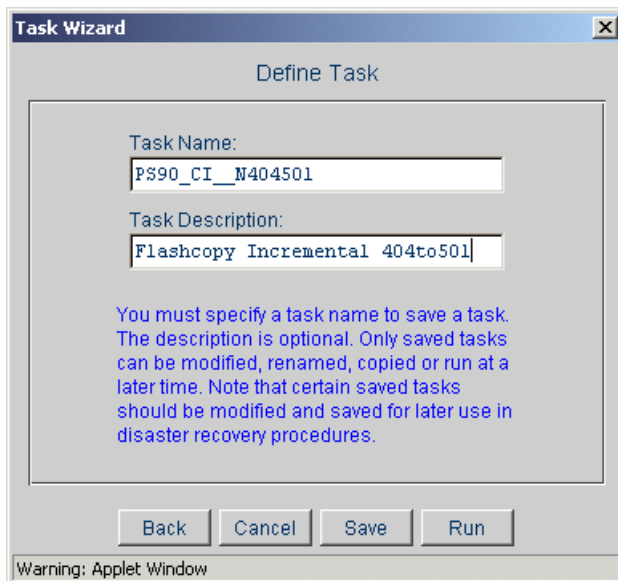

Figure 6-60   Task Wizard: Define Task

6.11.7 Reversing a FlashCopy relationship

Reversing a FlashCopy is possible with LIC 2.2.0 or higher and the FlashCopy V2 feature installed. When used within an established FlashCopy pair, this option allows you to restore the source volume with the data from the target volume. To get more details about how the reverse restore functions operates, refer to 3.3.3, "Incremental FlashCopy and the Reverse Restore option" on page 36.

You can perform a **Reverse restore** FlashCopy using two volumes in a FlashCopy relationship with **Start Change Recording** enabled and only if the background copy between the two volumes is finished.

To verify that the **Start Change Recording** is enabled on a FlashCopy pair and that the background copy is finished, you can select the source volume and click on **Information Panel** to check for the FlashCopy properties. The number of **Sectors out of sync** must be 0, indicating that the background copy if finished (Figure 6-61).
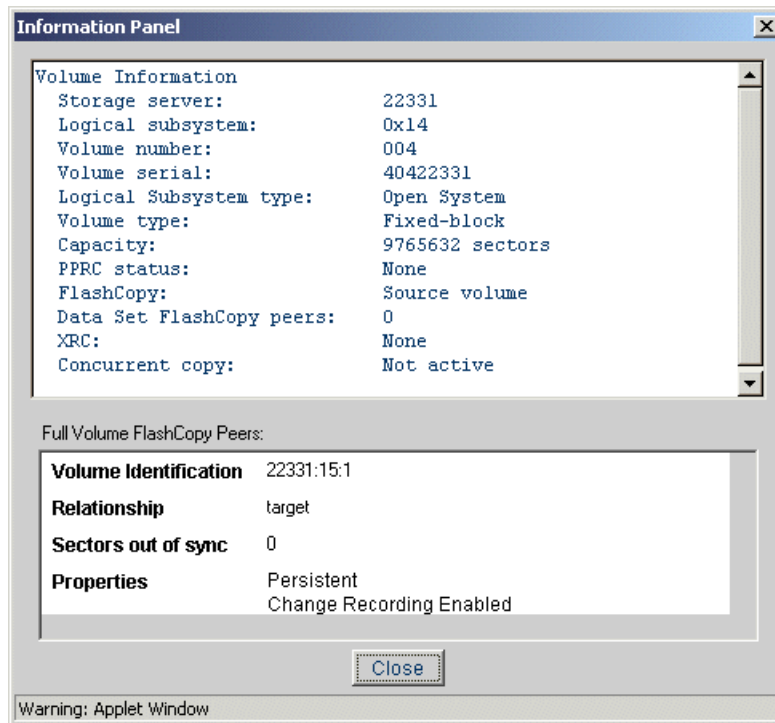
Figure 6-61   Change Recording enabled: Background copy finished

Open the Volumes panel and select the source LSS in the left panel, then select the target LSS in the right panel. Select the source volume and, with a right-click, select the target volume. With another right-click on the target volume, you will open the Task Wizard. Select the **Establish FlashCopy pair** option (Figure 6-62 on page 223).
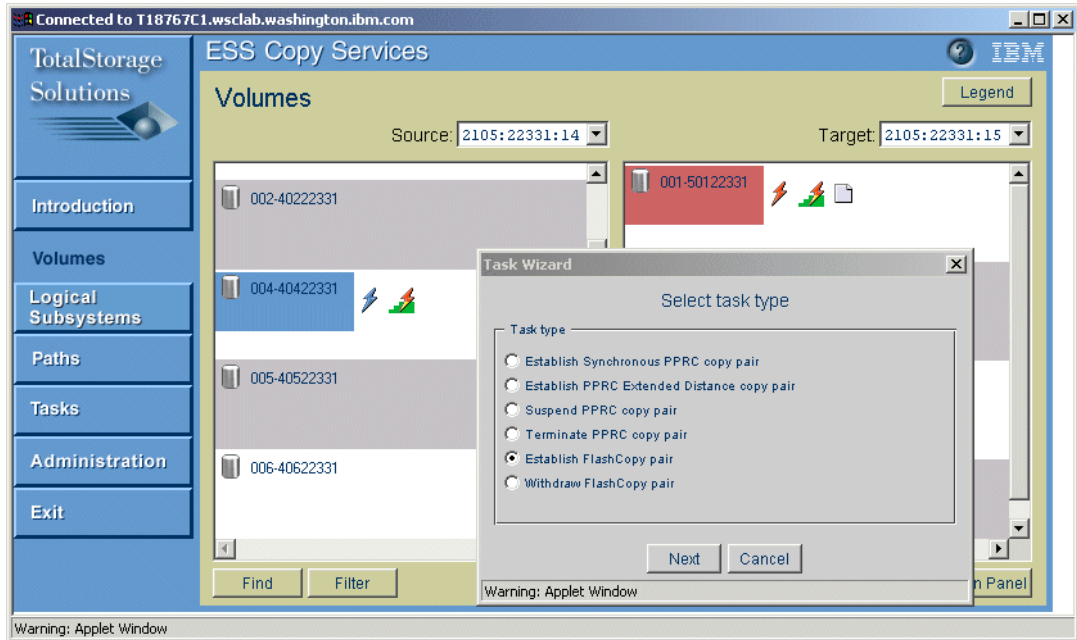
Figure 6-62   Volumes panel: Selecting volume for a Reverse Restore FlashCopy

Once the **Establish FlashCopy pair** task is selected, click on **Next**.

Select the **Reverse Restore** and **Start Change Recording** options. The **Persistent FlashCopy** option is automatically selected when you click on **Start Change Recording** (Figure 6-63).

> **Important:** If you do not select the **Start Change Recording** option, you will not be able to run another incremental or another reverse restore FlashCopy in the future after you run this task.
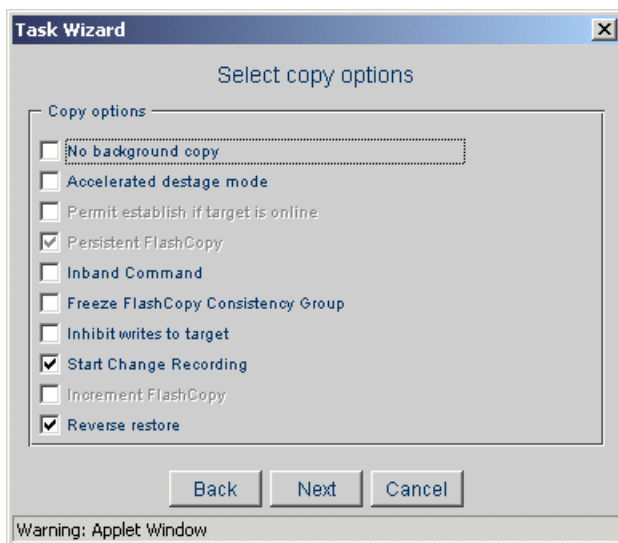
Figure 6-63   Task Wizard: Select copy options

Click on **Next** and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-64).
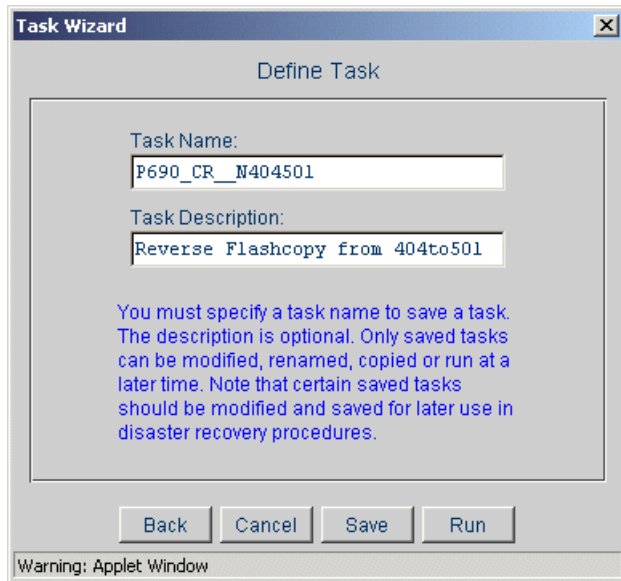
Figure 6-64   Task Wizard: Define Task

6.11.8 Creating a FlashCopy consistency group

The FlashCopy V2 feature also has two new options: the **Freeze FlashCopy pairs** and **FlashCopy Consistency created** options. You can combine those two options to create a consistency group across several volumes in different LSSs and even in different ESSs in the same Copy Services domain. Refer to 3.4.5, "Consistency Created task with FlashCopy V2" on page 43 to get more details about how this options works.

Here are the steps required to build a FlashCopy consistency group:

1. Create a task to Freeze a FlashCopy consistency group. To do this, you will:

   a. Create and save tasks with the Freeze options for each FlashCopy pair.

   b. Group all the tasks into one single Freeze task.

2. Create a task to Thaw the FlashCopy consistency group. To do this, you will:

   a. Create a task with **FlashCopy Consistency group created** option for each LSS containing source volumes.

   b. Group all the tasks into one single Consistency created task.

3. Invoke the Freeze task to freeze the FlashCopy consistency group. You can invoke the task using the Copy Services Web user interface or the CLI.

4. Invoke the Consistency created task. You can invoke the task using the Copy Services Web user interface or the CLI.

6.11.9 Freezing a FlashCopy consistency group

The FlashCopy Establish command can freeze write activity to FlashCopy source volumes in a consistency group. The freeze option ensures data consistency across multiple FlashCopy volumes in an LSS, across LSSs, or across Model 2105 ESSs.

Considerations

► You create a FlashCopy consistency group by grouping several FlashCopy tasks with the **Freeze FlashCopy consistency group** option specified within one single group task and running this group task. If you want to create different FlashCopy consistency groups, you will create different group tasks.

► When you establish a FlashCopy operation (with the **Freeze FlashCopy consistency group** option specified), the freeze is issued to each volume in the consistency group. This causes the volumes to remain in a Queue Full condition (volumes are not available to the application) until the condition is reset (see 6.11.10, "Thawing a FlashCopy consistency group" on page 228) or the time-out value expires (The default is two minutes. You can use the LSS properties panel to modify the default value. See 6.6.4, "LSS properties" on page 190).

► When source volumes are in a frozen state, and after a FlashCopy operation (with the freeze option specified) has been successfully issued to each volume in the consistency group, the FlashCopy target volumes will be in a consistent state.

To create tasks that will freeze a FlashCopy consistency group:

1. Use the volume panel to build the tasks that will establish the FlashCopy relationship with a **Freeze FlashCopy consistency group** option. Build one task for each FlashCopy pairs. You can use the multiple selection mode when selecting several volumes in the same LSS.

2. Use the task panel to group the FlashCopy tasks with the freeze option into one single task for all options that need to be in the consistency groups.

Establishing a FlashCopy with freeze option

In the volume panel, in the source area, select the LSS that contains the source volume. To select the LSS, you will use the source drop-down list. When the source LSS is selected, select the LSS that contains the target volume in the target area.

Click on the source volume. Once the source volume is selected, right-click on the target volume. You can use the Multiple Selection Mode to select multiple source and target volume pairs (Figure 6-65 on page 226).

Figure 6-65   Volumes panel: Selecting volumes using Multiple Selection Mode

To open the Task Wizard, right-click again on the last target volume you selected.

In the Select Task type panel, select the **Establish FlashCopy** task type (Figure 6-66).
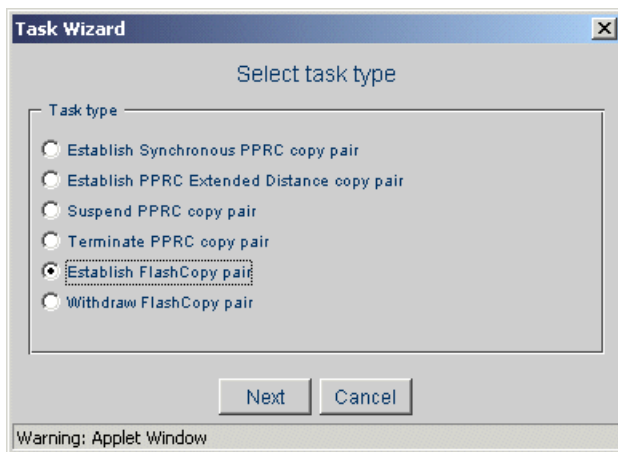
Figure 6-66   Task Wizard: Select task type

In the select copy options panel, click on **Freeze FlashCopy Consistency group**.
Eventually, select other copy options for the task (Figure 6-67 on page 227), and then click on
Next.
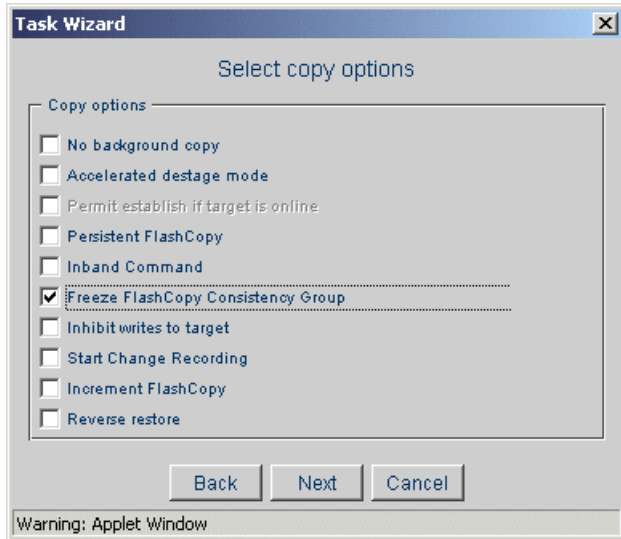
Figure 6-67   Task Wizard: Select copy options panel

In the Define task panel (Figure 6-68), either give a task name (mandatory) and a task description (optional), then click on **Save** to run the task later, or click on **Run** to run the task immediately.
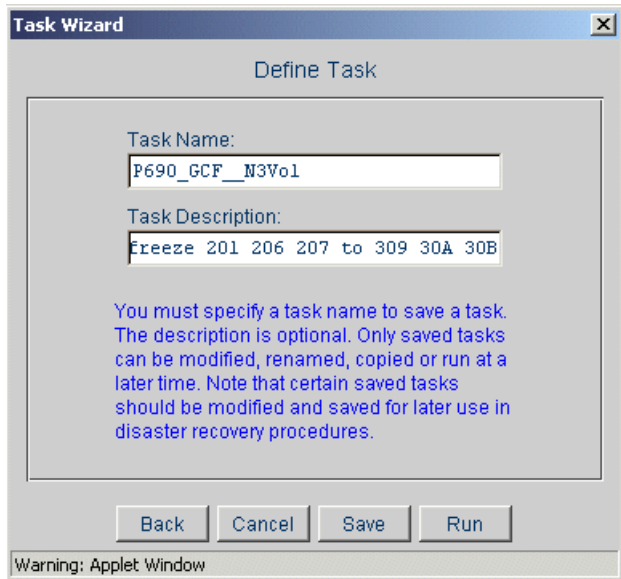
Figure 6-68   Task Wizard: Define Task panel

Grouping tasks

1. In the task panel, select all the FlashCopy tasks you want to group. Use the Shift and Ctrl keys to select multiple tasks.

2. Click on a group.

3. Give the task a name (mandatory) and a description (optional), and then save the task.

6.11.10 Thawing a FlashCopy consistency group

The thawing of a FlashCopy consistency group (consistency group created operation) follows the successful freeze for all FlashCopy volumes in the consistency group. It allows normal I/O operations from the host to the ESS to resume by removing the Queue full condition on a group of FlashCopy Source volumes.

Considerations

The Consistency created task is performed at the LSS level. This task will resume the Queue Full condition for all the volumes in the LSS. When the source volume in the consistency group are spread across LSS or across ESSs, you will need to perform the following actions:

► Create a **Consistency Created** task on each LSS containing a FlashCopy source with the Queue-full condition.

► Group all the tasks into one single task

► Run the task.

Creating a Consistency Created task

In the LSS panel, select the ESS using the **Select** drop-down menu. Click on one LSS and right-click on the same LSS. The LSS number will turn gray. Right-click again on the LSS.

In the Select task type, select **Consistency Created** (Figure 6-69), and then click **Next**.
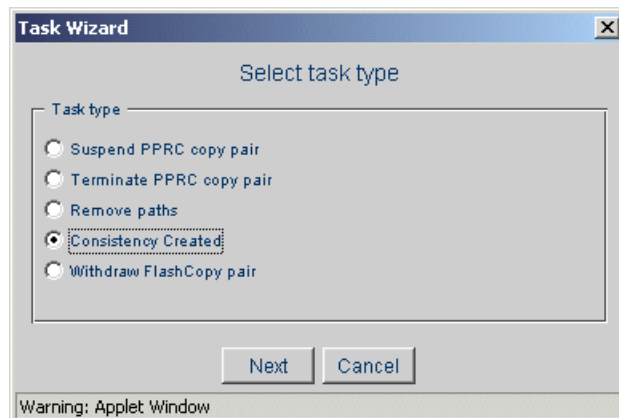
Figure 6-69   Task Wizard: Select task type

In the Select copy options panel, select **FlashCopy Consistency group (One LSS selected)** (Figure 6-70), and then click **Next.**
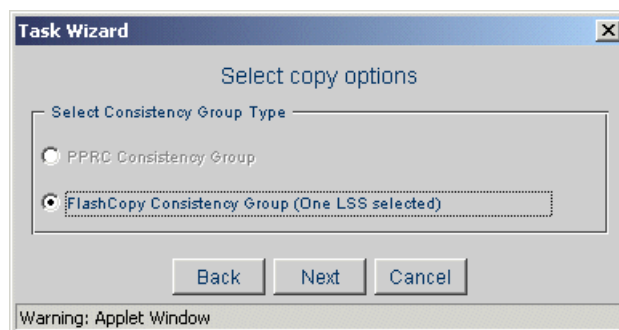
Figure 6-70   Task Wizard: Select copy options

In the Define task panel (Figure 6-71), either give a task name (mandatory) and a task description (optional) and then click on **Save** to run the task later, or click on **Run** to run the task immediately.
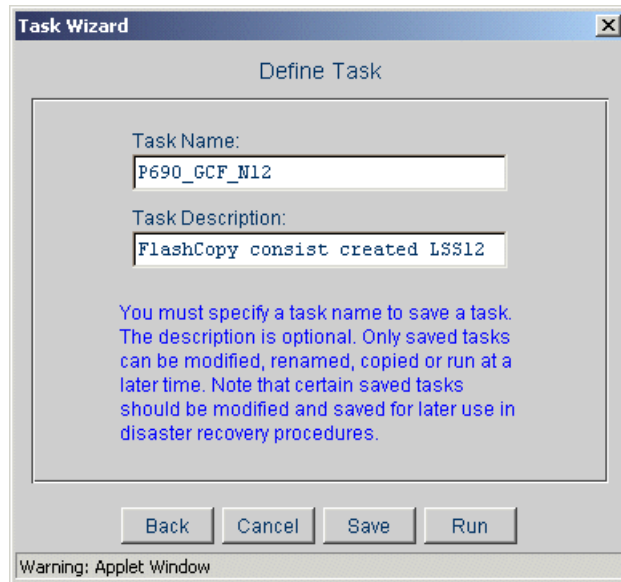
Figure 6-71   Task Window: Define Task

Grouping tasks

1. In the task panel, select all the FlashCopy tasks you want to group. Use the Shift and Ctrl keys to select multiple tasks.

2. Click on a group.

3. Give the task a name (mandatory) and a description (optional), and then save the task.

6.11.11 Inband FlashCopy

You can establish a FlashCopy pair at a remote ESS using inband commands. Inband commands allow you to establish and withdraw a FlashCopy pair at the remote site, without having a Copy Services Web interface connection to the remote site (this connection requires a LAN connection).

Inband commands are issued to a primary volume of a PPRC pair at a local ESS and sent across PPRC paths (acting as a conduit) to a remote ESS to enable a FlashCopy pair to be established at the remote site. The source volume you select for the FlashCopy at the remote site must be the secondary PPRC volume of the PPRC pair. (The PPRC secondary volume becomes the source volume in a FlashCopy pair.)

> **Important:** When you create the FlashCopy tasks with the Inband option, the remote ESS must be connected to the ESS Web domain. The connection is not mandatory when you issue the command (when you run the task) but is mandatory when you define (create) the command.

The task definition process is the same as described in 6.11.1, "Establishing a FlashCopy pair" on page 204. To create a task that will use an inband command to invoke the task at the remote site, you must select the **Inband Command** option.

Here are the steps for creating an inband FlashCopy:

1. In the volume panel, select the source volume (with a right-click) and target volume (with a left-click) for your FlashCopy relationship. Then right-click the target volume again to open the Task Wizard and select the establish FlashCopy pair task (Figure 6-72).

2. When you run the task, the source volume must be the secondary volume of a PPRC or PPRC-XD pair (Figure 6-72).
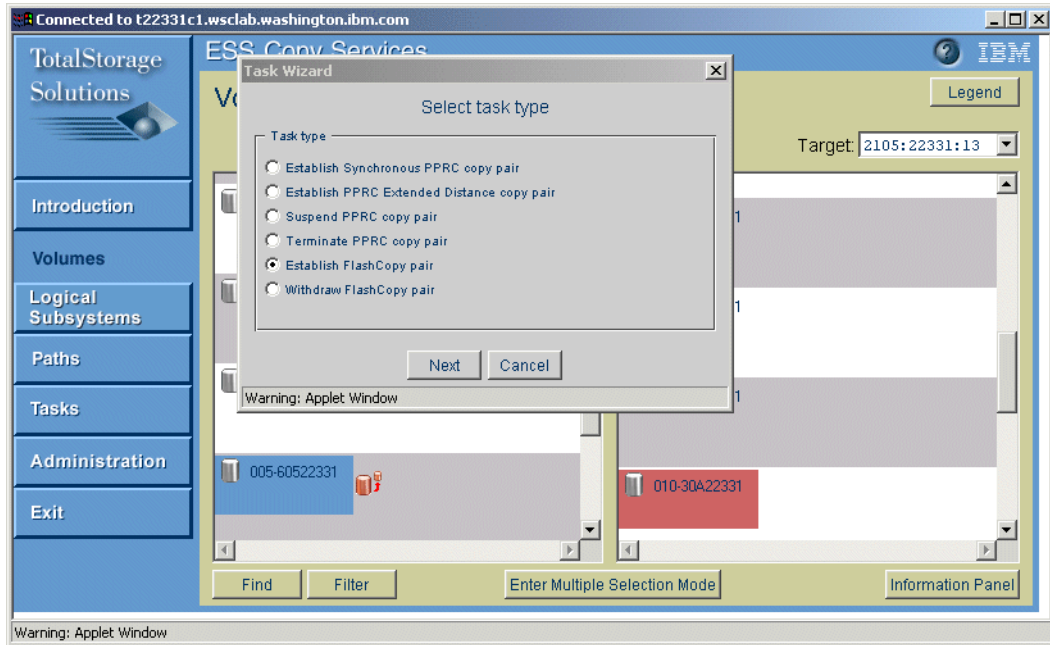
Figure 6-72   Volumes panel: Volume selection for a FlashCopy with Inband command option

3. Select the **Establish FlashCopy pair** option and click **Next**. Within the next panel, you select the **Inband Command** option and the other options for your FlashCopy relationship (in our example,Figure 6-73, we also select the **No Copy** option).
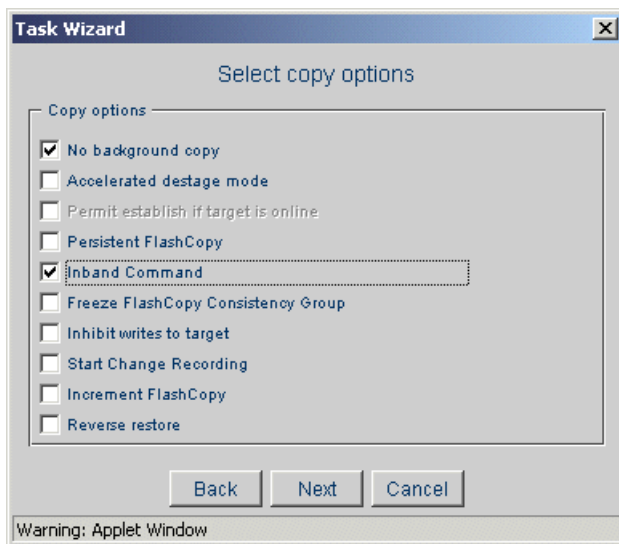
Figure 6-73   Task Wizard: Select copy options

4. Click the **Next** button and choose from which LSS you will issue the inband command (Figure 6-74). The LSS you select is on the Local site. It has to be the LSS where the primary volume of the PPRC relationship resides.
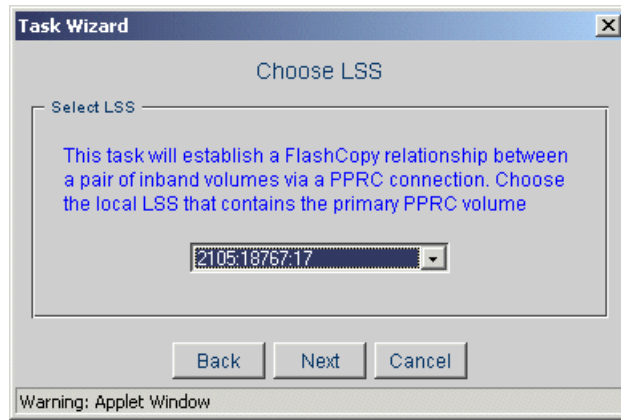
Figure 6-74   Task Wizard: Choose LSS

5. Click the **Next** button and decide whether to **Save**, **Run**, or **Cancel** the task (Figure 6-75).
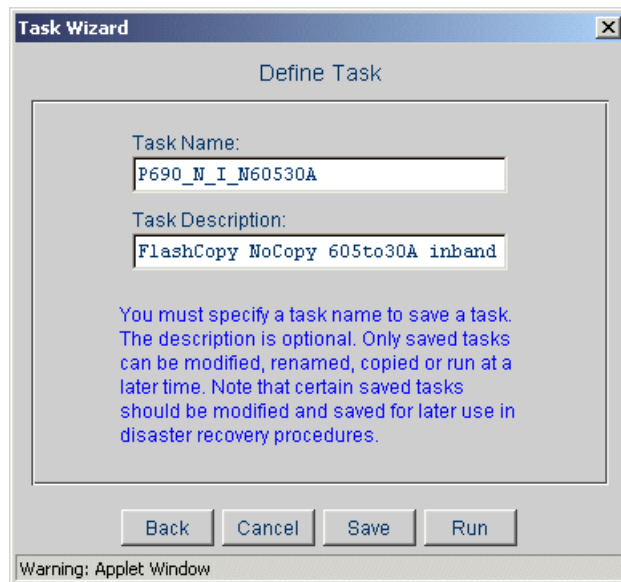
Figure 6-75   Task Wizard: Define Task

6. Once you run the task, the active server for Copy Services will issue the command to the LSS selected when creating the task. The order to run the task will be sent to the remote site from this LSS using an active path to the remote ESS. After the FlashCopy is established, the lightning bolts will appear on the source and the target volumes (Figure 6-76 on page 232).

> **Important:** You must have a PPRC relationship running between the primary and secondary site. The FlashCopy with inband command will only run successfully if the source volume for the FlashCopy is the secondary (target) volume of an established PPRC session (it can be suspended, as long as the links are available). The FlashCopy inband command will not run successfully if the source volume is in simplex state.

Figure 6-76   Volumes panel: Result of a FlashCopy with Inband Command option

6.12 PPRC setup examples using the Web user interface

In this section, we explain how to set up and configure different PPRC environments using the ESS Copy Services Web user interface. The Web user interface panels illustrated in these sections are the panels that a user at a minimum LIC 2.3.0 will see.

In general, there are two steps needed to successfully establish PPRC relationships:

► Setting up paths between the PPRC source and target LSSs

► Establishing the PPRC pairs, either on single volumes or on entire logical subsystems

> **Note:** To set up all PPRC pair related functions, either the Volumes panel or the Logical Subsystems panel (all volumes in the source and the target LSS) can be used. In this book, only the usage of the Volume panel is explained. For further information about how to set up a PPRC pair related function on the Logical Subsystem panel, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

Be aware of these requirements of the PPRC functionality:

► Paths for PPRC have to be available and need to be defined and established first.

► There are some rules when establishing paths (see 4.7, "PPRC connectivity" on page 82).

► The target volume has to be the same size as the source or larger.

6.12.1 Establishing paths

Before you can establish any PPRC pairs, you first have to set up the paths between the source and the target LSS where the primary and secondary volume reside. The paths are needed for communication between the PPRC pairs and to copy data from the source to the target.

Use the Paths panel of the ESS Copy Services Web use interface to set up paths for PPRC and select **Work with Paths**, as shown in Figure 6-77. Then follow this sequence:

1. Select the source ESS. The LSSs are then displayed

2. Select an LSS. Once the LSS is selected, you can then select the target ESS.

3. From the drop-down menu of the **Logical Subsystem Target** column, select the target ESS.

4. Select the target LSS. Notice that you can now select ESCON or FCP.

5. Select FCP or ESCON from the **Local Port/Remote Port** column.

6. Click on the first FCP or ESCON adapter and right-click on the subsequent adapters.

Note: Multiple adapters could be selected with a right-click after the first adapter was selected. If you choose the wrong adapter(s), just click on the correct adapter again to delete the selection.
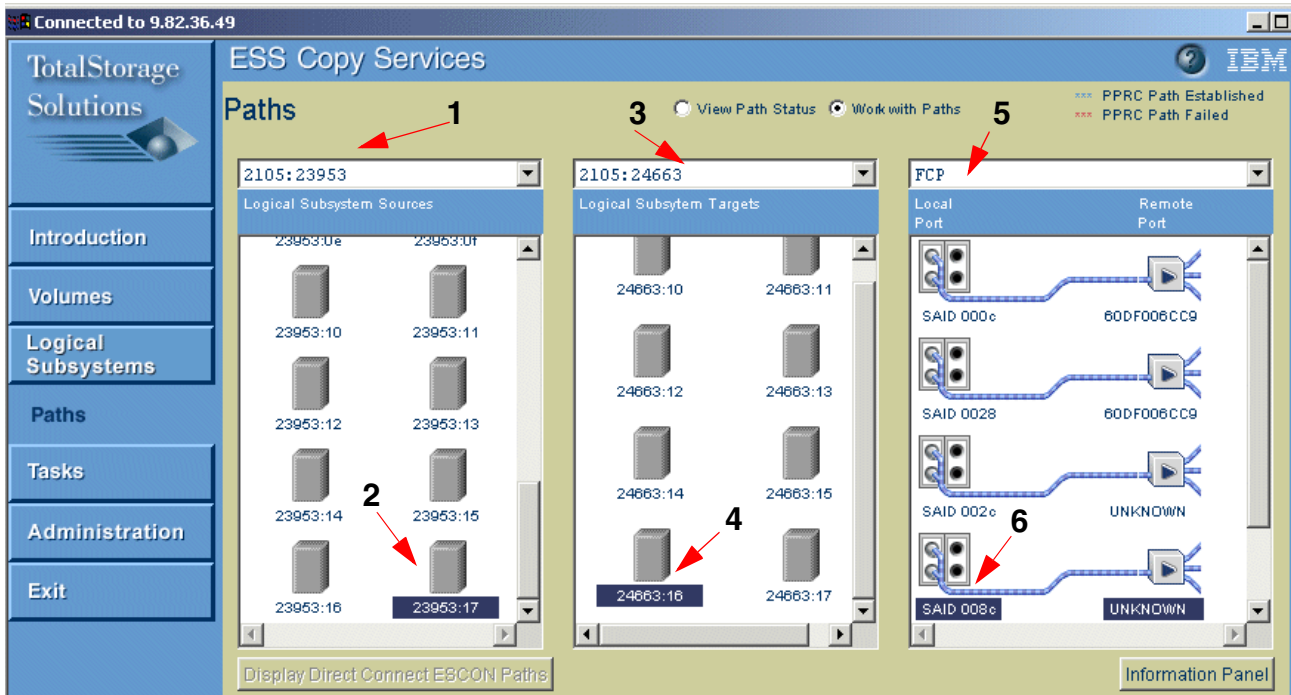
Figure 6-77   Establishing paths

The ESSs that are connected are automatically displayed in the **Logical Subsystem Targets** area.

Note: If you have some ESSs that are directly connected (that is, no switch between them), you can click **Display Direct Connect Paths** to display those paths.

You must perform a refresh function (see 6.9.2, "Refreshing path and volume information on an ESS" on page 202) for the following conditions:

▶ If this is the first time you are accessing ESS Copy Services
▶ After an update of the ESS microcode
▶ After a concurrent service action
▶ If the cabling for the PPRC links is modified

In the example shown in Figure 6-77 on page 233, we will establish an FCP path between the source LSS 17 in ESS 23953 and the target LSS 16 in ESS 24663.

Once the target and source of the PPRC path have been selected, right-click one of the highlighted adapters to bring up the Task Wizard (Figure 6-78). Select the **Establish Path** option and click **Next**.
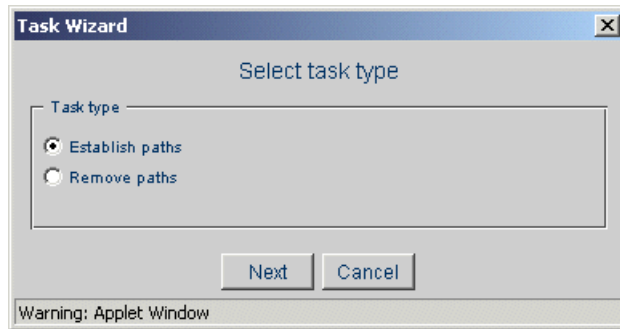
Figure 6-78   Task Wizard: Select task type

If a switch is part of the path, the Select outgoing ports panel of the Task Wizard opens, as shown in Figure 6-79, when using ESCON links.
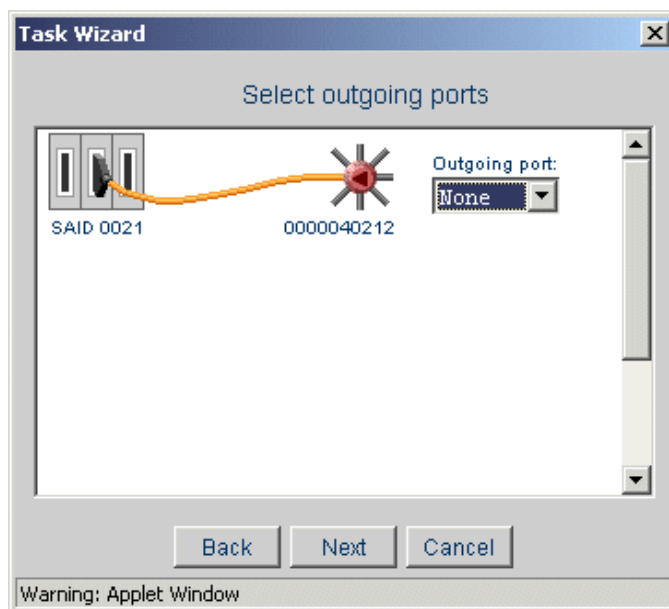
Figure 6-79   Select outgoing port: ESCON link

If you are working with Fibre Channel links and a switch is present, then the Select outgoing ports Task Wizard will be similar to Figure 6-80 on page 235.
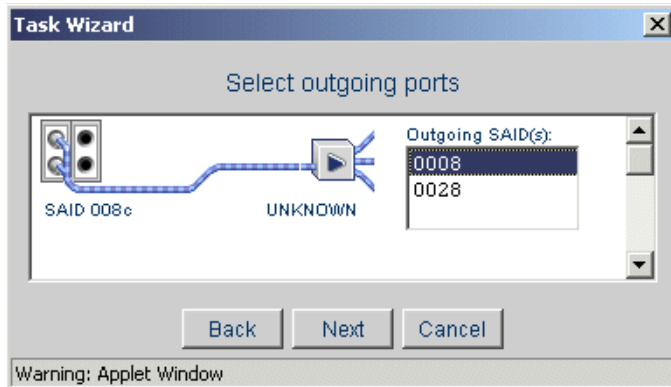
Figure 6-80   Select outgoing port: FCP link

In the outgoing ports list, select the desired port on the switch. Click **Next**. In the next panel you specify the path options, as Figure 6-81 illustrates, when working with ESCON links:

▸ **Do not establish paths if they already exist:** If this option is checked and there is already a path defined from the source to the target, the operation of establishing paths will not be executed. Does not apply to FCP paths.

▸ **Force removal of existing paths:** You must select this option if the selected path goes through a switch. The selection causes any logical paths between a host system and the selected ESCON adapter to be removed before establishing logical paths. Does not apply to FCP paths.

> **Attention:** The **Force removal of existing paths** option removes all logical paths from the host to the port. This is potentially disruptive to host activity, since there may be applications using the port. By host command under operator control, you should vary offline all logical paths from the host to the port. Only as a last solution or when the operator is absolutely sure there would be no negative impact should the **Force removal of existing paths** option be used.
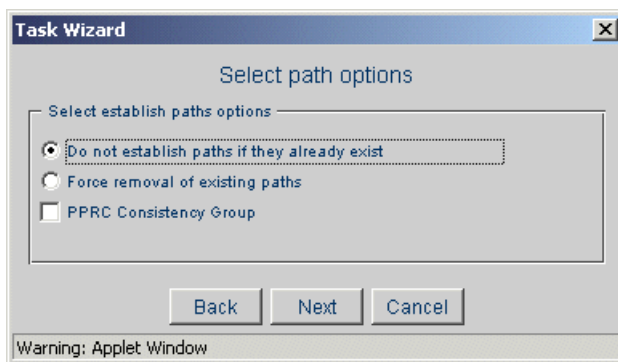
Figure 6-81   Select path options: ESCON links

If working with FCP links, the paths options you will see are shown in Figure 6-82 on page 236.
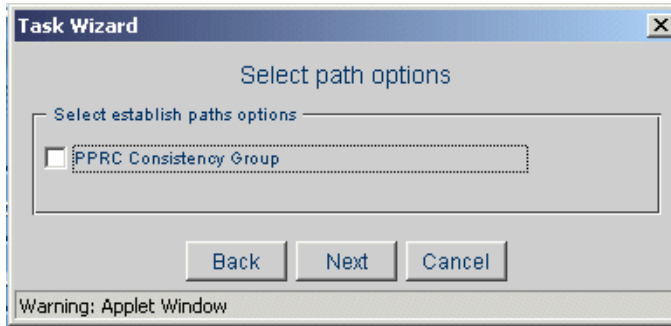
Figure 6-82   Select path options: FCP links

In either case, ESCON and FCP, you can also select **PPRC consistency group** to create a consistency group. For information on consistency groups, see 4.5, "PPRC data consistency" on page 72.

Click **Next** and decide whether to **Save**, **Run**, or **Cancel** the task. Click **Save** if you want to save the task in the task repository. If you save the task, it can be executed form the **Tasks** panel any time.
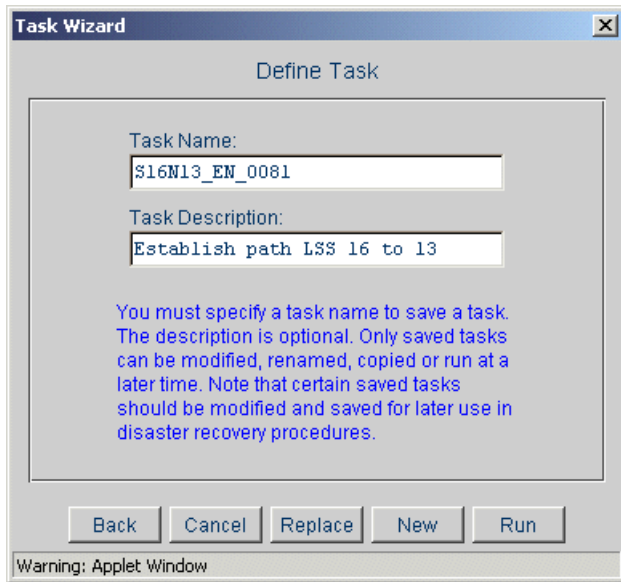
Figure 6-83   Task Wizard: Define Task

Tips:

► If you save the task, it can be executed from the Tasks panel at any time (see 5.7.3, "Running a task and viewing an error about a failed task" on page 126).

► After you create and save the task to establish paths, we strongly suggest that you immediately create the corresponding task to remove the path. You can modify the task within the Tasks panel so that it removes the path (see 6.8.5, "Modifying a task" on page 199). This way, you will save time when creating your tasks for disaster recovery.

When you run the task, and if the PPRC path is successfully established then three blue asterisks appear right below the adapter in the Local Port/Remote Port column of the Paths panel, as illustrated in Figure 6-84 on page 237, when using an ESCON link.
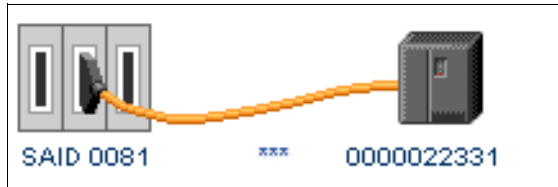
Figure 6-84   ESCON path successfully established

Figure 6-85 illustrates the display when a path is successfully establish using an FCP link.
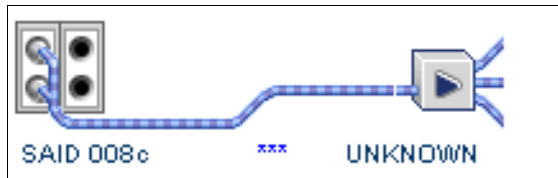
Figure 6-85   FCP path successfully established

Note: You can see information about the established path using the Information Panel. Refer to 6.7.5, "Path Information Panel button" on page 195.

6.12.2 Removing paths

You can remove paths using the Paths panel. The blue asterisks below the path-connection symbol in the Local Ports/Remote Ports column mean that a path is already established. This is how you remove a path (refer to Figure 6-86 on page 238):

1. Select **Work with Paths**.
2. From the drop-down menu, select the ESS source.
3. Click on the desired LSS source from the **Logical Subsystem Sources** panel.
4. From the drop-down menu, select the ESS target.
5. Click on the desired LSS target from the **Logical Subsystems Target** panel.
6. From the drop-down menu in the adapter column, select either ESCON or FCP.
7. Click on the path to be removed (it should have ***).
8. Right-click on the path to remove.
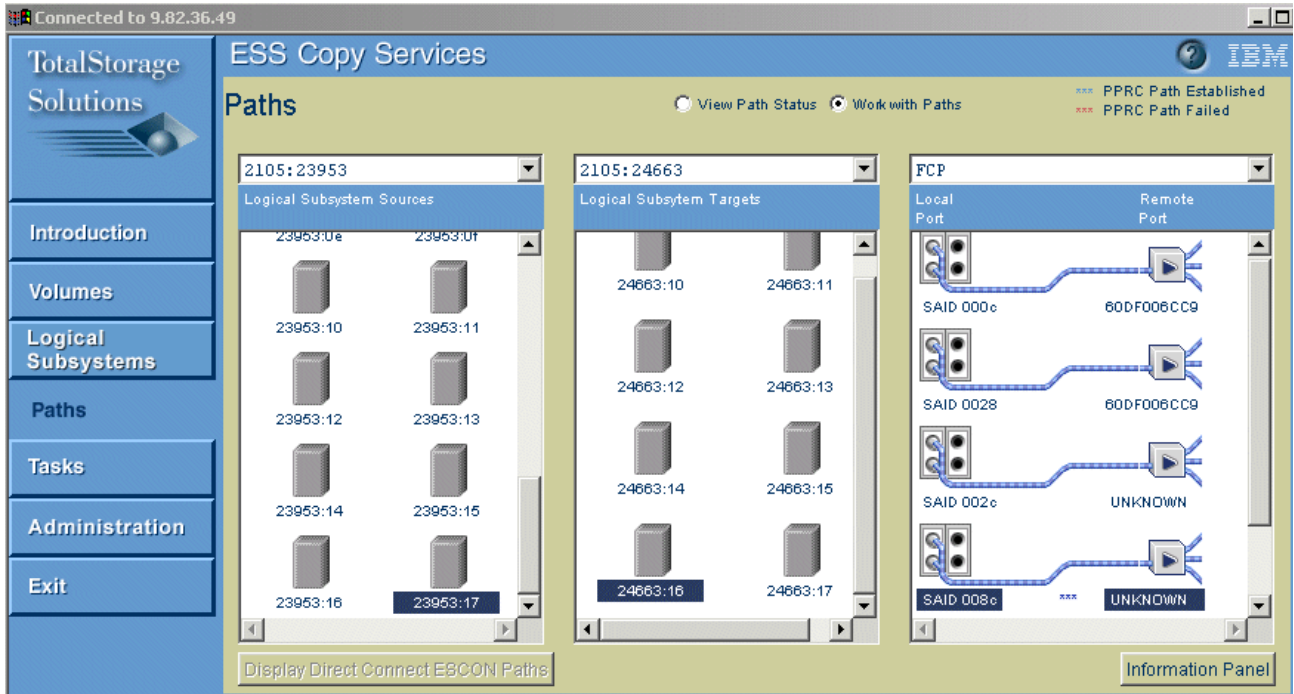
Figure 6-86   Selecting an established path to be removed

The Task Wizard opens (see Figure 6-87).
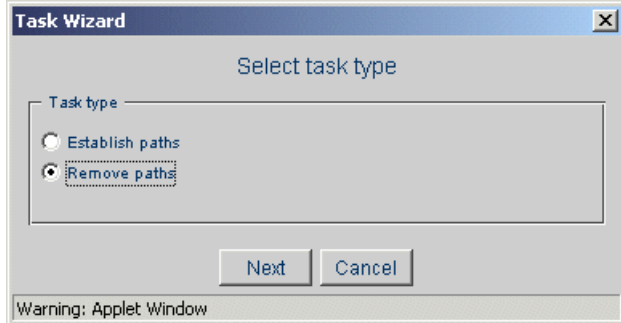
Figure 6-87   Task Wizard: Select task type

Select the **Remove paths** option and click **Next**. In the next panel, you can specify the task options (see Figure 6-88 on page 239).
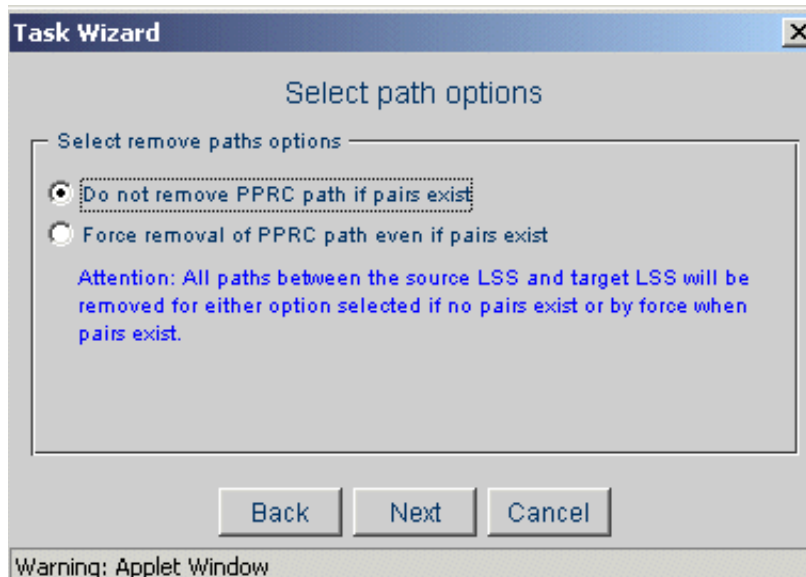
Figure 6-88   Task Wizard: Select path options

The options are:

►   **Do not remove PPPC path if pairs exists:** If you choose this option, the path will not be removed if there is a PPRC pair using this path.

►   **Force removal of PPRC path even if pairs exists:** This option enables you to force the removal of existing PPRC paths even if there are PPRC pairs established that could use this path. This option is helpful for the Failover/Failback Mode. For further information, refer to 4.7, "PPRC connectivity" on page 82.

From the next panel, you can either **Save**, **Run**, or **Cancel** the task. As mentioned before, once a task is saved, you can run it from the Tasks panel at any time.

Once you have run the task, all PPRC paths from the source to the target LSSs will be removed and the blue asterisks below the Path Connection symbol in the Path Origin column will disappear.

> **Tip:** To remove a path, you can either use the Paths panel, the Tasks panel (by modifying an Establish Paths task or by running a Remove Paths task), or by using the **Logical Subsystems** panel. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.12.3 Establishing a synchronous PPRC pair

Use the Volumes panel to establish PPRC pairs. On the left side, select the source LSS, and on the right side, the target LSS. This is done using the drop-down menus at the top of the Volumes menu.

The source and target logical subsystem are specified the following way: device type (4 digits): ESS Serial number (5 digits): LSS number (2 digits).

In our example shown in Figure 6-89 on page 240, we have selected 2105:18767:16 as the source LSS and 2105:22331:13 as the LSS target of our PPRC pair. The volume 000 in LSS 16 (serial number of 60018767) is selected as a primary (source) volume. The volume 009 in LSS 13 (Serial number of 30922331) is selected as a secondary (target) volume.
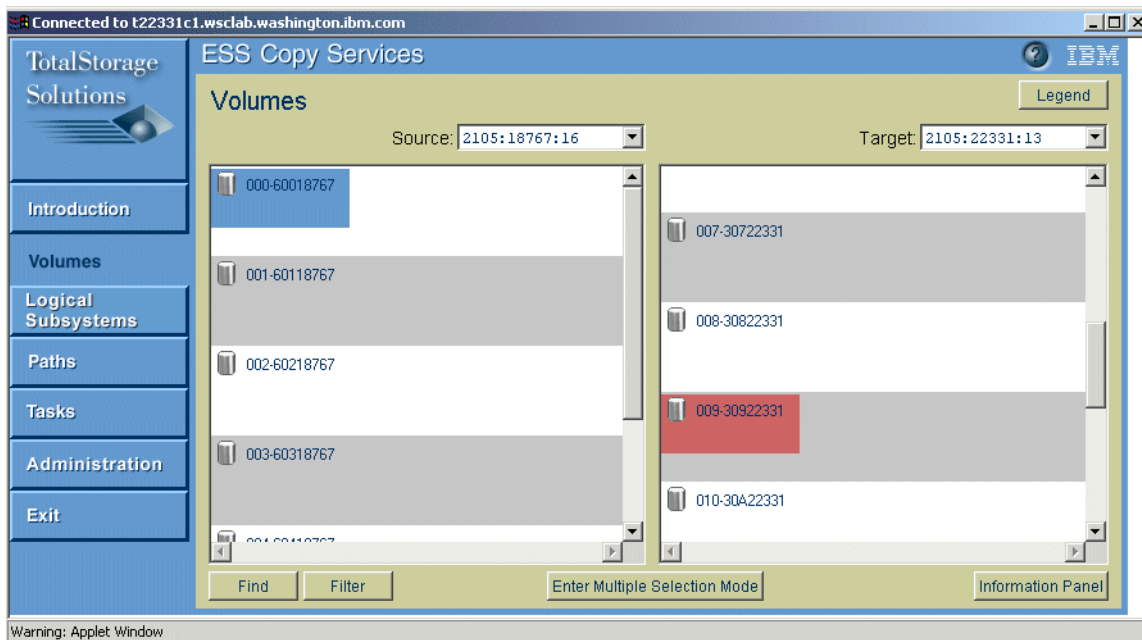
Figure 6-89   Volumes panel: Selecting source and target volumes for PPRC

You always need to have two components to establish a PPRC pair, a primary volume (source volume) and a secondary volume (target volume). Select the primary, and, with a right-click, the secondary. If you have selected the wrong volumes, just click on the correct primary volume again to clear the selection (Figure 6-89).

Once you have selected the primary and the secondary, do a second right-click on the secondary (target) to bring up the Task Wizard (Figure 6-90). Select the **Establish PPRC copy pair** option and click **Next**.
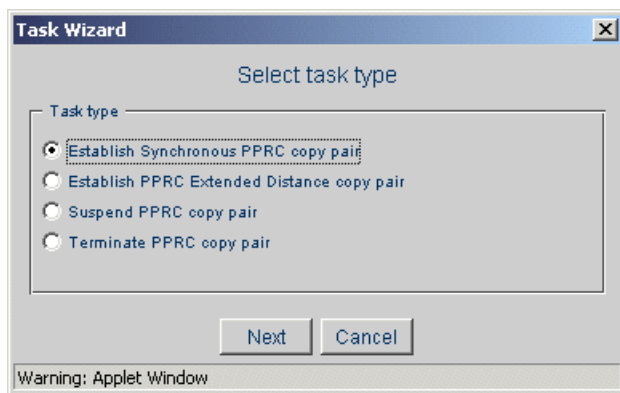
Figure 6-90   Task Wizard: Select task type

Within the next panel, you can specify the copy options of the PPRC pair (Figure 6-91 on page 241). Click **Next** when you have finished the selection.
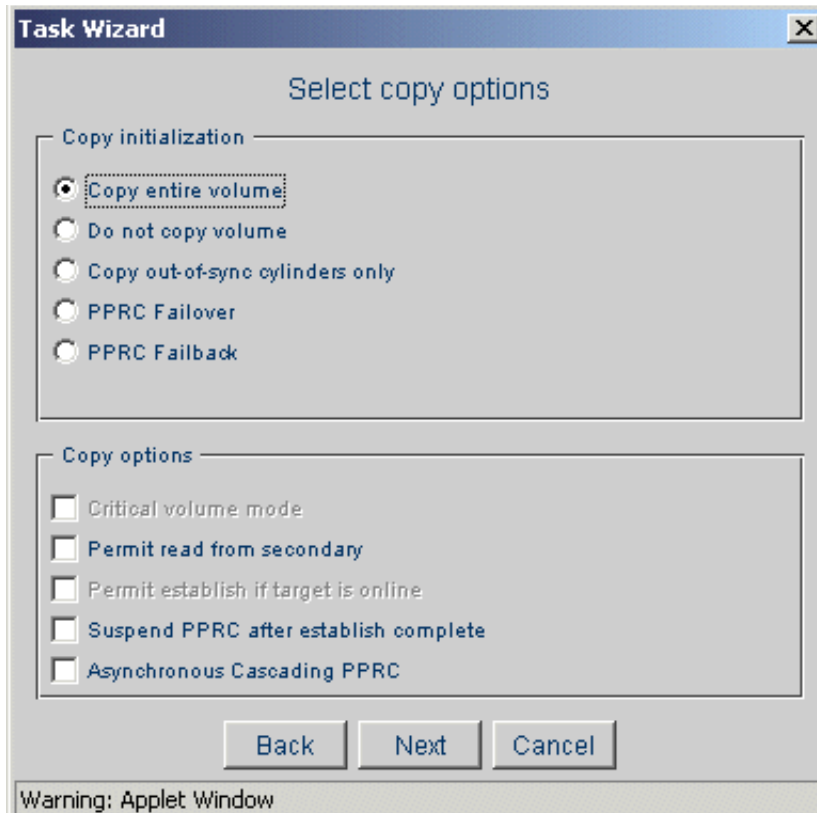
Figure 6-91 Task Wizard: Select Copy options

Options for synchronous PPRC

► Copy initialization:

- **Do not copy volume:** If this option is checked, the PPRC pair relationship is established without copying any data from the source to the target. This option is used when source and target contain exactly the same data and are consistent while in simplex state. Choosing this option causes the volumes to turn duplex immediately on the assumption that both source and target contain identical data.

- **Copy entire volume:** If this option is checked, all data is copied from the source to the target volume. This option has to be used the first time a PPRC relationship is going to be established and is needed to guarantee that source and target contain the same data.

- **Copy out-of-sync cylinders only:** This option copies only the data that was updated on the source volume since a PPRC copy pair was suspended. The option is used to resynchronize a PPRC synchronous or a PPRC-XD pair.

► PPRC modes:

- **PPRC Failover:** Select this option to reverse the direction of a PPRC pair. Use this option to make your current secondary volume at the recovery site become a primary volume (your current primary volume at the production side being the secondary volume of the PPRC relationship). For further information, refer to 6.12.13, "Performing a PPRC Failover" on page 266.

- **PPRC Failback:** Select this option to resynchronize the PPRC volumes at the recovery site with their peer volumes at the production site. For further information, refer to 6.12.14, "Performing a PPRC Failback" on page 269.

► Copy options:

– **Permit read from secondary:** Selecting this option allows host servers to read from the PPRC secondary volume. The PPRC pair must be in a full-duplex state in order for the host server to read the volume. This option is helpful for operating systems that use a Logical Volume Manager that reads data structures from the header of a secondary volume to update internal databases, such as the ODM under AIX, or /etc/lvmtab under HP-UX. For further information, refer to Chapter 8, "Open systems specifics" on page 327.

– **Suspend PPRC after establish complete:** If you select the **Copy entire volume** or **Copy out-of-sync cylinders only** option described above, this means that after the data is completely copied to the target volume, the target volume goes into the suspended state. Selecting this option saves you the effort of creating a separate task to suspend the pair. This is useful in a PPRC-XD catch-up operation over extended distances.

– **Asynchronous Cascading PPRC:** Select this option to allow a PPRC secondary volume to also be a PPRC primary volume for a different relationship. You can use this option for creating a consistent copy of your data at a tertiary site. For further information, refer to 6.12.15, "Establishing an Asynchronous Cascading PPRC" on page 273.

From the next panel, you can either **Save**, **Run**, or **Cancel** the copy task, as shown in Figure 6-92. Click **Save** to run the task and run it later. To run the task immediately, click **Run**.
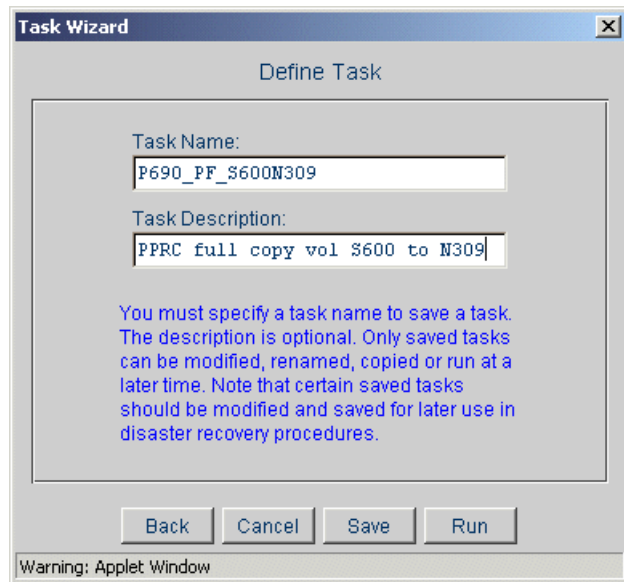
Figure 6-92   Task Wizard: Define Task

> **Tip:** After you create and save the task to establish synchronous PPRC copy pairs, we recommend that you create the tasks that suspend the pairs and terminate the pairs. You can save time in creating tasks if you use the task you just saved as a template to create the suspend task. Then use the suspend task as a template to create the terminate task. In the **Task panel**, select the task you created and click on **Modify** button to change the options of the task. Then, give it a new name and save it as a new task.

Once a synchronous PPRC pair is established, the display of the primary and secondary volumes from the Volumes panel changes, indicating that the initial copy is in progress (Figure 6-93).
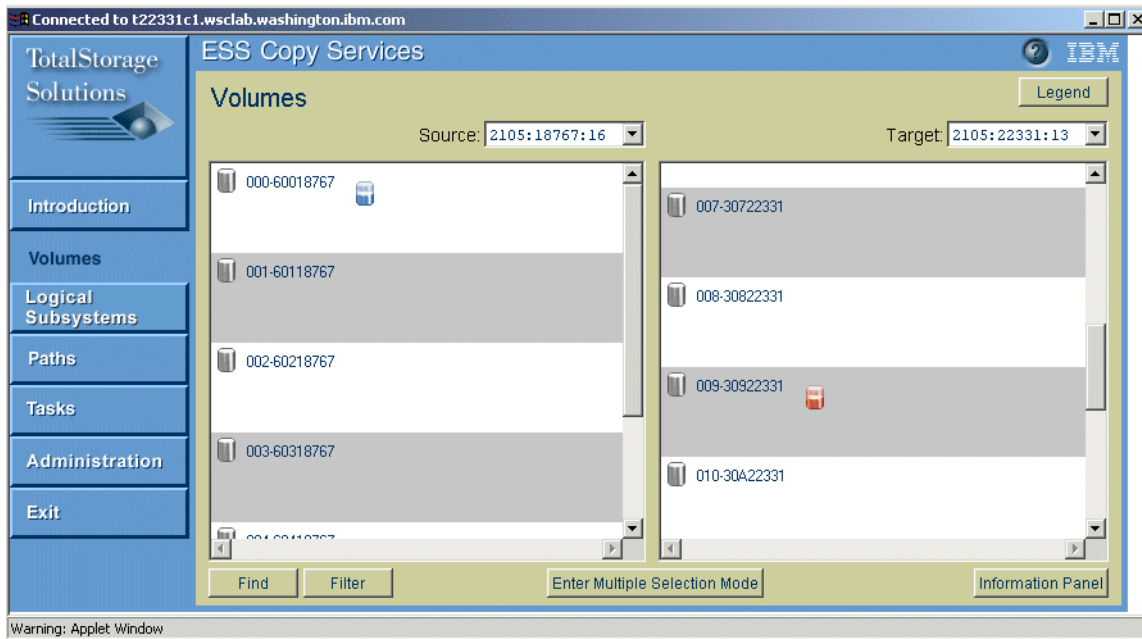
Figure 6-93   Two volumes in a PPRC relationship and in copy pending status

Select a volume and click the **Information Panel** button to retrieve more information about the status. If the source of a PPRC pair is selected, the number of out-of-sync cylinders that are still left to copy are displayed. Those are the tracks that need to be copied from the source to the target to achieve full copy mode (Figure 6-94).
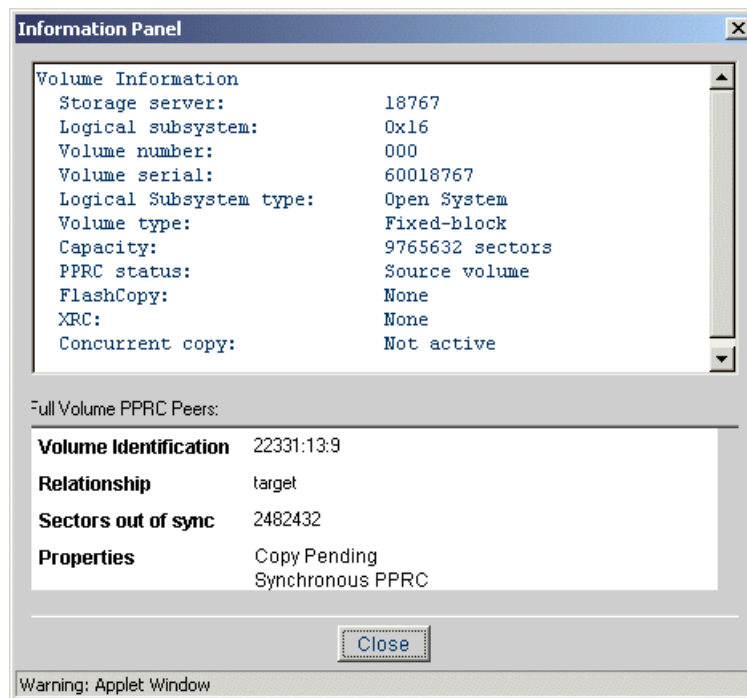
Figure 6-94   Information panel on a PPRC volume

Once the volumes are in sync, the icons of the volumes icons will change to solid. (Figure 6-95).
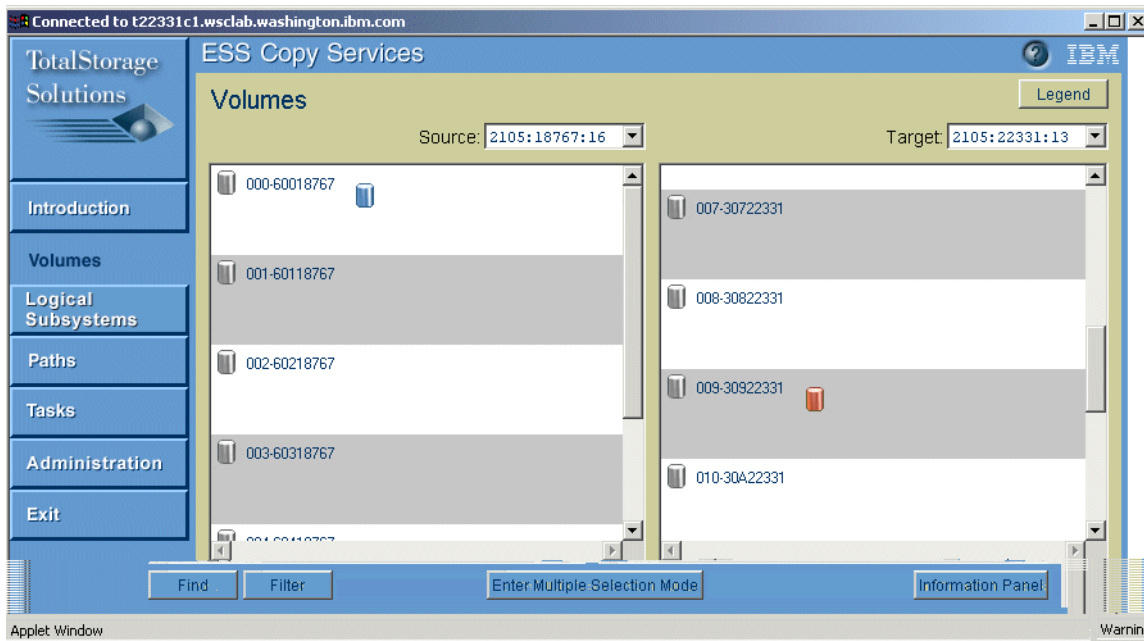
Figure 6-95   PPRC volume in full duplex mode

> **Tip:** Alternatively, you can use the Logical Subsystem panel to select **Establish Synchronous PPRC Pairs** for an entire LSS. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.12.4 Establish a PPRC Extended Distance copy pair

PPRC Extended Distance provides the following advantages over the standard synchronous PPRC mode:

► Minimal response-time delays for application hosts that are writing to a PPRC primary device

► Copy distances that can be increased by thousands of miles with channel extenders

► Data migration without the response-time penalty of the synchronous PPRC copy mode

► No FlashCopy of the primary volume required to eliminate the synchronous-write penalty

Refer to 4.3, "PPRC extended distance (PPRC-XD)" on page 59 to get more information about how PPRC-XD works and what benefits you can get from it in your production environment.

To establish the PPRC Extended Distance copy pair, use the Volumes panel. This procedure is almost identical to the previous one, which you use to create a synchronous PPRC copy pair, except that you select the **Establish Extended Distance PPRC** option in the Select task type panel of the Task Wizard (Figure 6-96 on page 245).
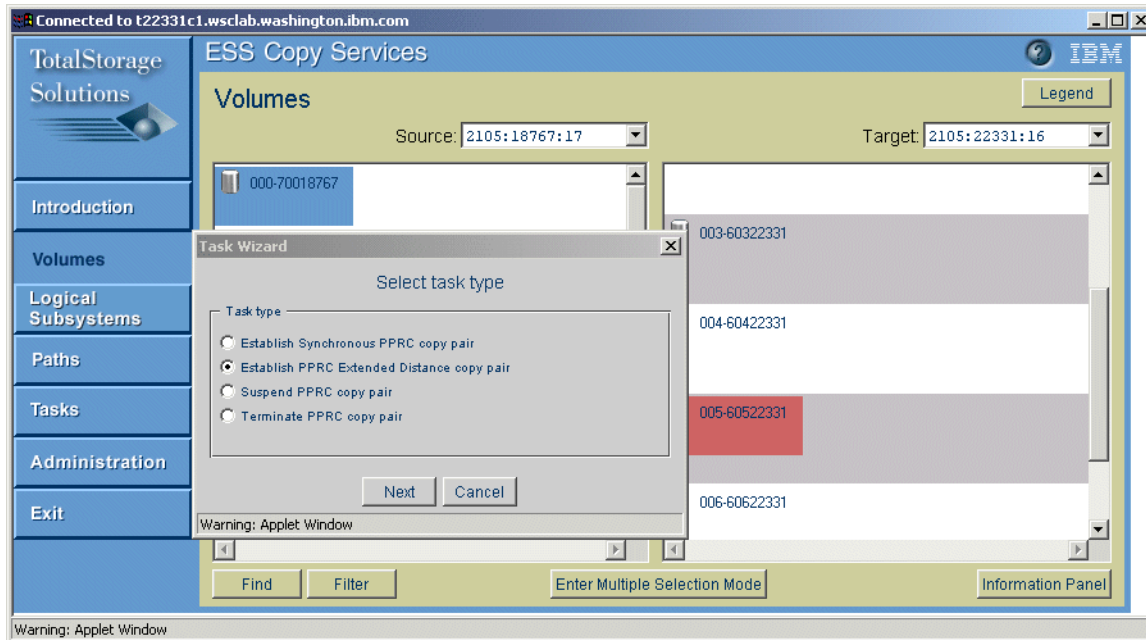
Figure 6-96   Establishing a PPRC-XD between two volumes

When the option is selected, click **Next** to select the options for this task (Figure 6-97).
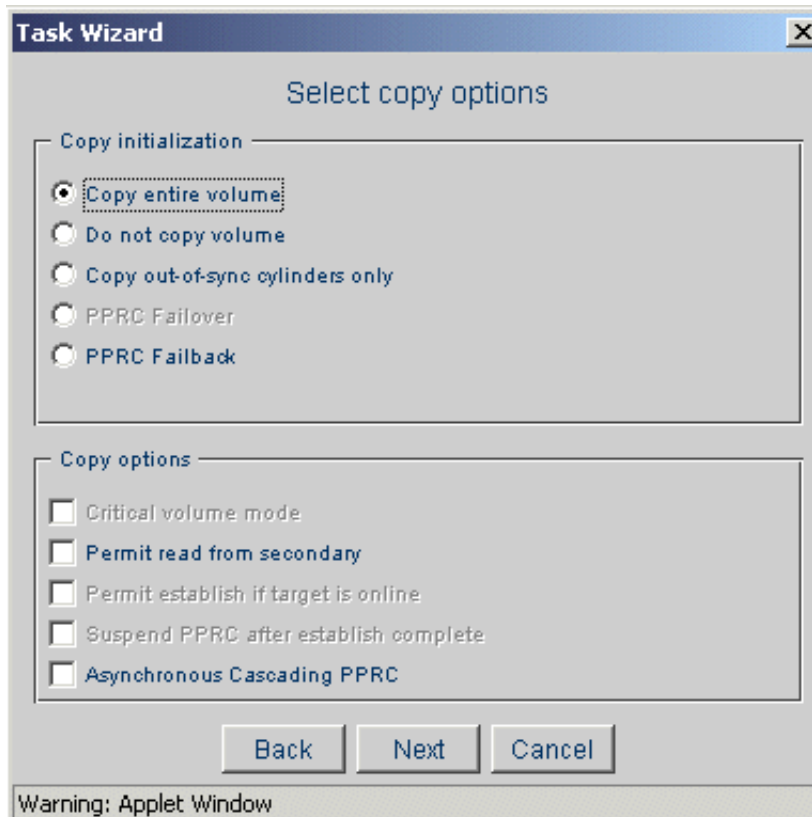
Figure 6-97   Task Wizard: Select copy option

Copy options for PPRC-XD

Here is a short description of the available options for PPRC-XD. To get a more detailed description of these options, refer to 4.3.2, "PPRC-XD tasks and options" on page 60.

- ► Copy initialization:
  - **Copy entire volume:** If this option is checked, the two volumes will be in a PPRC-XD duplex pending mode, considering that all the data from the primary volume need to be sent to the secondary volume. This option is the default option. It will be used to initiate the PPRC-XD relationship.

  - **Do not copy volume:** If this option is checked, when the PPRC-XD pair relationship is established, the primary and secondary volumes will be in PPRC-XD duplex pending mode, but the ESS will consider that no data need to be sent to the secondary. Only new updates to the primary volumes will be sent to the secondary volume.This option is used when the source and target contain exactly the same data and are consistent while in simplex state, or, when establishing an Asynchronous Cascading PPRC to avoid transferring unnecessary data (see 6.12.15, "Establishing an Asynchronous Cascading PPRC" on page 273).

  - **Copy out-of-sync cylinders only:** This option copies only the data that was updated on the target volume since a PPRC-XD copy pair was suspended. The option is used to resynchronize a PPRC-XD pair.

- ► PPRC modes:
  - **PPRC Failover:** This option is not available with PPRC-XD.

  - **PPRC Failback:** Select this option to resynchronize the PPRC volumes at the recovery site with their peer volumes at the production site. For further information, refer to 10.4, "Types of PPRC Failover and Failback procedures" on page 405.

- ► Copy options:
  - **Permit read from secondary:** Selecting this option allows host servers to read from the PPRC secondary volume. The PPRC-XD pair does not need to be in a full-duplex state in order for the host server to read the volume. This option is helpful for operating systems that use a Logical Volume Manager that reads data structures from the header of a secondary volume to update internal databases, such as the ODM under AIX, or /etc/lvmtab under HP-UX. For further information, refer to Chapter 8, "Open systems specifics" on page 327.

  - **Suspend PPRC after establish complete:** This option is not available for PPRC Extended Distance. The reason this requirement is not imposed is that the duplex state is not enforced by PPRC Extended Distance.

  - **Asynchronous Cascading PPRC:** Select this option to allow a PPRC secondary volume to also be a PPRC primary volume for a different relationship. You can use this option for creating a consistent copy of your data at a tertiary site.

From the next panel, you can either **Save**, **Run**, or **Cancel** the copy task, as shown in Figure 6-98 on page 247. Click **Save** to run the task and run it later. To run the task immediately, click **Run**.
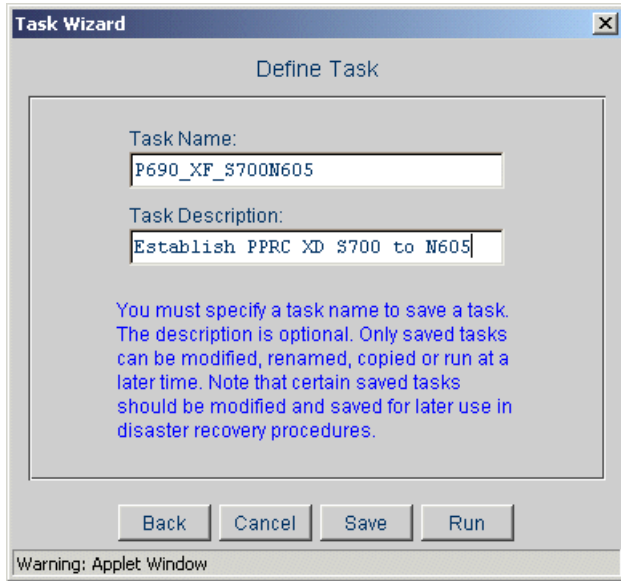
Figure 6-98   Task Wizard: Define Task

Once a synchronous PPRC Extended Distance pair is established, the PPRC-XD icons will appear next to the source and target volumes (Figure 6-99).
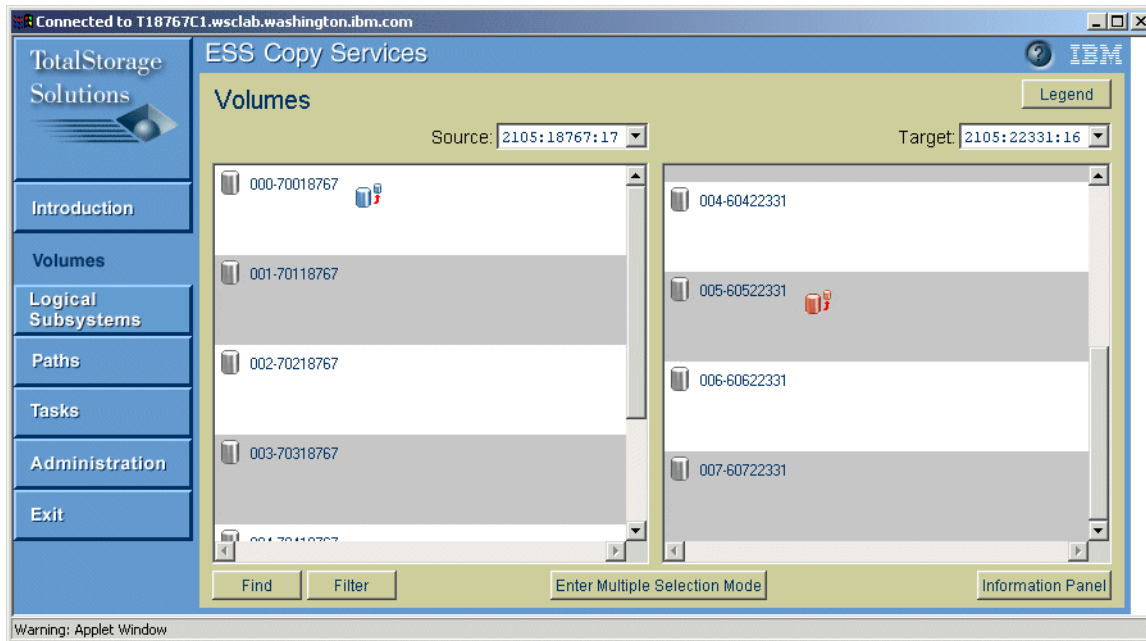
Figure 6-99   Volumes panel: PPRC-XD established between volumes 700 and 605

Tip: Alternatively, you can use the Logical Subsystem panel to select Extended Distance PPRC for an entire LSS. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.12.5 Converting a PPRC-XD copy pair to synchronous PPRC

There are two common situations when you would convert a Peer-to-Peer Remote Copy (PPRC) copy pair from PPRC Extended Distance (PPRC-XD) mode to synchronous PPRC:

► Situation 1: You have used PPRC-XD to complete the bulk transfer of data in the creation of many copy pairs, and you now want to convert some or all of those pairs to synchronous PPRC mode.

► Situation 2: You have PPRC-XD copy pairs for which you want to make FlashCopy backups on the remote site. You convert the pairs temporarily to synchronous PPRC mode in order to obtain a point-in-time consistent copy.

For an in-depth discussion of the use of PPRC Extended Distance, see the redbook *IBM TotalStorage Enterprise Storage Server: PPRC Extended Distance*, SG24-6568.

You can convert a PPRC Extended Distance copy pair to a synchronous pair using the **Volume** panel. Select the LSS within which you want to convert the PPRC Extended Distance copy pair to a synchronous pair. Click on the primary PPRC volume. In the Target column of the Volumes panel, the secondary Volume of the PPRC pair will be displayed (Figure 6-100).

Once you have selected the source volume, right-click on the target volume, then right-click it again to open the Task Wizard (Figure 6-100).
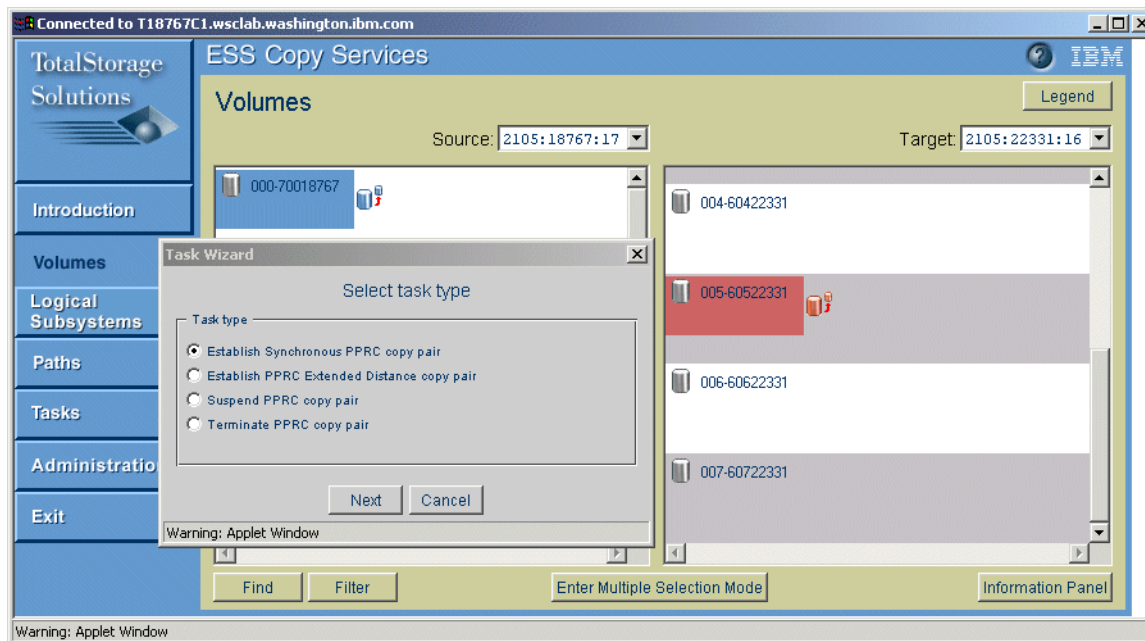
Figure 6-100   Volumes panel and Task Wizard

Select the **Establish synchronous PPRC copy pair** option and click **Next**. On the next panel, you can select the copy options (Figure 6-101 on page 249).
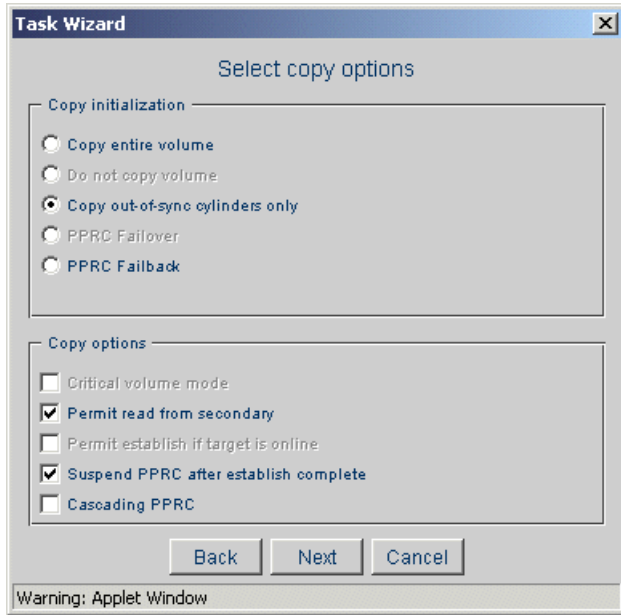
Figure 6-101   Task Wizard: Select copy options for PPRC-XD

Select **Copy out-of-sync cylinders only option.** Refer to 6.12.4, "Establish a PPRC Extended Distance copy pair" on page 244 to get details about these options. Other options can also be selected.
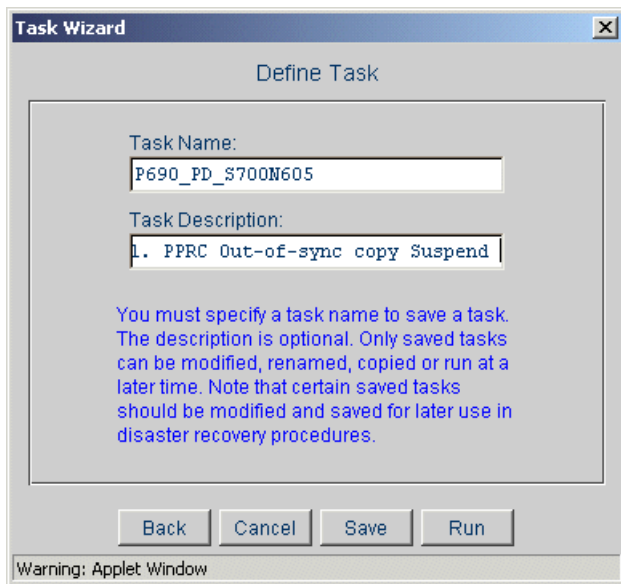
Figure 6-102   Task Wizard: Define Task

Click **Next**. From the next panel (Figure 6-102), you can either **Save**, **Run**, or **Cancel** the task.

Once you run the task, all out-of-sync cylinders will be copied from the primary to the secondary volume. Once the copy starts, the PPRC-XD icons will be removed and will change to copy-pending icons. The volumes will now be in "Copy pending" mode in a Synchronous PPRC relationship (Figure 6-103 on page 250).
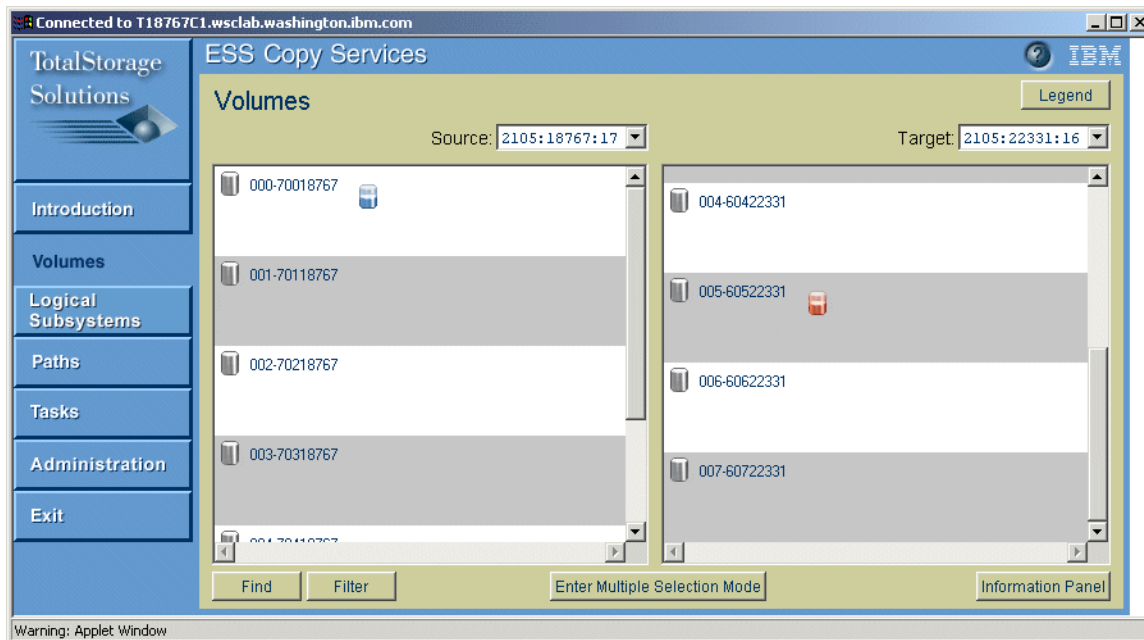
Figure 6-103   Volumes panel: Resynchronizing volumes

In our example, we selected the **Suspend after establish** option. So, at the end of the copy of all out-of-synchronous tracks, when the two volumes are in "Duplex" mode, they change immediately to the "suspended" mode. The icons will automatically change to show the suspended state when the copy of out-of-synch tracks is finished (Figure 6-104).
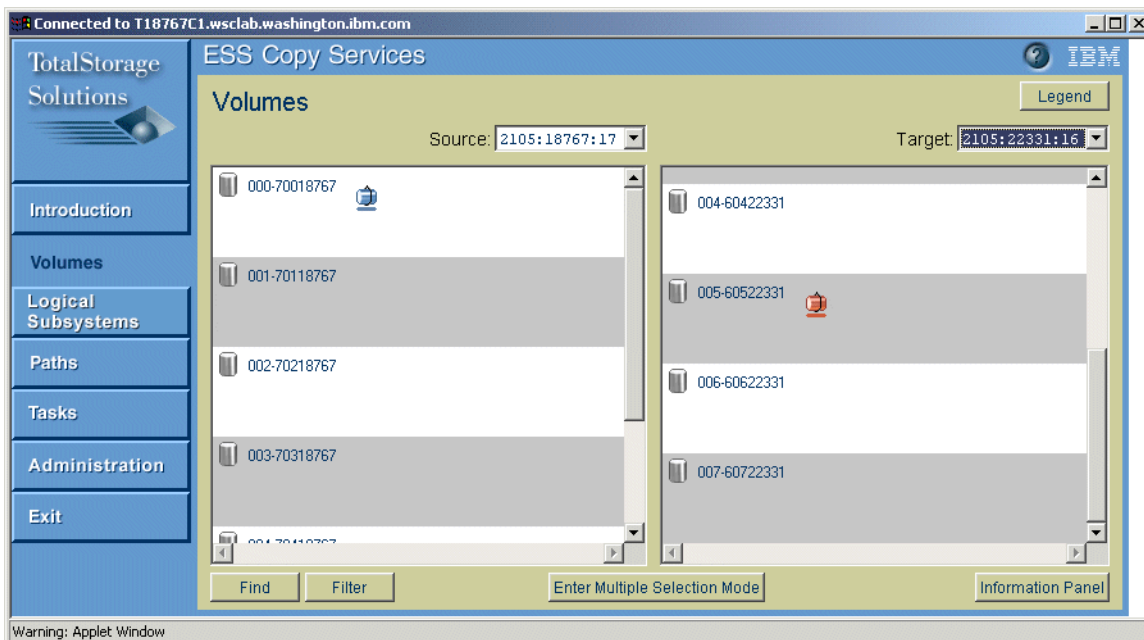
Figure 6-104   Volumes panel: Volumes are in suspended state after a fullcopy is established

Tip: Alternatively, you can use the Logical Subsystem Panel to select convert PPRC Extended Distance copy pairs to synchronous PPRC. Running the tasks at the LSS level will perform the task on several volumes. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.12.6 Creating a backup copy of volumes via PPRC Extended Distance

When you have used the PPRC Extended Distance (PPRC-XD) mode to transfer data, and the distance between the local ESS and the remote ESS is close enough for you to establish a synchronous PPRC copy pair that can achieve the full duplex state, you can temporarily convert from PPRC Extended Distance mode to synchronous PPRC mode when you want to make a consistent copy of the volumes.

This process would be useful as a backup copy and for data migration purposes, but would not be particularly useful for disaster recovery if the source data set is subject to rapid change. For further information, refer to 4.3, "PPRC extended distance (PPRC-XD)" on page 59.

Perform the following steps to obtain a backup point-in-time consistent copy of the PPRC-XD volumes that are at the remote site:

1. Change the copy mode for the PPRC Extended Distance copy pair to synchronous PPRC. Select the **Suspend after establish** option in the Task Wizard in creating the synchronous PPRC. This avoids the step of creating a separate task to just suspend the pairs. See 6.12.5, "Converting a PPRC-XD copy pair to synchronous PPRC" on page 248.

2. Alternatively, if you did not select the **Suspend after establish** option, suspend the synchronous PPRC copy pair after it reaches the full duplex state (consistency). See 6.12.7, "Suspending a PPRC pair" on page 251.

3. Perform a FlashCopy on the secondary volumes. See 6.11.1, "Establishing a FlashCopy pair" on page 204.

4. Resume PPRC Extended Distance mode for the copy pair. If you have saved the task that created the original PPRC-XD relationship (see 6.12.4, "Establish a PPRC Extended Distance copy pair" on page 244"), you can invoke the saved task from the Tasks panel.

6.12.7 Suspending a PPRC pair

When you suspend a copy pair, PPRC stops transferring data to the target volume. Because the primary ESS keeps track of all changed cylinders on the source volume, you can resume PPRC operations at a later time. To resume operations, run a task that resynchronizes the PPRC copy pair (using the **Copy out of sync cylinders only** option). See 6.12.3, "Establishing a synchronous PPRC pair" on page 239.

ESS Copy Services now provides an alternative to creating a separate task that suspends a PPRC pair. When you establish the pair, the Task Wizard provides the option **Suspend PPRC after Establish**. For details, see 6.12.3, "Establishing a synchronous PPRC pair" on page 239.

To suspend a PPRC copy pair, you use the Volumes panel. This procedure is almost identical to the one described in 6.12.3, "Establishing a synchronous PPRC pair" on page 239, except that you select the **Suspend PPRC copy pair** option in the Select task type panel of the Task Wizard (Figure 6-105 on page 252).

Figure 6-105   Task Wizard: Selecting Suspend PPRC copy pair option

If you click **Next** within the next panel, you can decide whether to schedule the task with the source volume or to schedule the task to the target volume (Figure 6-106).
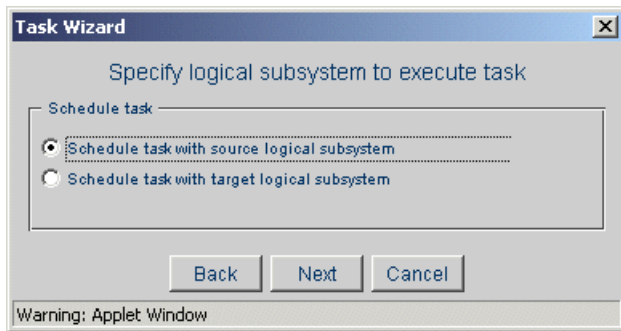
Figure 6-106   Task Wizard

The task runs on either the source or target ESS. For example, if the source LSS is not available, you can schedule a suspension of a PPRC pair with the target LSS. This is useful in a disaster situation when the primary site has gone down. Refer to Chapter 10, "Disaster recovery" on page 399 to get more details.

> **Note:**
>
> ► If you schedule the task with the source logical subsystem, the source will become suspended. As a result, the target will become suspended immediately after (Figure 6-107).
>
> ► If you schedule the task with the target logical subsystem. The target will become suspended. But the source might not become suspended. You may use this option only in a disaster situation when the source LSS is not available anymore. (An example of the result of this option is shown in Figure 6-108.)

The result of a suspend task with the **Schedule with the source** option selected is shown in Figure 6-107 on page 253.
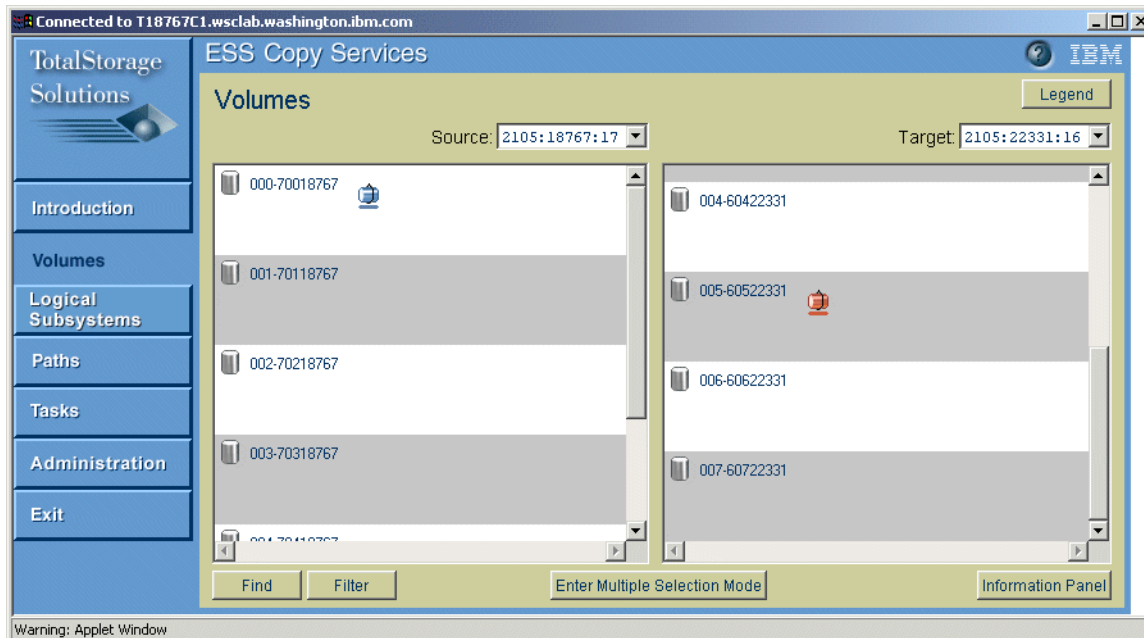
Figure 6-107   Volumes panel: Result of a suspend task with the Schedule with the source option

The result of a suspend task with the **Schedule with the source** option selected is shown in Figure 6-108.
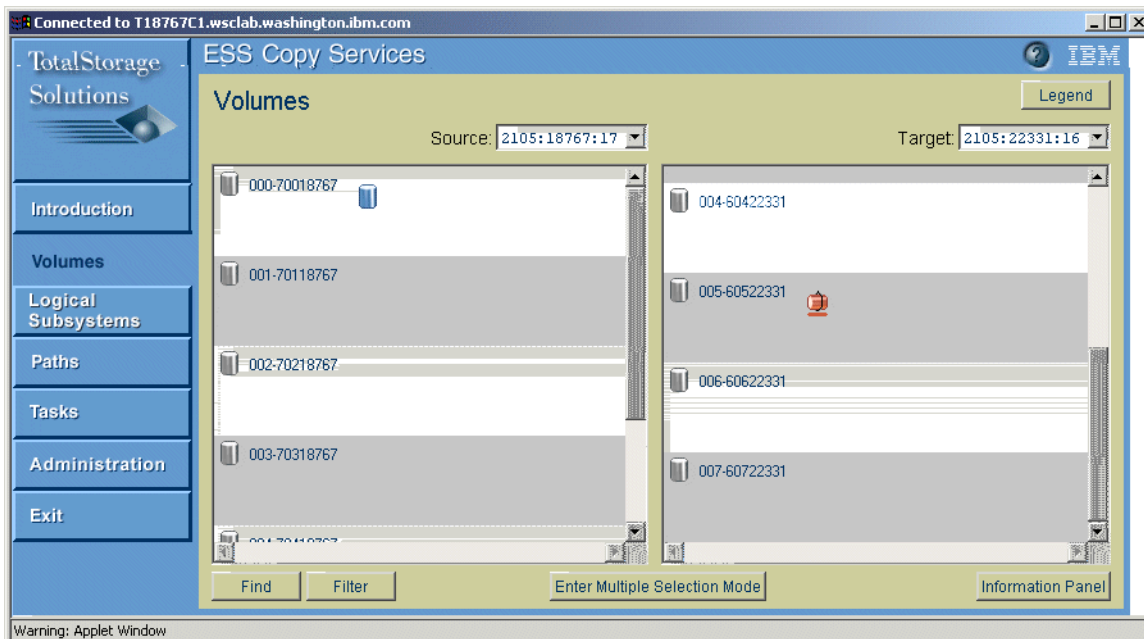
Figure 6-108   Volumes panel: Result of a suspend task with the Schedule with the target option

6.12.8 Terminating a PPRC copy pair

When you terminate a PPRC copy pair, the PPRC relationship between the source and target volume ends. The source and target volumes are removed from the PPRC configuration. You can either use this function to terminate a synchronous PPRC pair or an Extended Distance PPRC pair.

You can terminate a PPRC copy pair in the following ways:

► By modifying a task used to establish or suspend a PPRC copy pair that has been saved (refer to 6.8.5, "Modifying a task" on page 199 for further information)

► By using the Volumes panel or the LSS panel

This procedure is almost identical to 6.12.3, "Establishing a synchronous PPRC pair" on page 239, you first select the primary (source) volume and the secondary (target) volume, as shown in Figure 6-109.
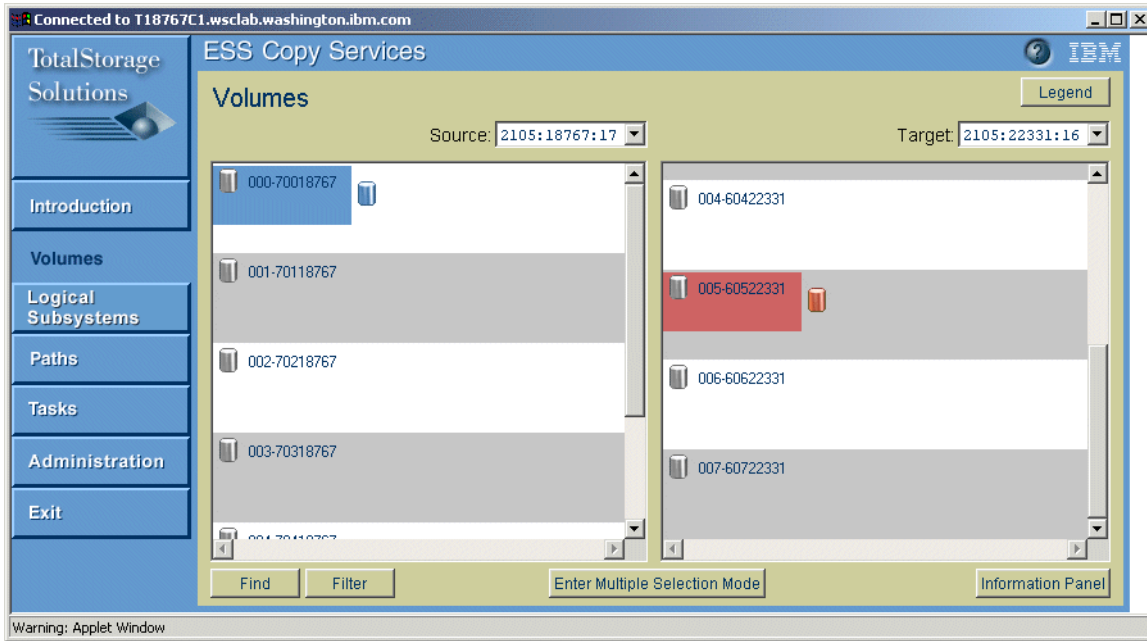
Figure 6-109   Volumes panel: Two volumes in a PPRC relationship selected

To terminate the pair, you select the **Terminate PPRC copy pair** option in the **Select task type** panel of the Task Wizard (Figure 6-110). Then click **Next.**
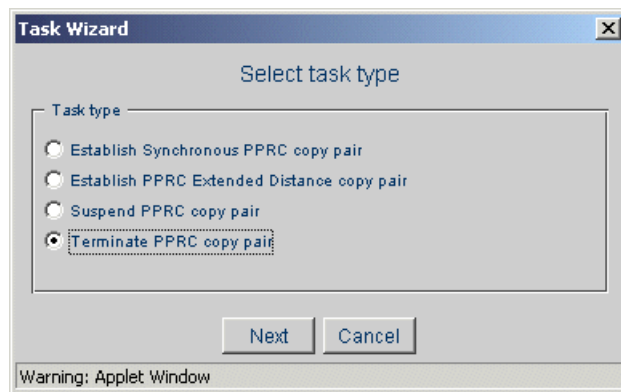
Figure 6-110   Task Wizard: Select task type

Within the next panel, you can decide whether to **Schedule task with the source logical subsystem** or to **Schedule task with the target logical subsystem** (Figure 6-111 on page 255).
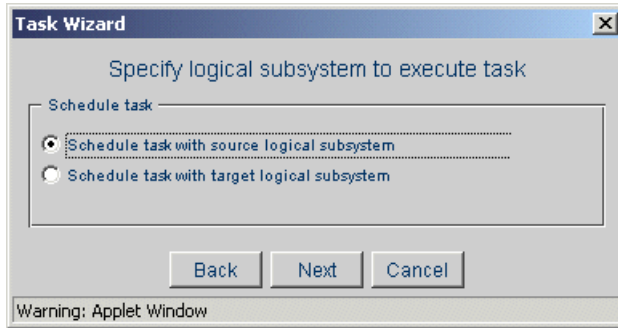
Figure 6-111   Task Wizard: Specify logical subsystem to execute the task

Note:

► If you schedule a **Terminate PPRC pair** task with the *source* logical subsystem, and you run the task successfully, the source and the target volume go to a simplex state (see Figure 6-112).

► If you schedule a **Terminate PPRC pair** task with the *target* logical subsystem, and the task has been completed successfully, the source volume is in suspended state, and the target volume is in the simplex state (see Figure 6-113). This option is useful in a disaster situation when the primary site has gone down.

The result of a terminate PPRC pair with the source Logical Subsystem selected is shown in Figure 6-112.
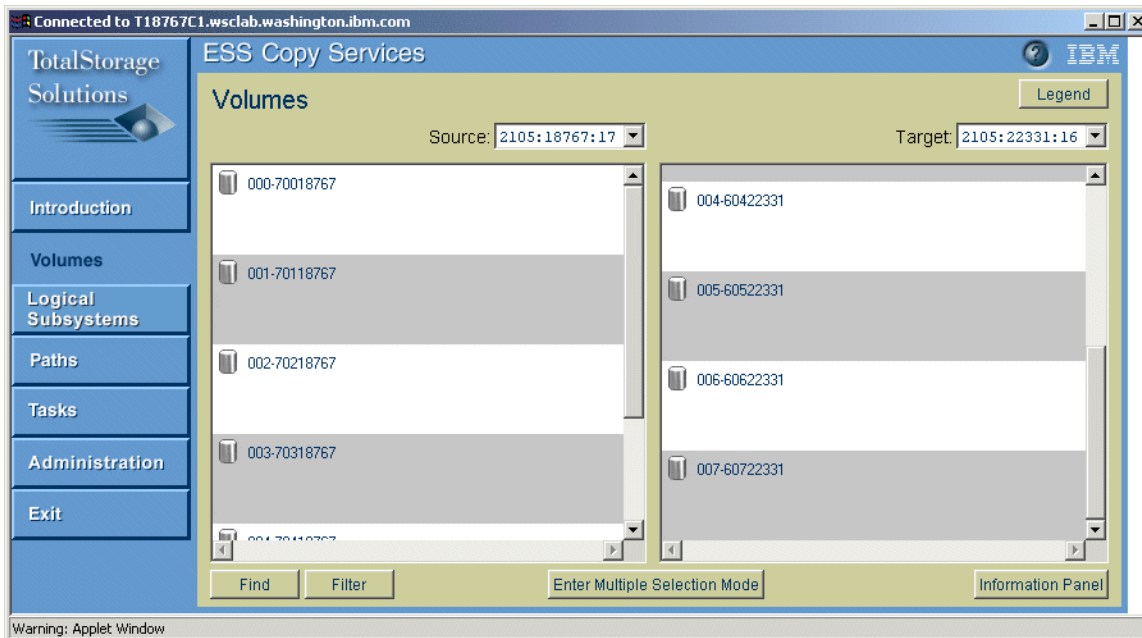
Figure 6-112   Volumes panel: Result of a terminate PPRC with Schedule with target option

The result of a terminate PPRC pair with the target Logical Subsystem selected is shown in Figure 6-113 on page 256.
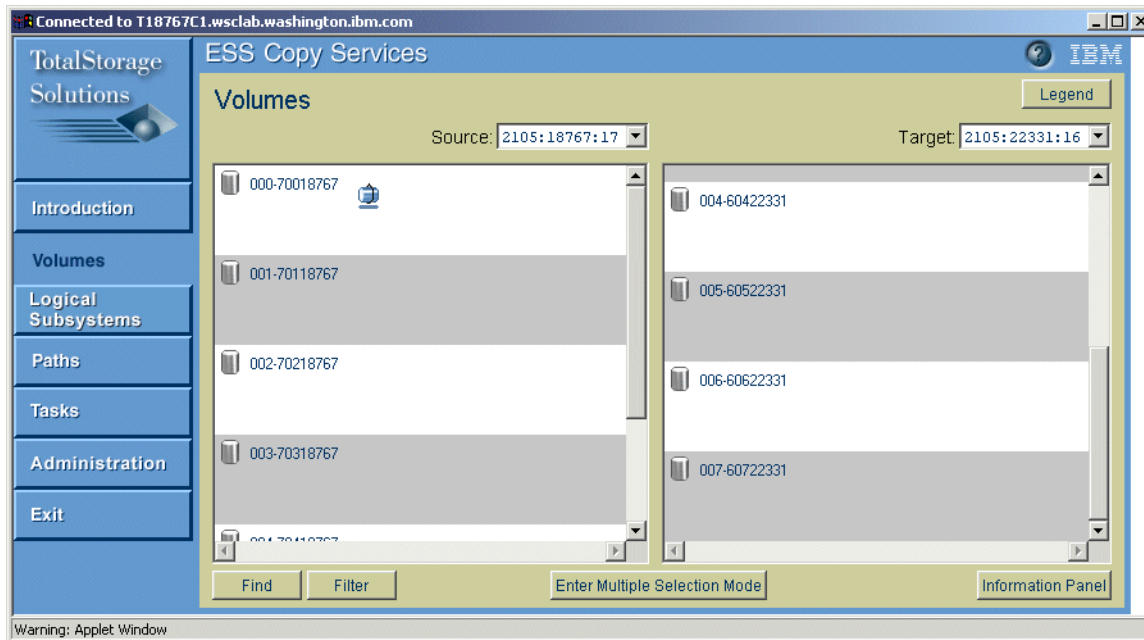
Figure 6-113   Volumes panel: Result of a terminate PPRC with Schedule with target option

6.12.9 Resynchronizing PPRC copy pairs

When you resynchronize a copy pair, all changed cylinders on the source volume are copied to the target volume. PPRC synchronous operations are then resumed for the pair. You will use this method to resynchronize two volumes in suspended mode. The method described here is the same method as 6.12.5, "Converting a PPRC-XD copy pair to synchronous PPRC" on page 248.

Use the Volumes panel to resynchronize a PPRC copy pair. On the left side, select the source LSS, and on the right side, the target LSS. This is done using the drop-down menu at the top of the Volumes menu. Select the source volume, and, with a right-click, the target. Once you have selected the source and the target, you do a second right-click on the target to bring up the Task Wizard (Figure 6-114 on page 257).
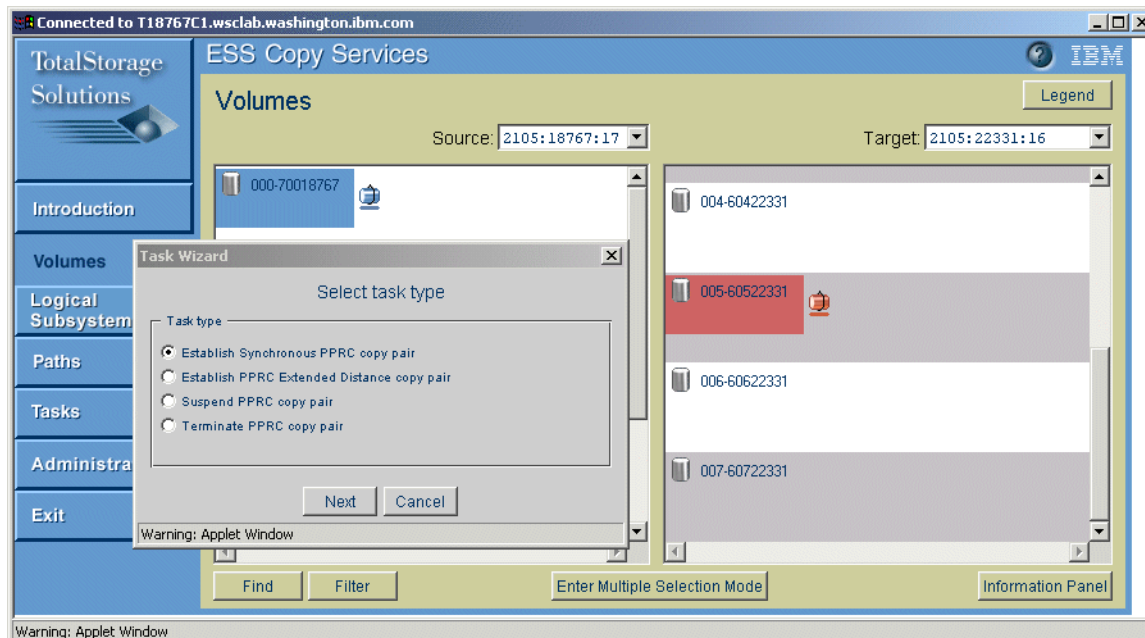
Figure 6-114   Volumes panel and Task Wizard

Select the **Establish synchronous PPRC copy pair** option and click **Next**. On the next panel, where you can select the copy options (Figure 6-115).

Figure 6-115   Task Wizard: Select copy options

Select the **Copy out-of-sync cylinders only** initialization option and decide whether to use the following two options (those options are described in 6.12.3, "Establishing a synchronous PPRC pair" on page 239):

► **Permit read from secondary**

► **Suspend PPRC after establish complete**

Click **Next**. From the next panel (Figure 6-116), you can either **Save**, **Run**, or **Cancel** the task. Once you run the task, all out-of-sync cylinders are copied from the primary to the secondary volume.
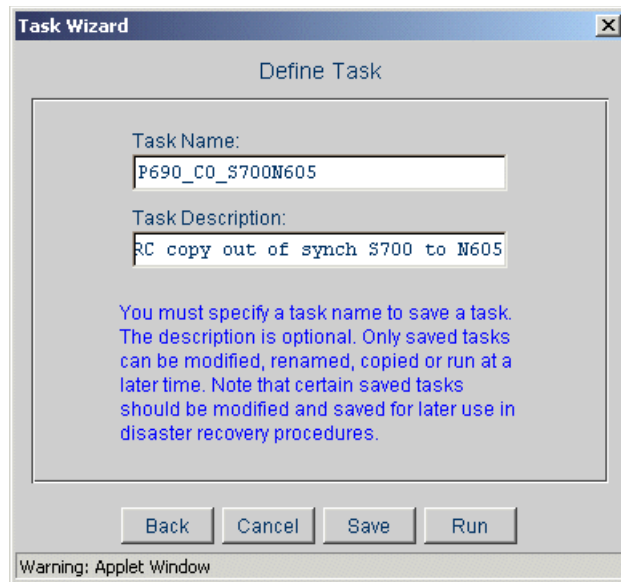
Figure 6-116   Task Wizard: Define Task

> **Tip:** Alternatively, you can use the Logical Subsystem Panel to select **Resynchronize PPRC Copy Pairs**. For further information, refer to the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448.

6.12.10 Creating a PPRC consistency group

The PPRC consistency group will be use to get consistent data at the remote site when there are unplanned outages and other problems. It can also be used to build consistent data at a remote site for backup purposes, for example. Refer to 4.5, "PPRC data consistency" on page 72 for more information.

You will define three different types of tasks related to consistency groups:

1. Tasks that Establish Path with the **PPRC consistency group** option. The implementation of this task is explained in detail in this section.

2. Tasks that **Freeze PPRC consistency group**. These tasks are established at the LSS level. The implementation of this task is explained in detail in 6.12.11, "Freezing a PPRC consistency group" on page 262.

3. Tasks that resume the I/O on the "frozen" volumes (**Consistency Created** tasks). These tasks are established at the LSS level. The implementation of this task is explained in detail in 6.12.12, "Thawing a PPRC consistency group" on page 264.

Here is a short description of how and under which circumstances those tasks will be executed. For more information, refer to 4.5, "PPRC data consistency" on page 72.

A PPRC consistency group is a set of PPRC volume pairs that have the same source and target LSS. ESS Copy Services creates the consistency group when you run a task to establish a path between the primary and secondary LSSs if this task has the **PPRC consistency group** option selected.

When an error occurs that affects any of the volumes in a PPRC consistency group, the ESS causes the volume where the error is detected to enter a suspended state. If the volume is participating in a consistency group, it enters a queue full state. The primary host temporarily queues all the updates to the primary volume. The primary volume will stay in the queue full state during two minutes (two minutes is the default value; you can modify this value using the LSS properties panel (see 6.6.4, "LSS properties" on page 190).

Within those two minutes, you can detect that the volume enters a queue full state (via SNMP traps, e-mail notification, or host I/O error messages) and trigger the execution of the task that will freeze all the volumes in the PPRC consistency group. This task will remove the established path between the LSS, causing the primary volumes to be in a suspended state. Furthermore, if the volumes in the LSS are in a consistency group, all primary volumes will be in a queue full state.

Once the Freeze command is executed, you will execute a **Consistency created** task to resume the queue full state of the primary volume. This will allow the primary hosts to perform write I/Os on the primary volumes. If you do not perform this **Consistency Created** task, the queue full state will be automatically resumed after a delay (default is two minutes).

To come back to the initial state, you must re-establish the path between the two LSSs (The freeze command removes the path between LSSs) and re-establish the PPRC relationship.

Establishing a path with the PPRC consistency group option

To create a PPRC consistency group, click **Paths** in the navigation frame of an ESS Copy Services panel. The Paths panel opens.

From the Logical Subsystem Sources list, select the logical subsystem (LSS) from which you want to create the paths. Click on the path in the Local Port/Remote Port column. Refer to Figure 6-117 on page 260 as you go through the steps listed here:

1. Select **Work with Paths**.
2. Select the source ESS from the **Logical Subsystem sources**.
3. Click on the source LSS.
4. Select the target ESS in the drop-down menu from the **Logical Subsystem targets**.
5. Click on the target LSS.
6. Click on the desired path in the third column labeled **Local Port/Remote Port**.
7. Right-click on the path highlighted.
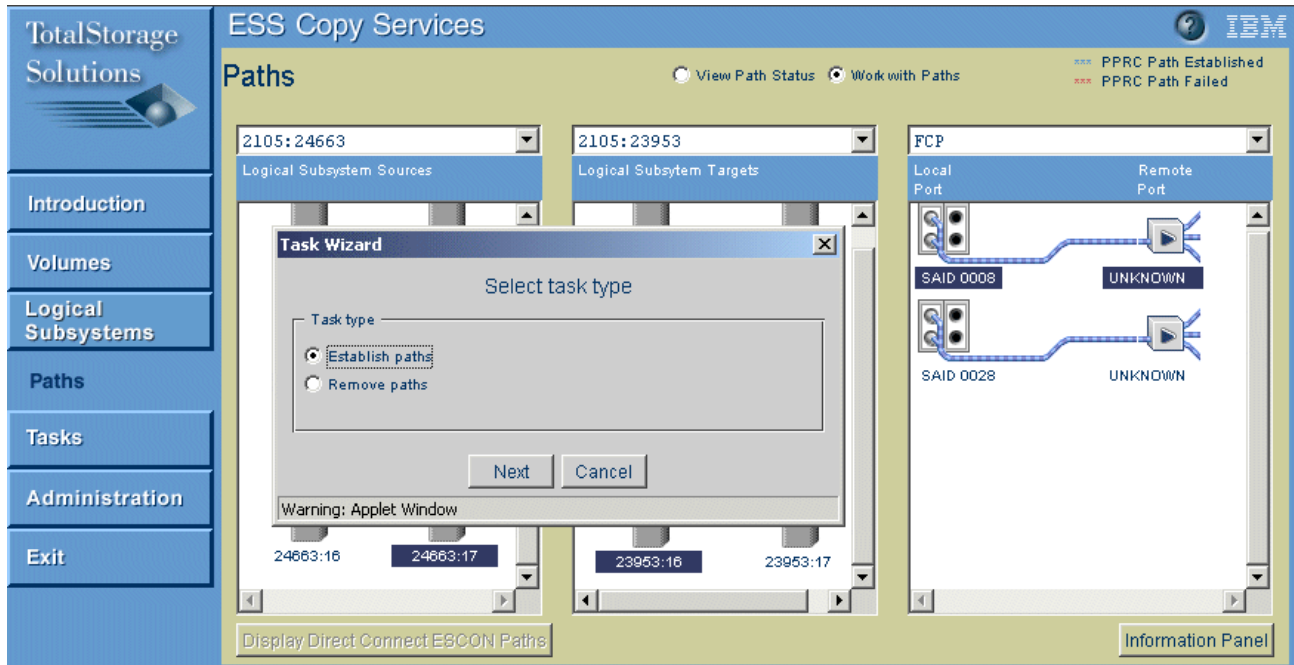
Figure 6-117   Paths panel: Establishing a path

8. Select **Establish path** and click **Next**.
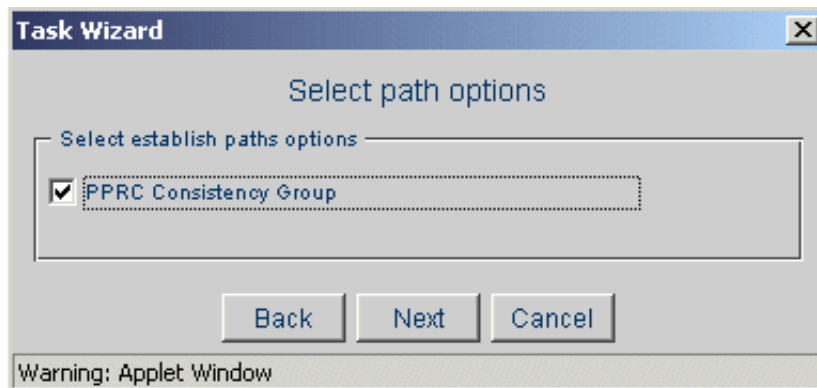
Figure 6-118   Task Wizard: Select path options - FCP link

9. Because we are working with an FCP link in our example, the Select path options Task Wizard will be like the one shown in Figure 6-118. Select **PPRC consistency group** and click **Next**. The Define task panel of the Task Wizard opens.

If you are working with an ESCON port, then you will see the Task Wizard shown in Figure 6-119 on page 261.
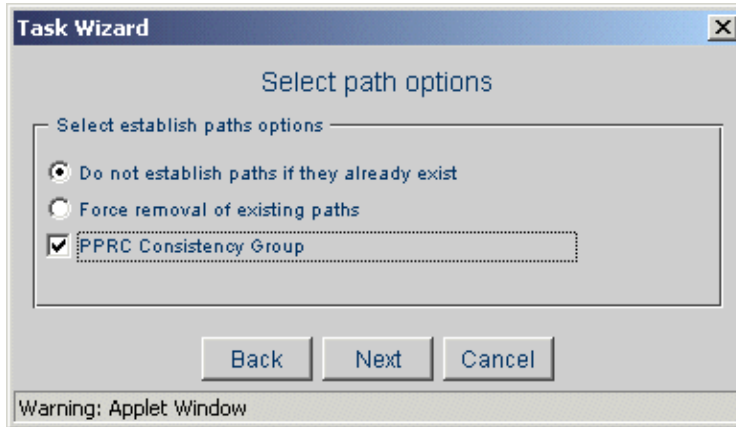
Figure 6-119   Task Wizard: Select path options - ESCON link

You can either click **Save** to save the task and run it later, or click **Run** to run the task immediately (see Figure 6-120).
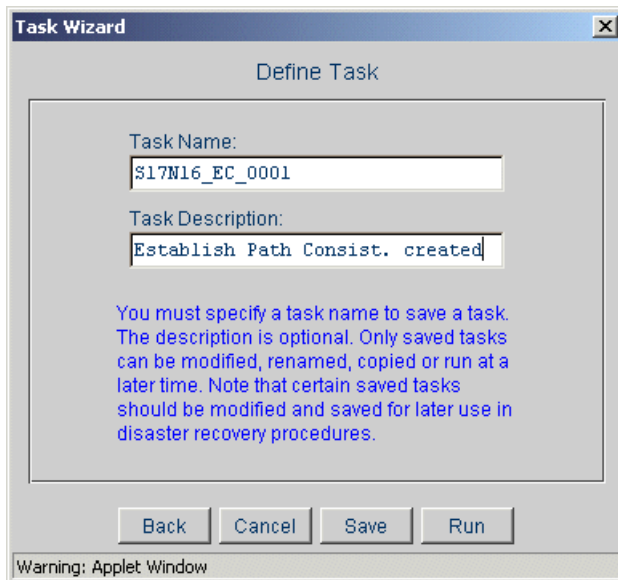
Figure 6-120   Task Wizard: Select path options

Attention: When you create paths for PPRC consistency groups, it is important to save the task. After you run the Freeze command and correct the problem that caused the suspension, you need to re-establish the paths. You must re-establish the paths before you run the task that resynchronizes the PPRC volumes.

Note: There is another way to create a consistency group: You can create a consistency group, including all the volumes, in one LSS by clicking on the **PPRC Consistency Group** check box in the LSS properties panel (see "PPRC consistency group" on page 191).

6.12.11 Freezing a PPRC consistency group

This task will remove the established path between two LSSs, causing the primary volumes to be in a suspended state. So performing a **Freeze PPRC consistency group** stops all operations from being propagated to the secondary volumes in the LSS for all volumes that were using this path. Furthermore, if any volumes in the LSS are in a consistency group, all primary volumes will be in a queue full state.

When you want to freeze several consistency groups of volumes spread across different LSSs, the freeze operation is best implemented by creating and saving tasks for all the associated LSS pairs and then grouping the tasks. When you need to perform a freeze operation, the group can be run from the Task panel or from the command line interface. See 6.8.2, "Grouping and ungrouping tasks" on page 197 for information on grouping tasks.

For the freeze operation to work and create a consistent image of the data at the remote site, the **Establish Paths** task must indicate that the LSS to LSS relationship is part of a consistency group (see 6.12.10, "Creating a PPRC consistency group" on page 258).

You can invoke the freeze task when a freeze trigger event is detected (for example, an e-mail notification, an SNMP alert, or a host-detected failure condition). Following the freeze operation, a *consistency-group-created* operation should be run to stop any queue full states that the volumes are in. See 4.5, "PPRC data consistency" on page 72 for more information about the freeze and consistency-group-created operations.

To create a task that freezes a consistency group, click the **Logical Subsystems** button in the navigation frame of an ESS Copy Service panel. The Logical Subsystems panel opens, and within the Select list in the upper right corner of the panel, select the source ESS. A panel showing all the logical subsystems (LSSs) is displayed, with the appropriate names and serial numbers indicated below each LSS icon. Click the source LSS. The LSS ID of the source LSS turns blue (Figure 6-121). Click the select list again and select the target ESS. The LSSs of this ESS are displayed; right-click on the target LSS. The LSS ID of the target LSS turns red (Figure 6-122 on page 263).
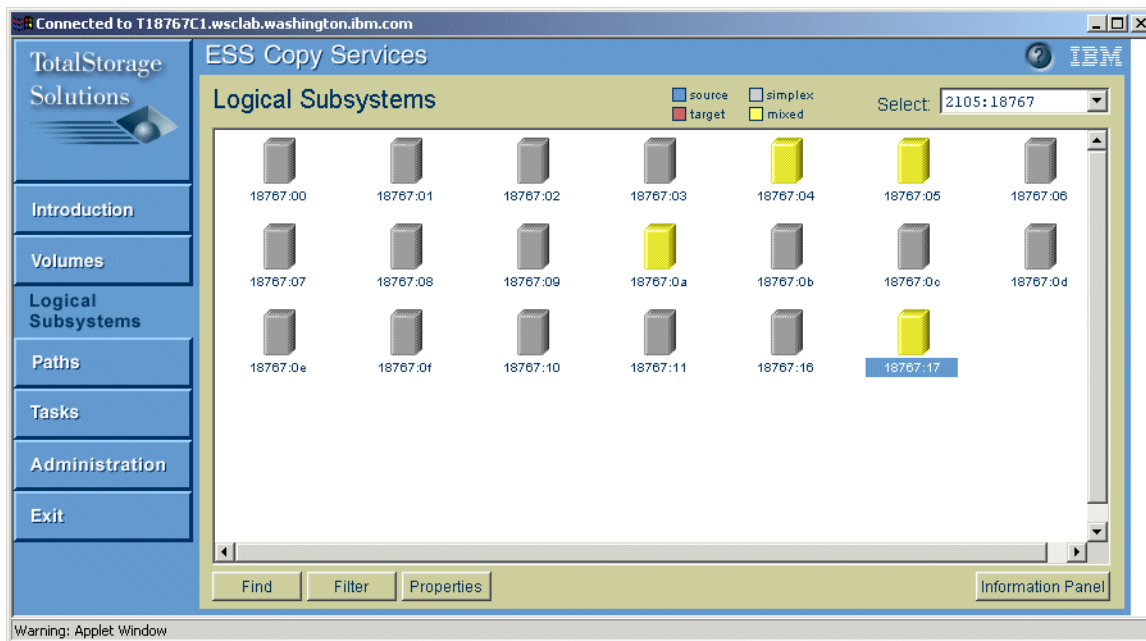
Figure 6-121 LSS panel: LSS source selected
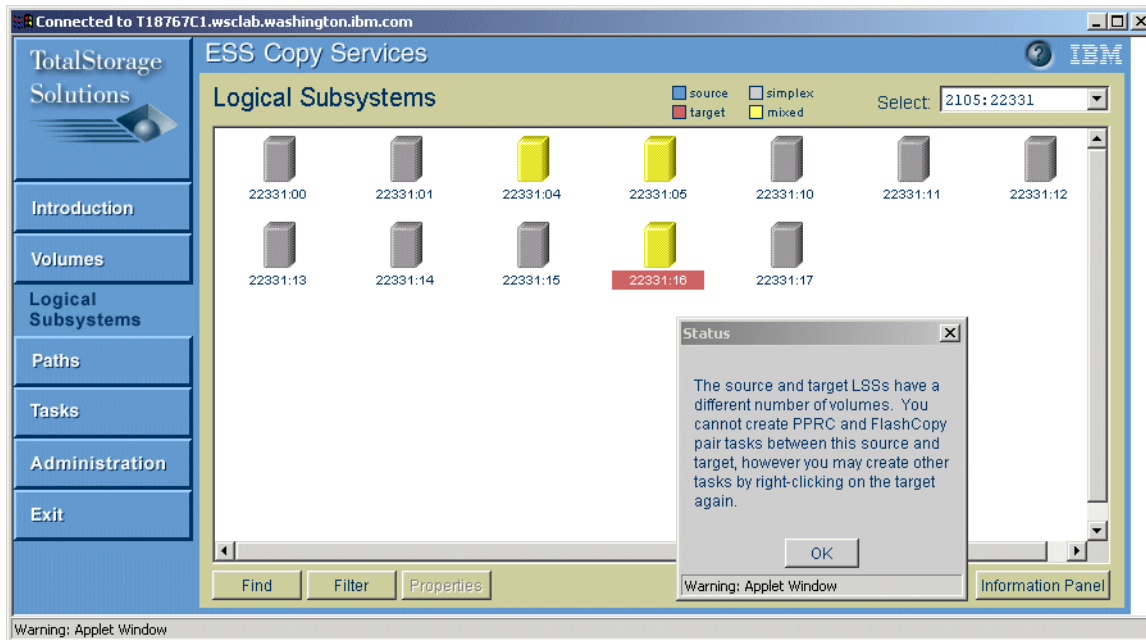
Figure 6-122   LSS panel: LSS target selected

Right-click again on the target LSS. A warning message may pop up if the configuration of the selected source and target LSSs are different (see Figure 6-122). If so, click **OK** on this message, and right-click again on the target LSS.

In the Task Wizard. Select **Freeze PPRC consistency group** (Figure 6-123). Then, click **Next**.
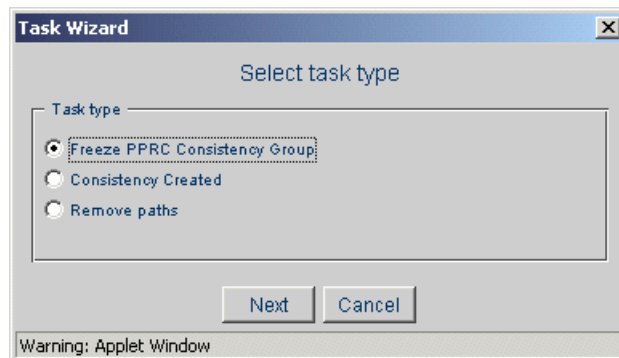
Figure 6-123   Task Wizard: Select task type

Within the next panel (Figure 6-124 on page 264) you can decide whether to **Run**, **Save**, or **Cancel** the task. If you want to save the task, type the task name (mandatory) and the task description (optional) in the appropriate fields. You can also click **Run** to execute the task immediately.
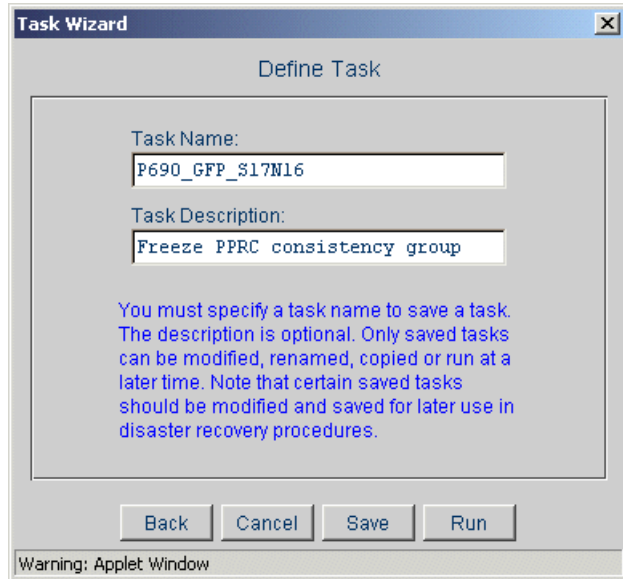
Figure 6-124   Task Wizard: Define Task

> **Attention:** After creating and saving the task to freeze a PPRC consistency group, you should create the corresponding consistency-group-created task. We recommend that you modify the task that you just created so that it runs a consistency-group-created operation, then name and save the task. For further information about modifying a task, refer to 6.8.5, "Modifying a task" on page 199.

6.12.12 Thawing a PPRC consistency group

The thawing of a PPRC consistency group (consistency-group-created operation) follows the successful freeze for all LSSs in the consistency group. It allows normal I/O operations from the host to the ESS to resume. For further information, refer to 4.5, "PPRC data consistency" on page 72).

You can either use the Task panel to modify a **Freeze PPRC Consistency Group** task to create an appropriate task to thaw the consistency group (refer to 6.8.5, "Modifying a task" on page 199 for further information), or you can use the Logical Subsystems panel to create a task to thaw a PPRC consistency group.

This procedure of using the Logical Subsystems panel is almost identical to the previous procedure, which you can use to freeze a PPRC consistency group, except that you select the **PPRC consistency created** option in the Select task type panel of the Task Wizard (Figure 6-125 on page 265).
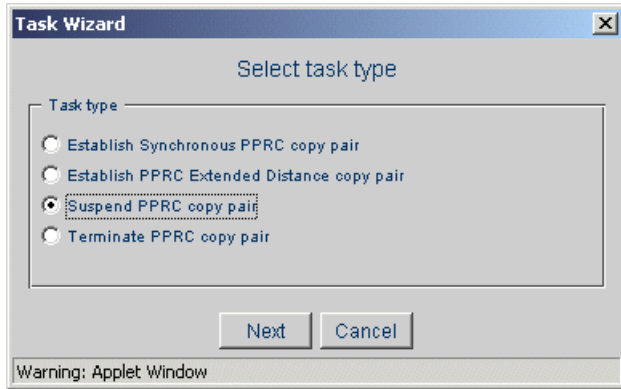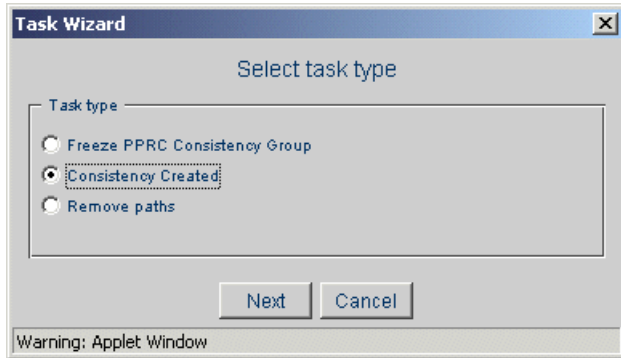
Figure 6-125   Task Wizard: Select task type

Click **Next**. In the next panel (Figure 6-127), you can decide whether to **Run**, **Save**, or **Cancel** the task. If you want to save the task, type the task name and the task description in the appropriate fields. Click **Run** to execute the task immediately.
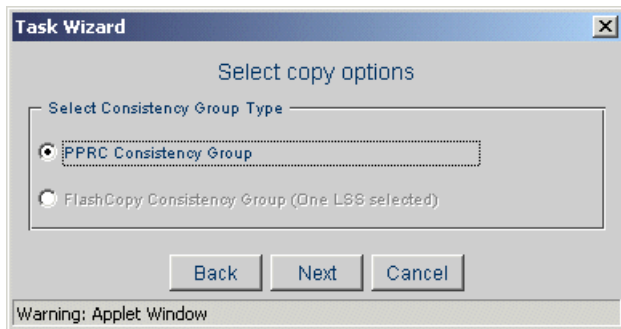
Figure 6-126   Task Wizard: Select copy options
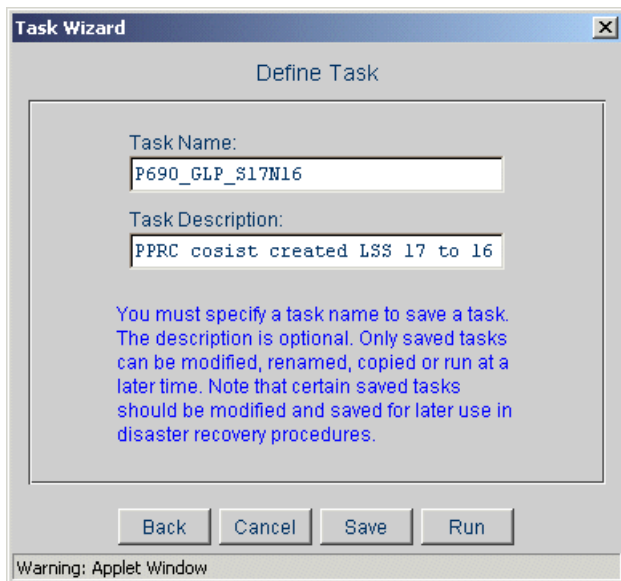
Figure 6-127   Task Wizard: Define Task

6.12.13 Performing a PPRC Failover

The PPRC Failover task can be used in a disaster recovery process. This task is part of the process used in Chapter 10, "Disaster recovery" on page 399 to restart a production environment using the volumes from the recovery site. In a disaster recovery process, the fail-over task is followed by a fail-back task once a path from the recovery site to the production site is established.

Use the Volumes panel to establish the PPRC pairs. The process is identical to the one used when establishing a PPRC relationship. Refer to 4.2, "Synchronous PPRC" on page 54 to get a full detailed description of the process.

When creating the PPRC Failover task, you must be very careful when choosing the source volume and the target volumes for the task:

► The secondary volume of the established PPRC pair should be selected as the primary (source) volume for the new PPRC Failover task. Click to select the volume.

► The Primary of the established PPRC pair should be selected as the secondary (target) volume for the new PPRC Failover task. Right-click to select the volume.

Figure 6-128 shows an example of how the volumes must be selected:

► A PPRC relationship is established and running between the primary volume 60018767 at the production site and the secondary volume 30922331 at the recovery site.

► The volume 30922331 is the secondary (target) volume for the established PPRC, so it is selected as the primary (source) volume.

► The volume 60018767 is the primary (source) volume for the established PPRC, so it is selected as the secondary (target) volume.
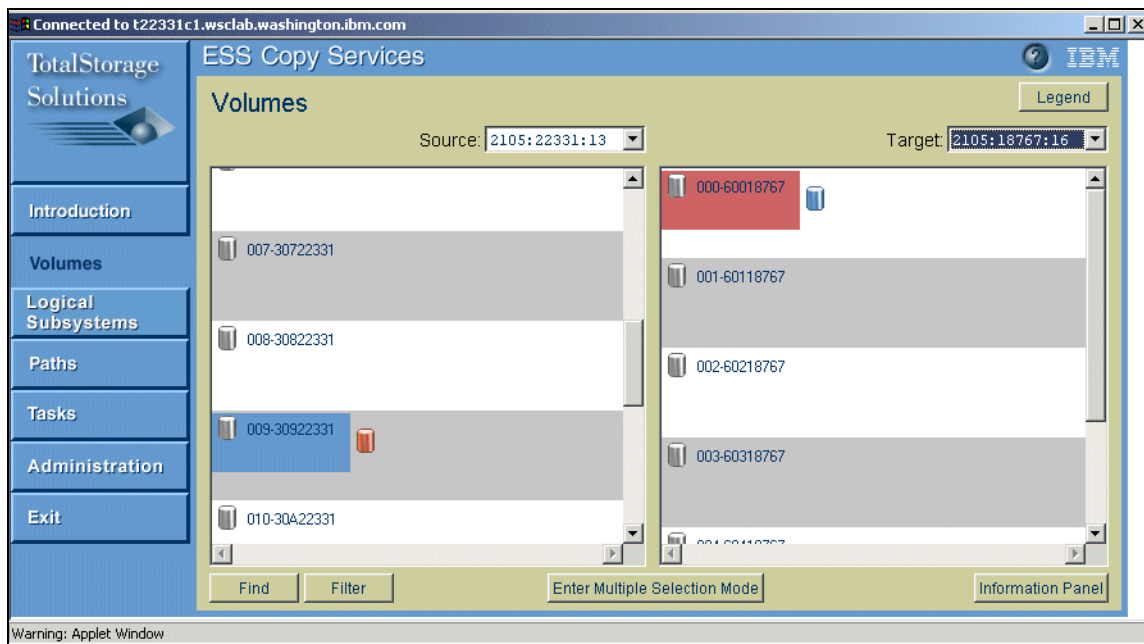
Figure 6-128   Volumes panel: Volume selection for a PPRC Failover

Once the two volumes are selected, another right-click on the target volume will open the Task Wizard.

In the Task Wizard, select the **Establish Synchronous PPRC copy pair** task (Figure 6-129) and click **Next**.
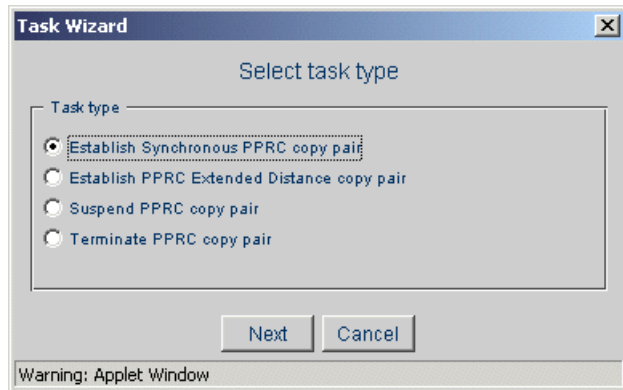
Figure 6-129   Task Wizard: Select task type

In the second Task Wizard panel (Figure 6-130), select the fail-over task and click **Next.**
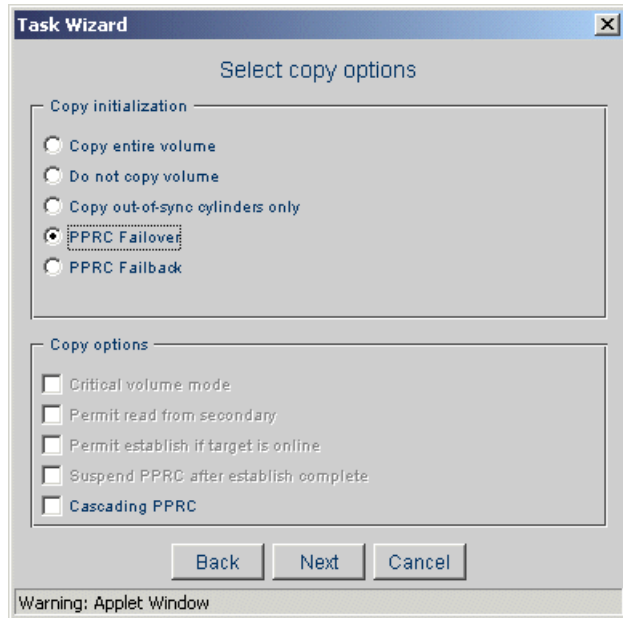
Figure 6-130   Task Wizard: Define Task

In the next panel (Figure 6-130), you can decide whether to **Run**, **Save**, or **Cancel** the task. If you want to save the task, type the task name and the task description in the appropriate fields. You can also click **Run** to execute the task immediately.
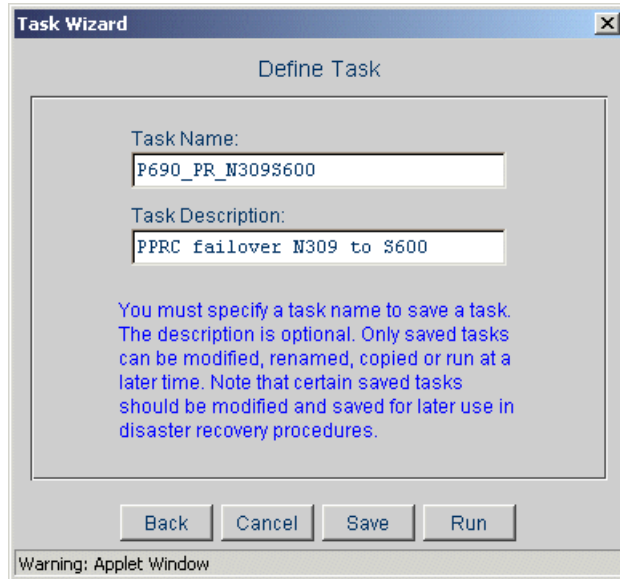
Figure 6-131   Task Wizard: Define Task

Once the PPRC Failover task is established on a synchronous PPRC pair, the volumes states change as shown in Figure 6-132 on page 269:

► The initial primary volume at the production site, remains unchanged. Its state is: Source volume of a PPRC relationship. The target volume of this relationship is the volume at the recovery site. The status of the volume at the production site is: Full duplex of Synchronous PPRC relationship. If you perform any write operation on this volume, its status will change from Full duplex to Suspended.

► The initial secondary volume at the recovery site changed from Target volume to the new state of: Source volume of PPRC relationship. The target of this relationship is the initial primary volume at the local site. The status of the volume at the recovery site is: Suspended of Synchronous PPRC relationship.

After the PPRC Failover task, the volume at the recovery site is available to be mounted and used by a host. You can use this volume to restart your application from the recovery site.
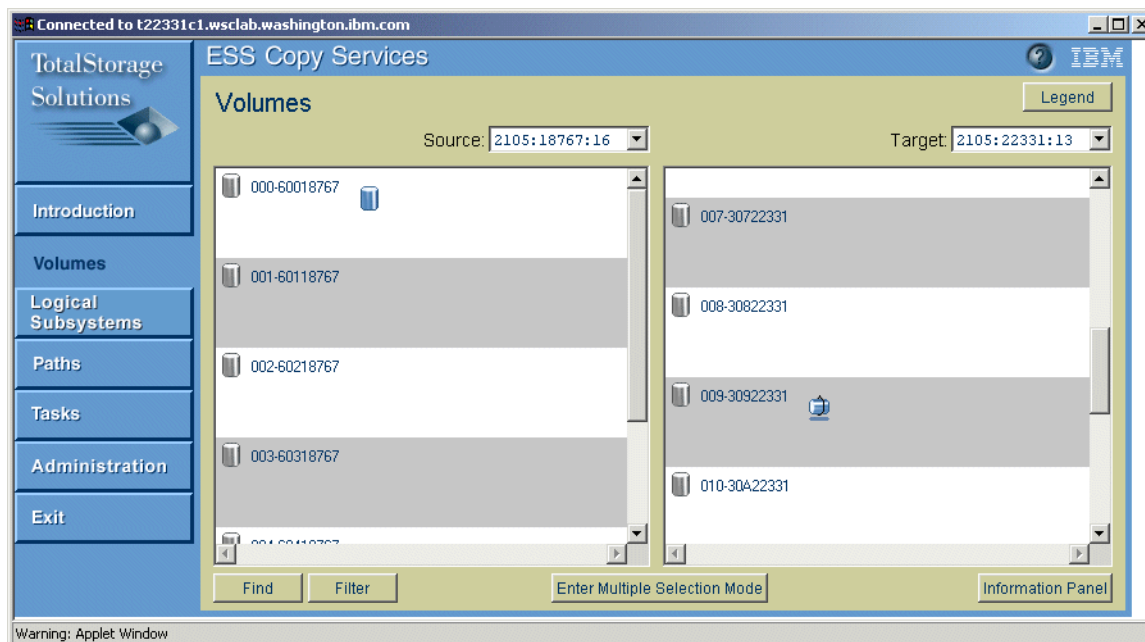
Figure 6-132   Volumes panel: Result of a fail-over task

Based on the volume IDs used in the screen shots, Table 6-4 presents a summary of the PPRC Failover rules when selecting the volumes and the result of the volume status after running the task.

Table 6-4   Summary of the PPRC Failover creation and results

	Production Site Volume 60018767	**Recovery Site Volume 30922331**
Initial status, before running the PPRC Failover task	Primary volume (source) of a PPRC. Full Duplex. The target (secondary) volume is volume 30922331.	Secondary volume (target) of the PPRC Full Duplex. The source (primary) volume is volume 60018767.
Selection rules, when creating the PPRC Failover task	Select as secondary (target) volume.	Select as primary (source) volume.
Final status, after running the PPRC Failover task	Primary volume (source) of a PPRC. Full Duplex. The target (secondary) volume is volume 30922331.	Primary volume (source) of a PPRC. Suspended. The target (secondary) volume is volume 60018767.

6.12.14 Performing a PPRC Failback

This task is part of the process used to restart a production environment using the volumes from the recovery site. See Chapter 10, "Disaster recovery" on page 399 for details. The fail-back task is generally performed after a fail-over task.

After establishing a PPRC Failover task, the volume at the recovery site is available to be mounted and used by a host. You can use this volume to restart your application from the recovery site (see previous section). Establishing a PPRC with the **Failback** option will

resynchronize the PPRC volumes at the recovery site with their peer volumes at the production site.

> **Important:** Before you run the PPRC Failback task, it is necessary to established a PPRC path from the recovery site to the production site between the concerned LSSs.

Use the Volumes panel to establish the PPRC pairs. The process is identical to the one used when establishing a PPRC relationship. Refer to 6.12.3, "Establishing a synchronous PPRC pair" on page 239 to get a full detailed description of the process.

When creating the PPRC Failback task you must be very careful when choosing the source volume and the target volumes for the task.

► The primary (source) volume of your new PPRC Failback task must be the volume at the recovery site. This volume was initially the secondary volume of the PPRC relationship. After the PPRC Failover task has run, this volume state is changed to become a primary (source) volume in a Suspended state. Click to select the volume.

► The secondary (target) volume of your new PPRC Failback task must be the volume at the production site. This volume is the primary volume of the initial PPRC relationship. After the PPRC Failover task has run, this volume state remains unchanged. If you perform any write operation to this volume after the PPRC Failover task was established, then its status changes to Suspended. Right-click to select the volume.

Figure 6-133 shows an example of how the volumes must be selected:

► A PPRC Failover task was performed on an established PPRC pair (primary volume was 60018767 at the production site and the secondary volume was 30922331 at the recovery site).

► The volume 30922331 at the recovery site is now a primary volume in a PPRC relationship and is in a suspended state, so it is selected as the primary (source) volume.

► The volume 60018767 at the production site is still the primary (source) volume for the established PPRC, so it is selected as the secondary (target) volume.
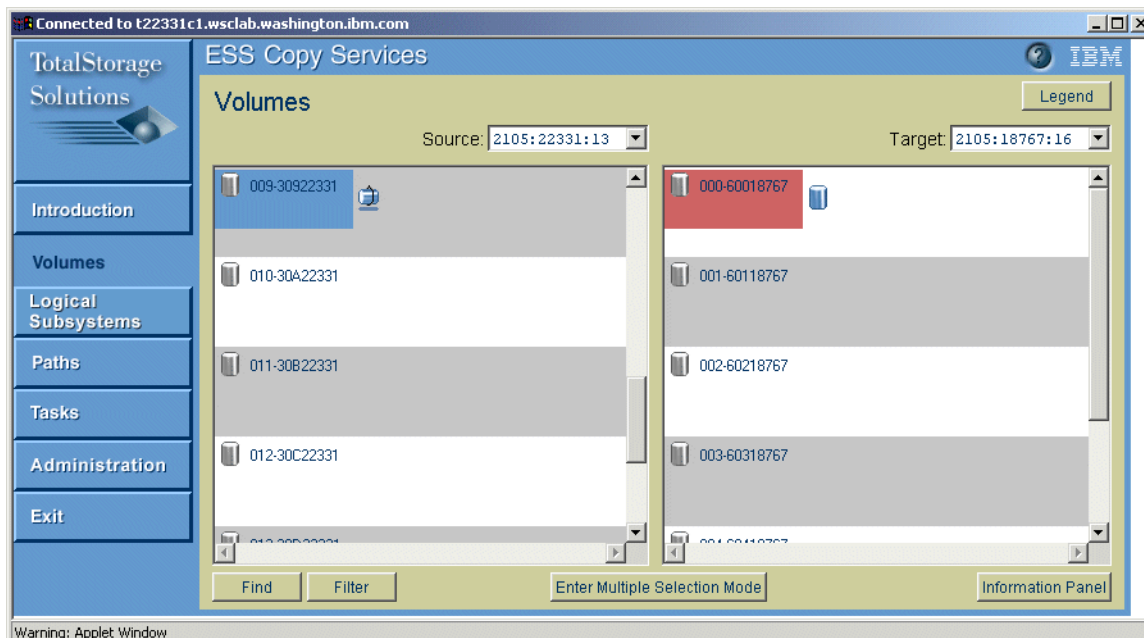
Figure 6-133   Volumes panel: Selecting the volume for a PPRC Failback task

Once the two volumes are selected, another right-click on the target volume will open the Task Wizard (Table 6-134).

In the Task Wizard, select the **Establish Synchronous PPRC copy pair** task (Figure 6-134) and click **Next**.
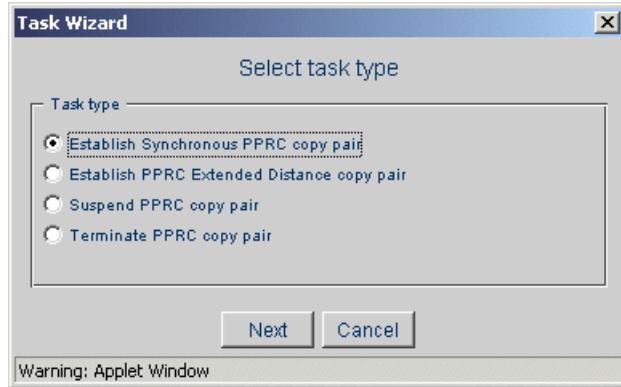
Figure 6-134   Task Wizard: Select task type

In the second Task Wizard panel (Figure 6-135), select the **Failback** option and click **Next.**
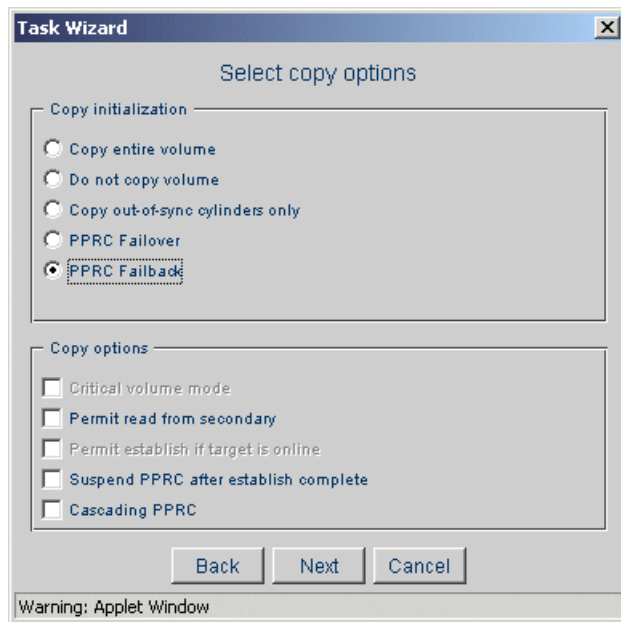
Figure 6-135   Task Wizard: Select copy options

In the next panel, you can decide whether to **Run**, **Save**, or **Cancel** the task. If you want to save the task, type the task name and the task description in the appropriate fields. Click **Run** to execute the task immediately.
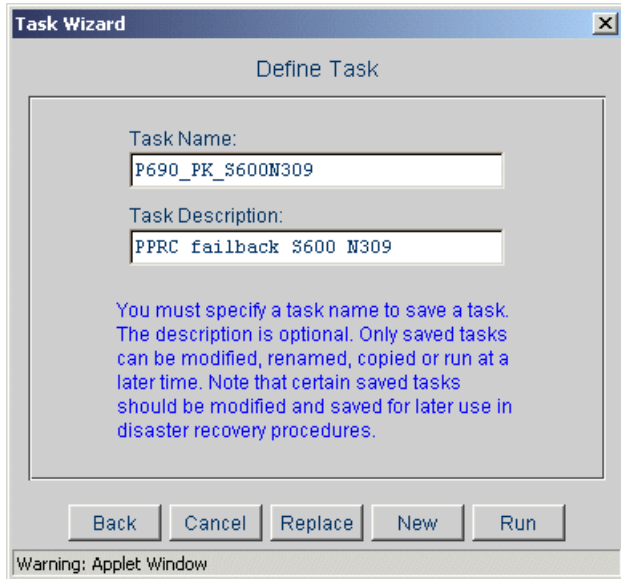
Figure 6-136   Task Wizard: Define Task

Once the PPRC Failback task is established, the PPRC relationship is resynchronized between the recovery site and the production site. The volumes' state changes are shown in Figure 6-137:

► The volume at the production site is now the secondary (target) volume of the PPRC relationship.

► The volume at the recovery site is now the primary (source) volume of the PPRC.

► The status of the two volumes will be first Copy pending, then Full Duplex, when the copy of the out-of-sync sectors is finished.
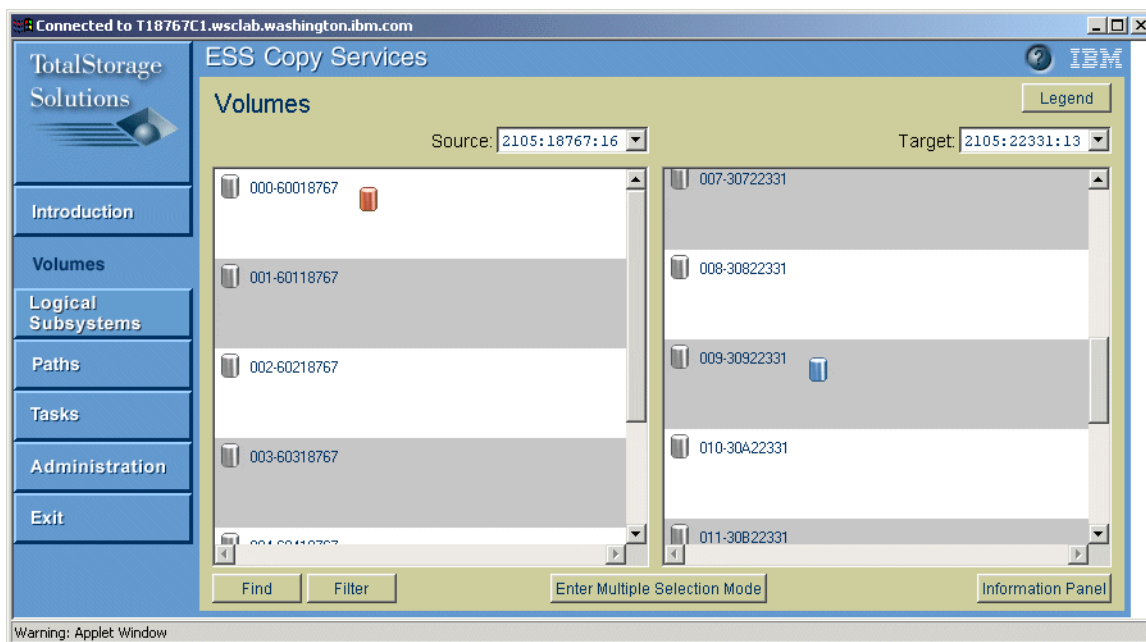
Figure 6-137   Volumes panel: Result of a fail-back task

Based on the volumes IDs used in the screen shots, Table 6-5 presents a summary of the PPRC Failover and PPRC Failback rules when selecting the volumes and the result on the volume status after running each tasks.

Table 6-5   Summary of the PPRC Failover and Failback creation and results

	Production Site Volume 60018767	**Recovery Site Volume 30922331**
Initial status, before running the PPRC Failover task	Primary volume (source) of a PPRC. Full Duplex. The target (secondary) volume is volume 30922331.	Secondary volume (target) of the PPRC Full Duplex. The source (primary) volume is volume 60018767.
Selection rules, when creating the PPRC Failover task	Select as secondary (target) volume.	Select as primary (source) volume.
Status, after running the PPRC Failover task	Primary volume (source) of a PPRC. Full Duplex. The target (secondary) volume is volume 30922331.	Primary volume (source) of a PPRC. Suspended. The target (secondary) volume is volume 60018767.
Selection rules, when creating the PPRC Failback task	Select as secondary (target) volume.	Select as primary (source) volume.
Final status, after running the **PPRC Failback** task	Secondary volume (target) of the PPRC Copy pending. The source (primary) volume is volume 30922331.	Primary volume (source) of a PPRC. Copy pending. The target (secondary) volume is volume 60018767.

6.12.15 Establishing an Asynchronous Cascading PPRC

PPRC Version 2 includes a new option, Asynchronous Cascading PPRC, when establishing a PPRC relationship. This option allows a PPRC secondary volume to also be a PPRC primary volume for a different relationship.

You will use this option to create an Asynchronous Cascading PPRC relationship across three ESS volumes (The primary volume, the secondary volume and the tertiary volume). To get details about the benefits of Asynchronous Cascading PPRC, refer to 4.4, "Asynchronous Cascading PPRC" on page 65.

Terminology

Here is the terminology that is used when discussing Asynchronous Cascading PPRC:

► A site where the production applications run is referred to as the *local site* or the *application site*.

► An *intermediate site* is also called a *bunker site* and represents a site where the application site data is mirrored by means of synchronous PPRC. The local site and the intermediate site are often two different locations, but they can also be the same location.

► A site where the local site data is mirrored via the intermediate site is referred to as the *remote* site.

► The synchronous PPRC pairs established between the volumes at the local host site and the intermediate site are called the *local PPRC pairs*.

► The PPRC pairs created between the intermediate site and the remote site are called the *remote PPRC pairs.*

Necessary tasks

In order to create an Asynchronous Cascading PPRC between three volumes, you have to create and run four tasks:

► Two tasks to establish the required paths:

    a. One path is needed from the local to the intermediate site

    b. One path is needed from the intermediate path to the remote site

► One task to establish a PPRC relationship between the Primary volume and the Secondary volume.

► One task to establish a PPRC relationship between the Secondary volume and the Tertiary volume. You must select the Asynchronous Cascading PPRC option for this task in order to allow the Secondary Volume of your Asynchronous Cascading PPRC to be both a secondary volume and a primary volume in two different relationships.

Allowed configurations

There are three available configurations in Asynchronous Cascading PPRC:

1. SYNC $\rightarrow$ XD

   This is a typical Asynchronous Cascading PPRC configuration. There is a synchronous PPRC relationship from the local site to the intermediate site, and then an asynchronous PPRC-XD relationship from the intermediate site to the remote site.

   In case of an unplanned outage at the local site, the intermediate site still contains all of the data, since it is an exact copy of local site. The intermediate site would complete sending the updates to the remote site and then the remote site would have a copy of data up to the time of failure.

2. XD $\rightarrow$ XD

   This configuration consists of a PPRC-XD relationship from the local site to the intermediate site, and PPRC-XD relationship from the intermediate site to the remote site. It may be used to make multiple copies of data, without impact on the application I/O. The intermediate and remote volumes are 'fuzzy', and it is necessary, from time to time, to resynchronize them in order to make point-in-time consistent backup.

3. SYNC $\rightarrow$ SYNC

   This allows a synchronous PPRC relationship between the local and intermediate sites as well as between the intermediate and remote sites. However, this Asynchronous Cascading PPRC configuration may have a severe impact on the application I/O, because each write I/O has to wait for two PPRC data transfers. The application has to wait for the confirmation that data is written to the intermediate and remote ESS cache and NVS.

   This configuration may be useful to get a consistent copy at the remote location, but only if there is a little or no I/O at the local site.

► SYNC $\rightarrow$ XD and SYNC $\rightarrow$ SYNC combination

   This configuration combines the above described SYNC $\rightarrow$ XD and SYNC $\rightarrow$ SYNC PPRC cascading solutions. Basically, the SYNC $\rightarrow$ XD configuration is used, but periodically you can switch to SYNC $\rightarrow$ SYNC mode in order to force the remote PPRC pairs to go to SYNC and take a consistent copy at the remote site. Switching remote PPRC pairs from PPRC-XD to PPRC SYNC mode is usually done when the application using the primary PPRC volume(s) is stopped or quiesced so there is no I/O at the local site.

Establishment order

There are basically two distinguished orders used to set up an Asynchronous Cascading PPRC environment, and both of them are valid:

► Recommended order

– Establish remote PPRC pairs in XD mode with the **Do not copy volume** option

– Establish local PPRC pairs in XD mode with the **Copy entire volume** option and then go to synchronous

► Alternative order

– Establish local PPRC pairs with the **Copy entire volume** option

– Wait for the volumes to reach DUPLEX state

– Establish remote PPRC pairs in XD mode with the **Copy entire volume** option

> **Note:** It is possible to do other combinations, for example, in the first case above, the remote pairs do not have to be established with the **Do not copy volume** option, but it is a waste of bandwidth to do so. However, the options listed above may be the most efficient.

Options to make a volume eligible for cascading

A pair can be made eligible to be a cascaded pair at the time the pair is established or at a later time:

► When a PPRC pair is initially established, the user can indicate that the pair is eligible for cascading.

► For an already established pair, the user can issue an establish command that makes it a cascaded pair without breaking the existing pair to do the re-establish.

Some examples are listed below:

► Intermediate Site $\rightarrow$ Remote site was already established with the indication that the volume at Intermediate site is eligible for cascading. When Local Site $\rightarrow$ Intermediate site is established, this now becomes an Asynchronous Cascading PPRC relationship.

► Intermediate Site $\rightarrow$ Remote Site was already established without the indication that it is eligible for cascading. An establish is issued again from the volume at Intermediate site to the volume at remote site, but this time with the Asynchronous Cascading PPRC option selected, and the pair remains established. Then Local Site $\rightarrow$ Intermediate site can be established, without breaking the Intermediate Site $\rightarrow$ Remote Site pair. This now becomes an Asynchronous Cascading PPRC relationship.

► Local site $\rightarrow$ Intermediate Site is already established. Then Intermediate Site $\rightarrow$ Remote Site can be established with an indication that the volume at the Intermediate Site is eligible for cascading, without breaking the Local site $\rightarrow$ Intermediate Site pair. This now becomes an Asynchronous Cascading PPRC relationship.

Establishing an Asynchronous Cascading PPRC with the WUI

You can invoke the PPRC cascading function using the ESS Copy Services Web user interface. More detailed descriptions of the ESS Copy Services Web user interface is in Chapter 5, "ESS Copy Services Web User Interface prior to LIC 2.2.0" on page 103. In addition, when you are planning to use the ESS Copy Services WUI for PPRC operation, refer to the *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7448.

An Asynchronous Cascading PPRC environment consists of local, intermediate, and remote sites. The following procedure should be applied when establishing an Asynchronous

Cascading PPRC environment (the example uses ESCON links and an ESS with LIC level 2.2.0):

1. Create the PPRC logical paths between the local and intermediate sites. As illustrated in Figure 6-138, the source LSS 2105:18767:16 is selected, where 18767 is the ESS serial number and the LSS number is 16 and represents the local site. A path is established to the target LSS 2105:18767:17.

> **Note:** The Paths panel that you see in Figure 6-138 and Figure 6-139 correspond to ESSs with LIC level 2.2.0. This panels look different when working with LIC level 2.3.0. Refer to 6.12.1, "Establishing paths" on page 232 for a description and discussion of the Paths panel use when using LIC level 2.3.0.
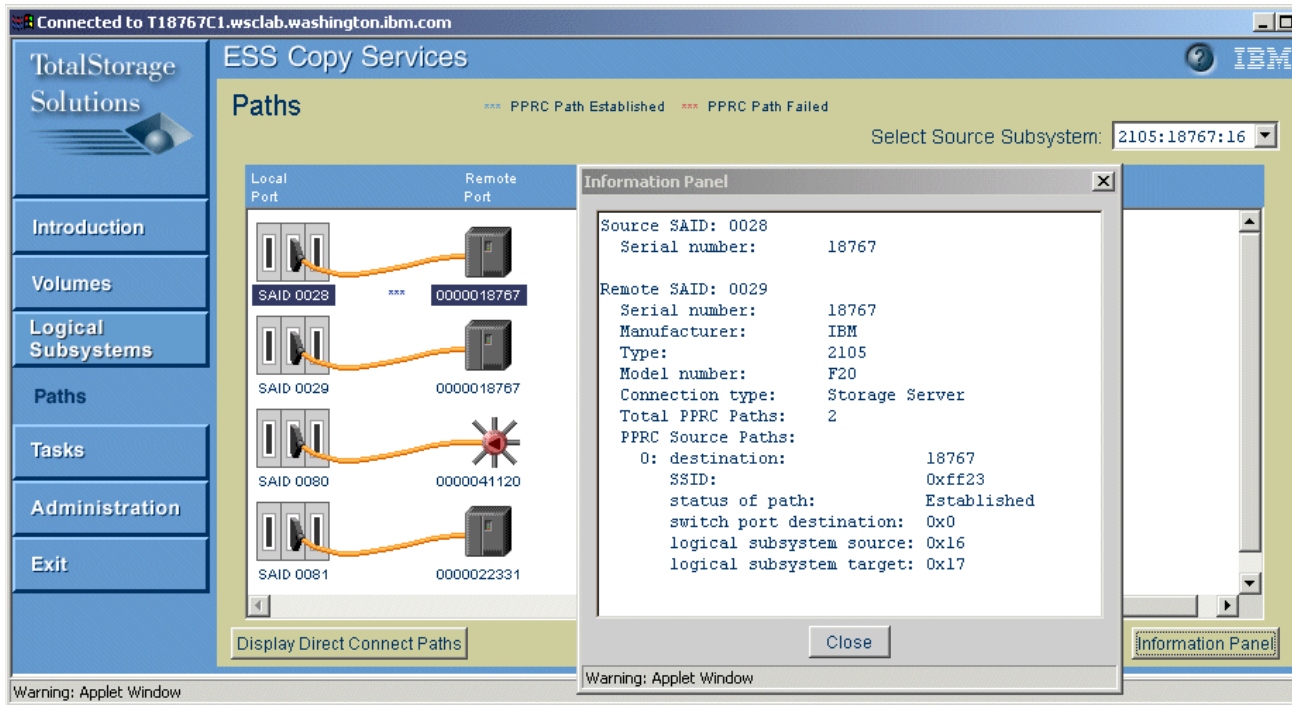
Figure 6-138   PPRC establish path: Local to intermediate site (ESS at LIC level 2.2.0)

In this example, local and intermediate volumes are placed in the same ESS disk subsystem but in different LSSs: LSS 17 is the intermediate. The path origin selected is the ESS ESCON adapter with the SAID 0028 (see 6.7.4, "Display Connection Paths" on page 195).

2. Establish the PPRC paths between the intermediate and remote sites, as illustrated in Figure 6-139 on page 277.
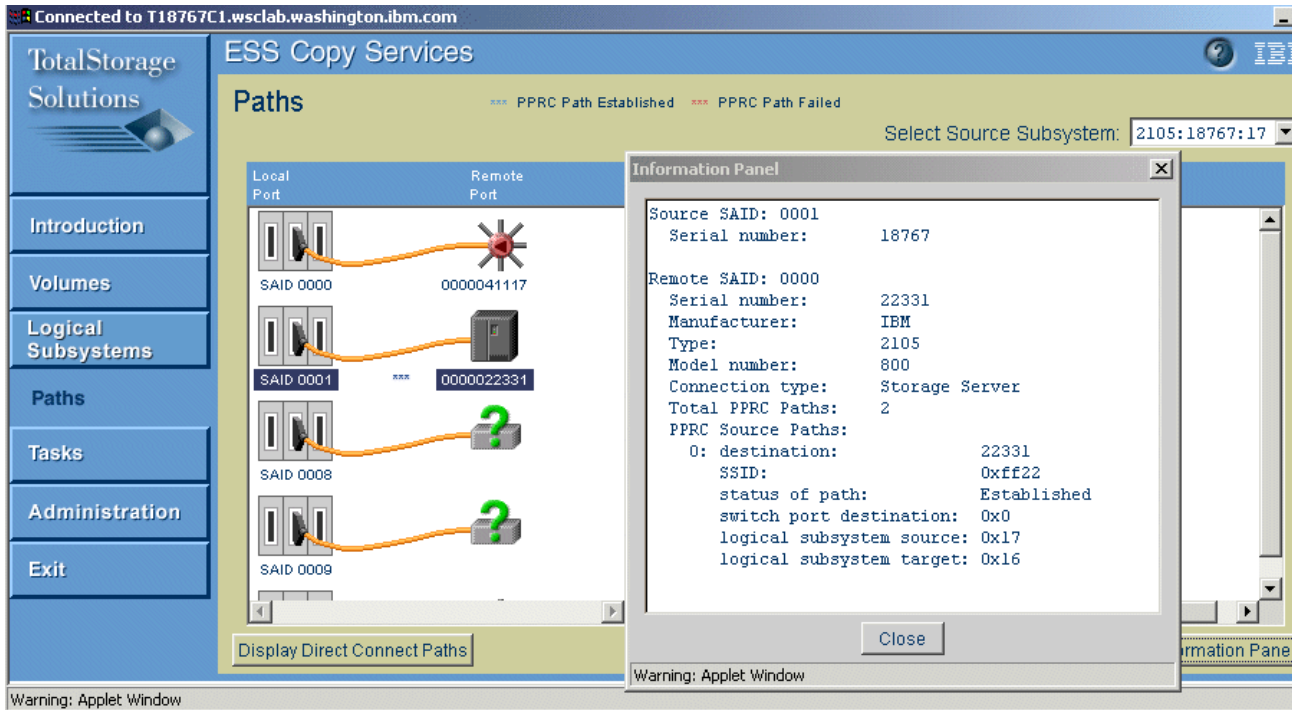
Figure 6-139   PPRC establish path: Intermediate to remote site (ESS at LIC level 2.2.0)

As you can see in the Information Panel in Figure 6-139, the remote site LSS is number 16 and belongs to another ESS disk subsystem with serial number 22331.

3. Once the logical paths are established, create the PPRC-XD asynchronous relationship between the volumes that resides in the intermediate site with the volume in the remote site. Select the appropriate remote PPRC primary and secondary (with a right-click) volumes, launch the Task Wizard (with another right-click), and then select **Establish PPRC Extended Distance copy pair**, as shown in Figure 6-140.
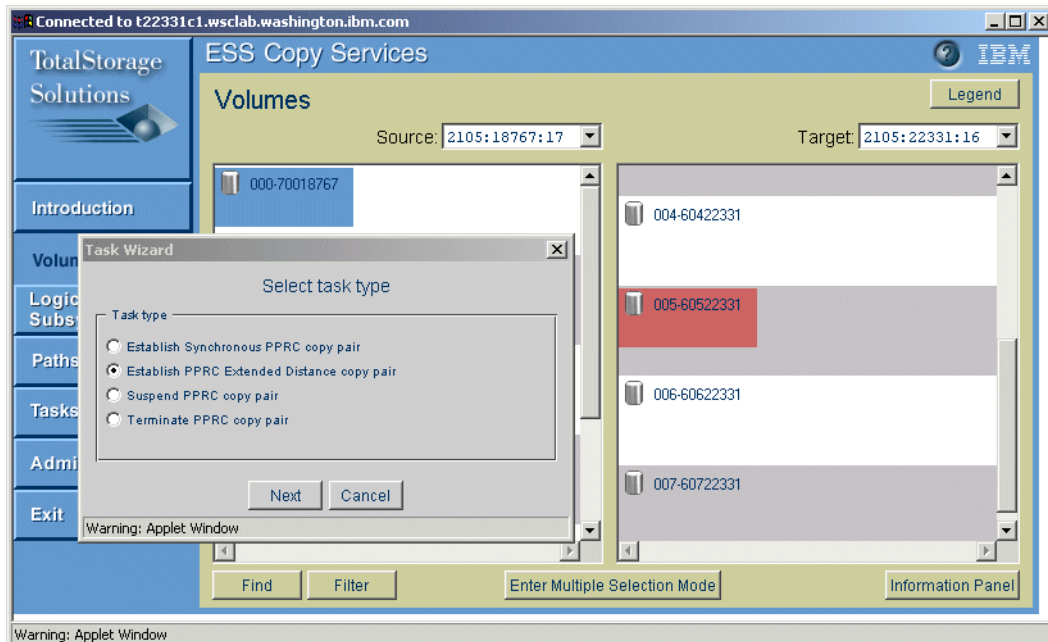
Figure 6-140   Establish PPRC-XD volume pair: Intermediate to remote site volume

In the above example, the remote PPRC pair consists of the intermediate volume 70018767 in LSS 17 in ESS 18767 (defined as an intermediate PPRC primary) and remote volume 60522331 in LSS 16 but in ESS 22331 (defined as a remote PPRC secondary).

In the Select copy options panel select **Do not copy volume** as a copy initialization mode and **Asynchronous Cascading PPRC** as a copy option (Figure 6-141).
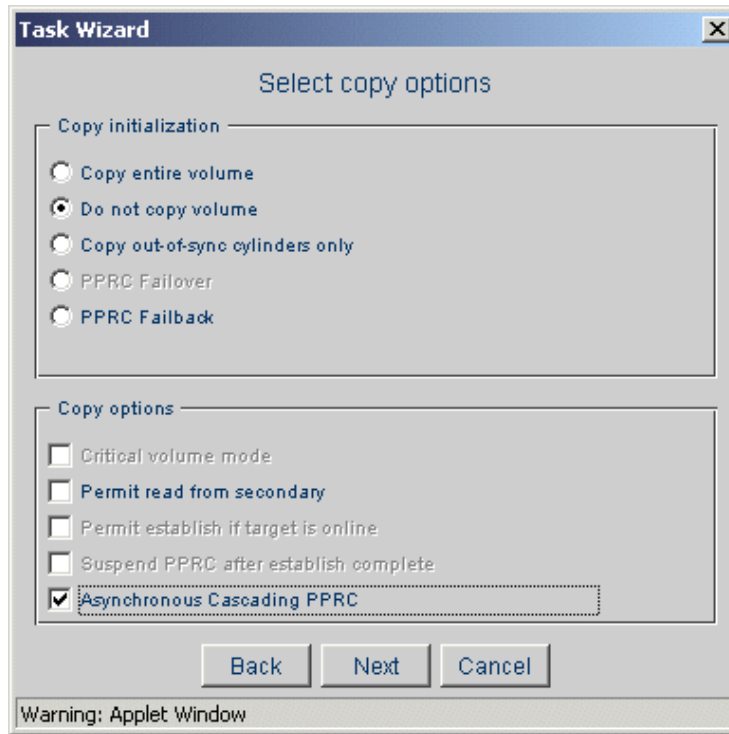
Figure 6-141   Select copy option panel: Asynchronous Cascading PPRC

The **Asynchronous Cascading PPRC** copy option allows a PPRC secondary volume to also be a PPRC primary volume for a different PPRC relationship. This means that the primary PPRC volume in the above defined task is eligible to be a target PPRC volume in a local PPRC pair relationship, from local site to intermediate site.

When the options are selected, click on **Next**. In the Define task panel (Figure 6-142 on page 279), decide whether you want to **Save**, **Run**, or **Cancel** the task.
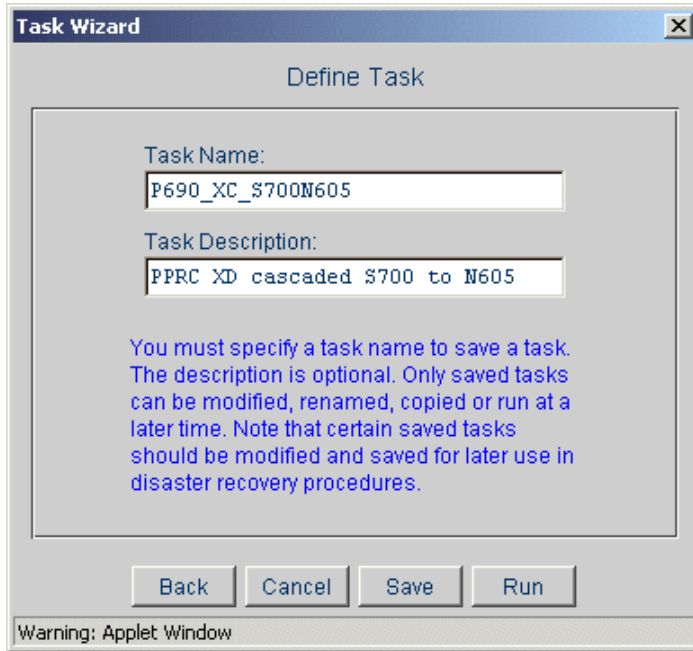
Figure 6-142   Task Wizard: Define Task

4. To check if the intermediate volume is eligible for cascading, select the volume and click on the **Information Panel** button at the left bottom of the Volumes panel. A Volume Information panel pops up (see Figure 6-143).
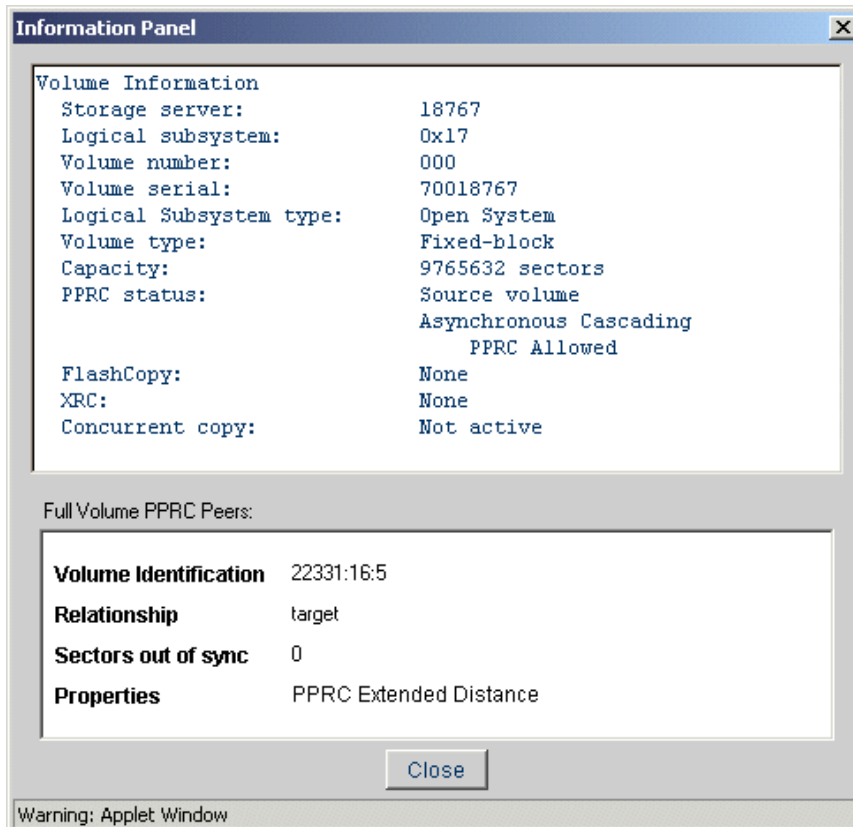
Figure 6-143   Volume Information Panel

Among other volume information, you can see if the cascading function for the selected volume is allowed (see PPRC status in Figure 6-143 on page 279).

5. As soon as the remote PPRC pair is established, create the PPRC synchronous relationship between the volume that resides in the local site with the volume in the intermediate site. Select the appropriate PPRC primary and secondary (with a right-click) volumes, then open the Task Wizard with another right-click. Then select **Establish PPRC Synchronous PPRC copy pair**, as shown in Figure 6-144.
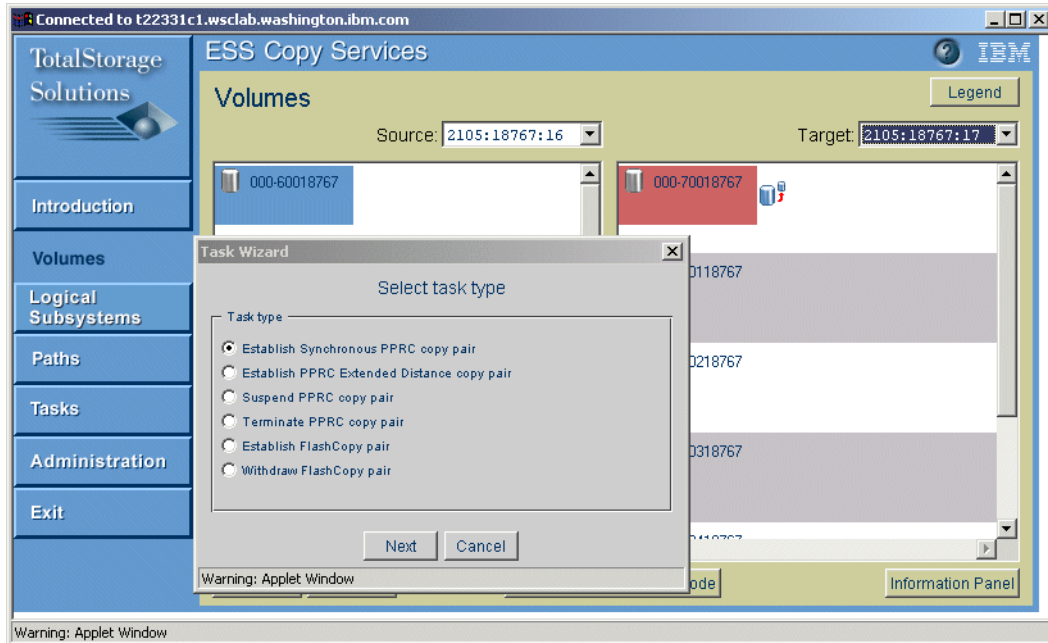
Figure 6-144   Intermediate volume as a PPRC primary and PPRC secondary

In this example, the local PPRC pair consists of the local and intermediate volumes in the same ESS (18767) but in different LSSs. The local PPRC primary volume is in LSS 0A and the local PPRC secondary volume in LSS 04.

In the Select Copy Option panel, select **Copy entire volume** as a copy initialization mode, as shown in Figure 6-145 on page 281.
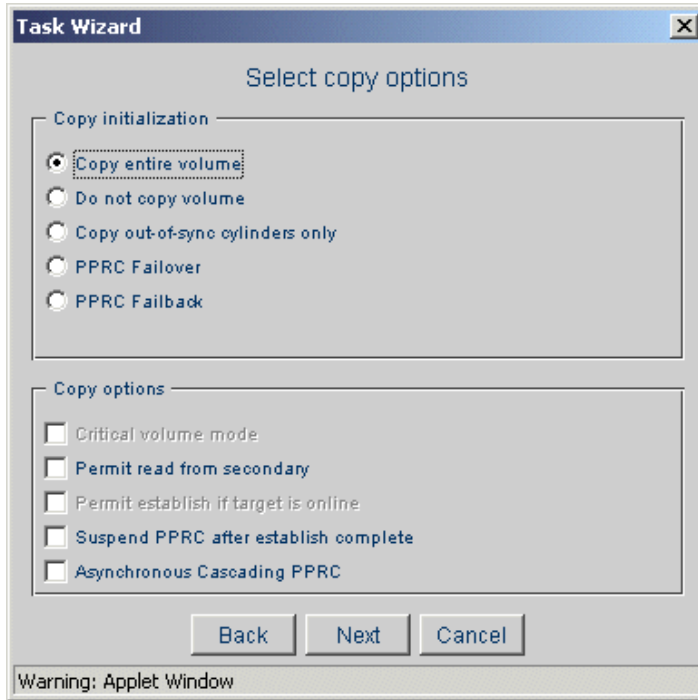
Figure 6-145   Select copy option panel: Copy entire volume

When the option is selected, click on **Next**.

In the Define task panel (Figure 6-143 on page 279), decide whether you want to **Save**, **Run**, or **Cancel** the task.
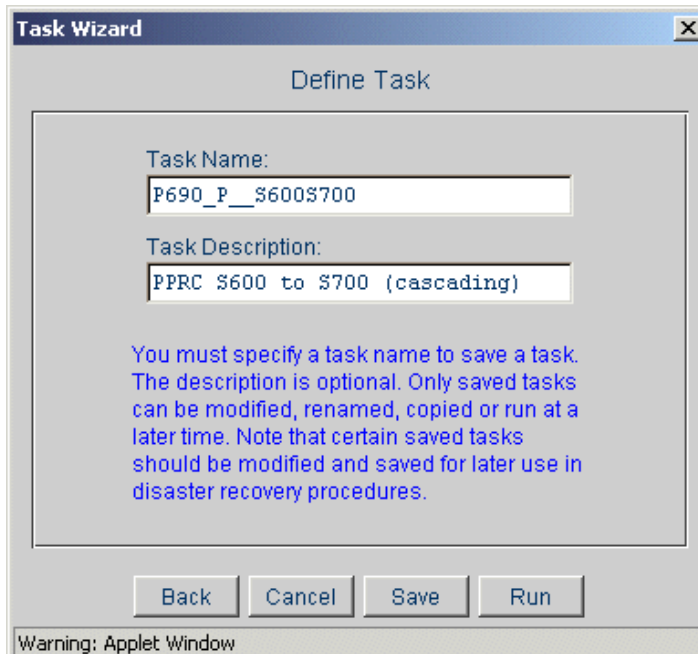
Figure 6-146   Define Task Wizard

6. Once the PPRC local pair is established, both the primary and secondary volumes are in a PPRC relationship, and in the copy pending state mode. The intermediate volume (ESS

18767, LSS 17, volume 70018767) has an associated icon combination, as shown in Figure 6-147.

Figure 6-147   Asynchronous Cascading PPRC: Local pair

In the example, the icons on the volume 70018767 indicate that the very same volume is in a PPRC-XD relationship as the primary volume of a remote PPRC pair (blue icon) and at the same time in a PPRC SYNC relationship as a secondary volume of the local PPRC pair, in copy pending mode (red icon volume, half empty). This is the typical Asynchronous Cascading PPRC configuration.

7. When the synchronization of the two volumes in the PPRC sync pair is complete, the icons will automatically change to show the duplex state of the volumes, as in Figure 6-148 on page 283.
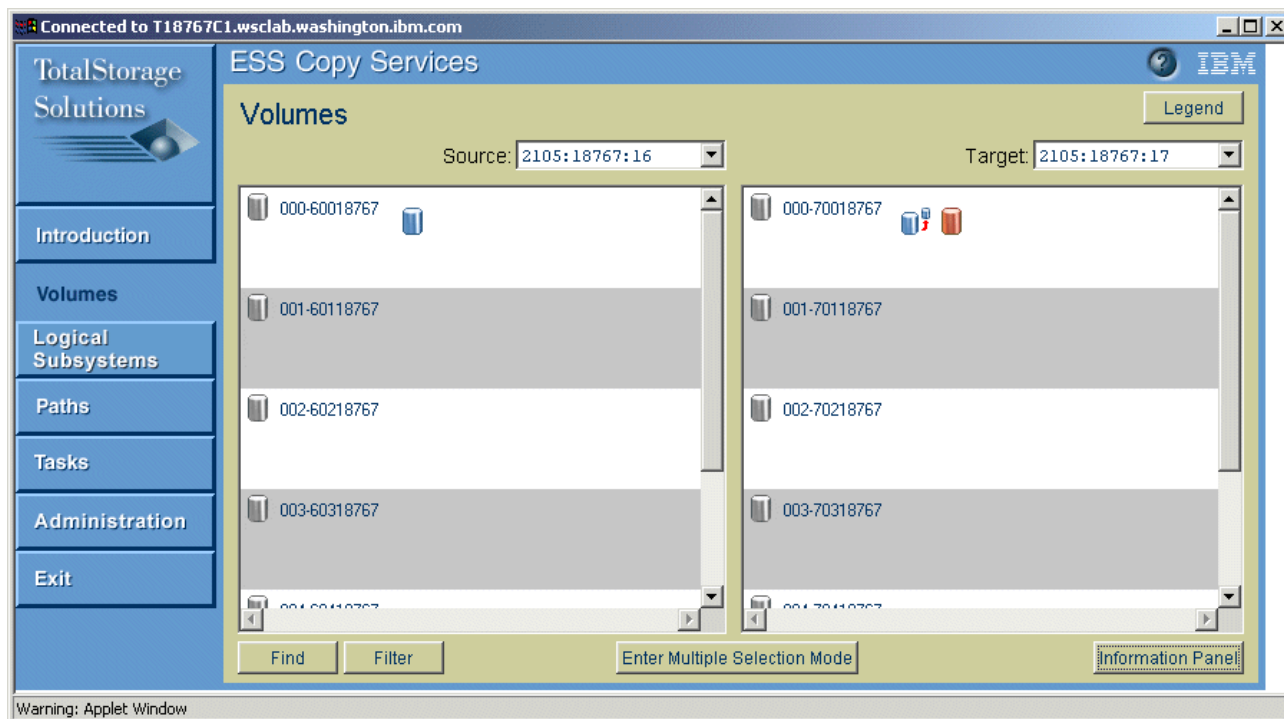
Figure 6-148   Asynchronous Cascading PPRC

6.12.16 PPRC Failover/Failback with Asynchronous Cascading PPRC

In this section, we present an example for setting up an asynchronous cascading PPRC environment and how to use Failover/Failback to switch application I/O from the local site (site A) to the intermediate site (site B) and back. This example is a detailed step-by-step procedure that follows the discussion in 10.8, "Failover/failback with Asynchronous Cascading PPRC" on page 429.

For the procedure, we use the ESS Copy Services Web user interface panels to:

1. Set up the asynchronous cascading PPRC environment using Fibre Channel links
2. Fail over the workload to site B
3. Fail back the workload to site A

Configuration setup and procedure steps

The following is the configuration used in our example:

► ESS A/B - Serial 24663, WWNN - 5005076300C09DEF
► ESS C - Serial 23953, WWNN - 5005076300C09B29

Figure 6-149 on page 284 illustrates the configuration used for this example.

We have a Fibre Channel link from SAID 000C on ESS-A/B to SAID 002C on the same ESS-A/B. This link is used to establish a PPRC path from LSS16 to LSS17 in ESS-A/B. This is used for the synchronous PPRC volumes.

We also have a Fibre Channel link from SAID 0008 on ESS-A/B to SAID 008C on ESS-C. This link is used to establish a PPRC path from LSS17 on ESS-A/B to LSS16 on ESS-C. This is used for our remote PPRC-XD volumes.

The configuration in our example is a two site asynchronous cascading PPRC configuration, as discussed in 4.4.3, "Asynchronous Cascading PPRC mode combinations" on page 68, where site A and site B volumes reside on the same physical ESS (designated as ESS-A/B), while site C volumes reside at a remote ESS (designated as ESS-C).

In our example, we use one volume in each LSS. The volume in LSS16 on ESS-A/B is referred to as the site A volume, the volume in LSS17 on ESS-A/B is referred to as the site B volume, and the volume in LSS16 on ESS-C is referred to as the site C volume.
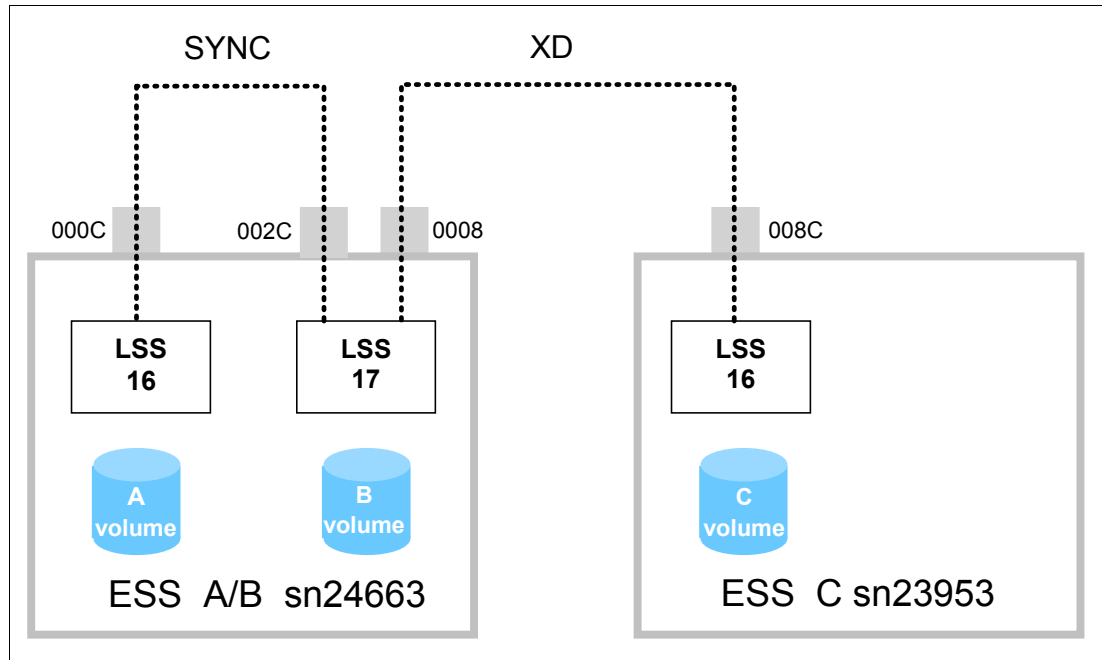
Figure 6-149   Asynchronous Cascading PPRC configuration used in this example

The procedure steps are the following:

► Establishing the local and remote paths
► Establishing the remote and local PPRC volume pairs
► PPRC Failover task to site B volumes
► Suspend the B → C PPRC sessions
► Establish PPRC paths B → A
► PPRC Failback task on site B volumes
► Establish PPRC paths A → B
► PPRC Failover task to site A volumes
► PPRC Failback to site A volumes
► Re-establish B → C sessions

Each of these steps are discussed in detail in the following sections.

Establishing the local and remote paths

The first step for setting up the asynchronous cascading PPRC environment is to establish the paths from LSS16 to LSS17 in ESS A/B (local path, site A to site B) and from LSS17 in ESS-A/B to LSS16 in ESS-C (remote path, site B to site C). Figure 6-150 on page 285 shows the Paths panel once the local path is defined (note the blue asterisks below the link between LSS16 to LSS17 within ESS A/B).
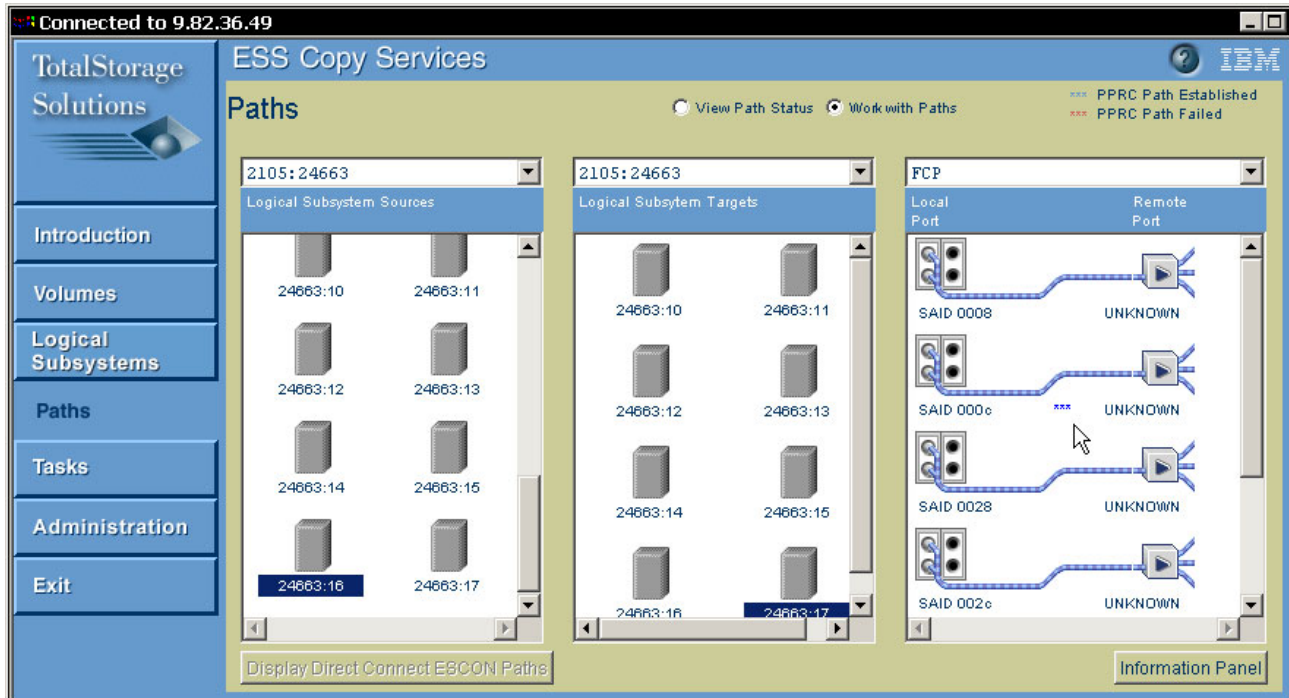
Figure 6-150   Path successfully defined between site A and site B LSSs

Figure 6-151 shows the Paths panel once the remote path is defined (note the blue asterisks below the link between LSS17 in ESS-A/B to LSS16 within ESS-C).
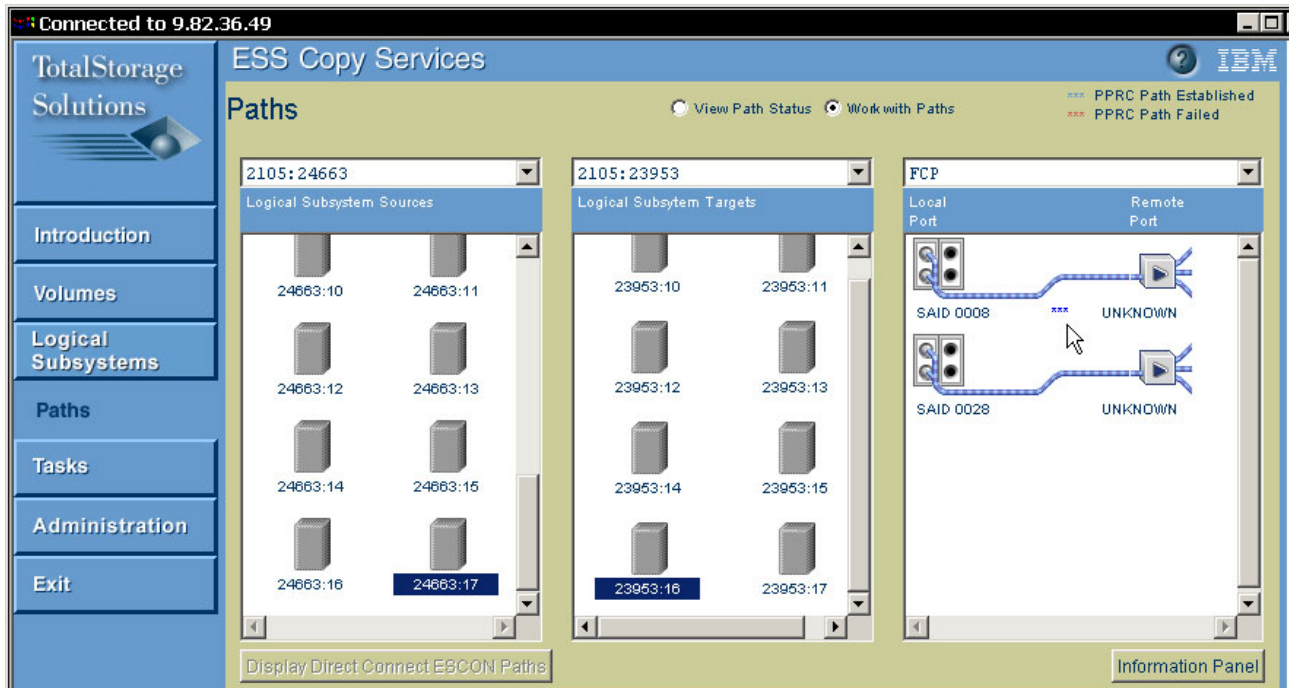
Figure 6-151   Path successfully defined between site B and site C LSSs

Establishing the remote and local PPRC volume pairs

Now we establish the PPRC remote volume pairs (one volume pair in our example) from LSS17 on ESS-A/B (intermediate site, primary XD volume) and LSS16 on ESS-C (remote

site, secondary XD volume). This is our remote PPRC-XD pair, and Figure 6-152 shows the options chosen when establishing the pair.
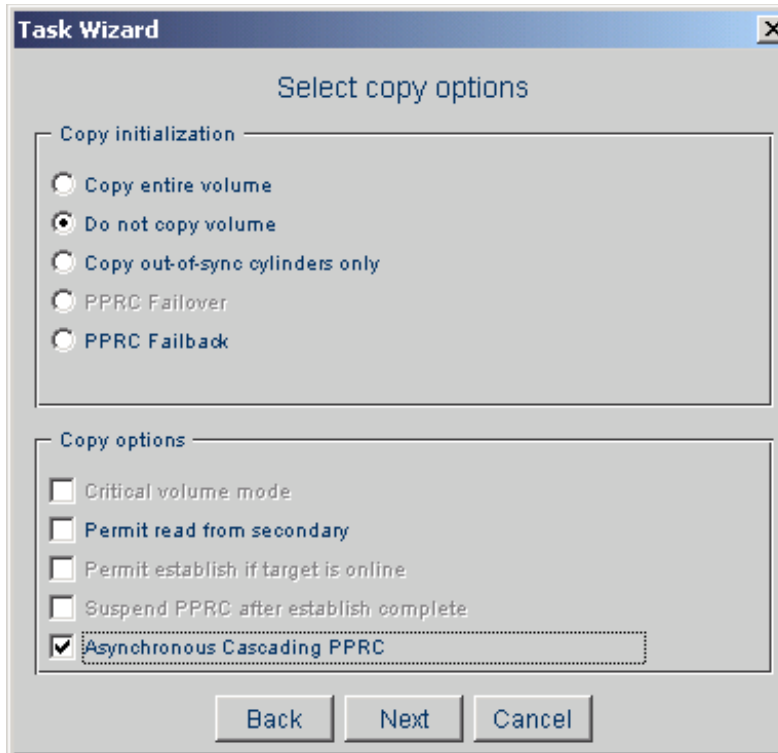
Figure 6-152   Options to create the XD pair with

Figure 6-153 shows the remote pair, once it has been established, as displayed in the Volumes panel of the ESS Copy Services Web user interface.
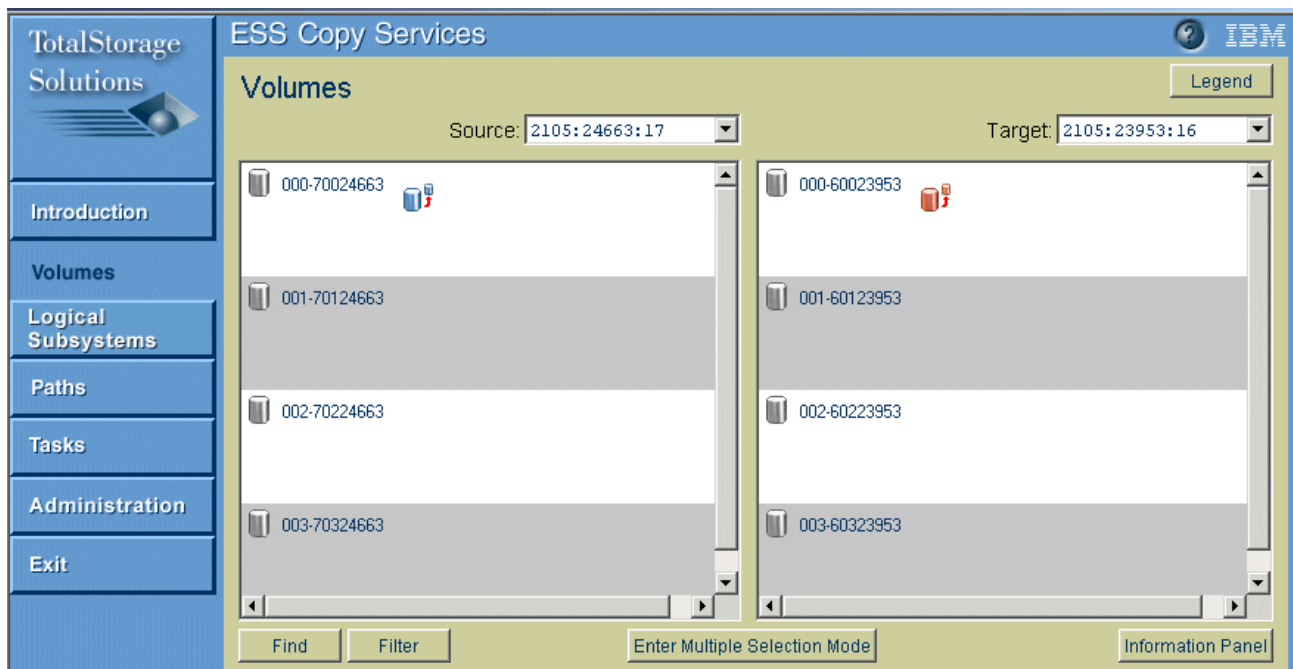
Figure 6-153   Remote PPRC-XD volume pair established

We then establish the PPRC local volume pairs (one pair in our example) from LSS16 on ESS-A/B (local site, primary PPRC-SYNC volume) and LSS17, also on ESS-A/B (intermediate site, secondary PPRC-SYNC volume). This is our local PPRC-SYNC pair, and Figure 6-154 shows the options chosen when establishing this PPRC volume pair.
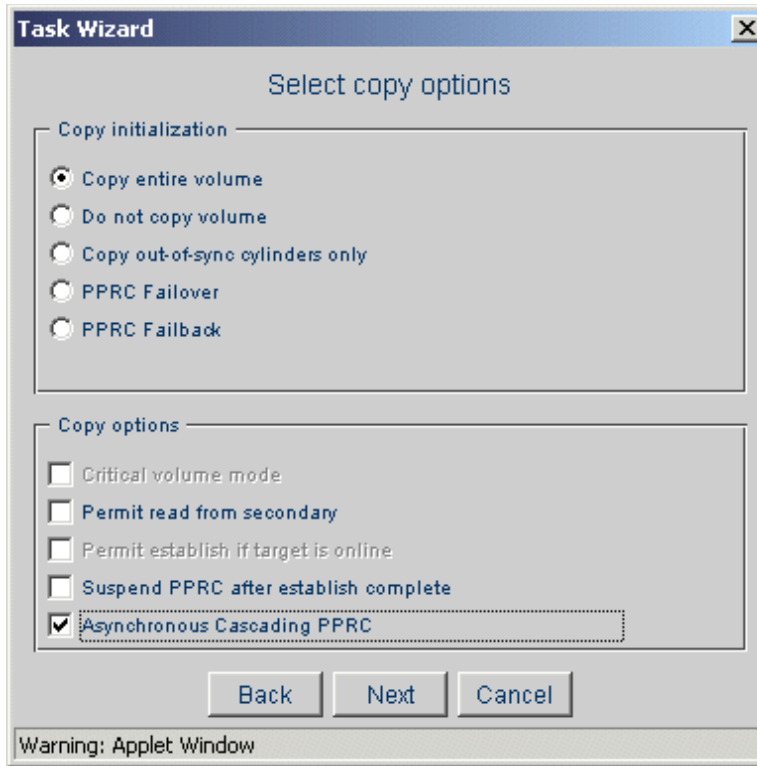
Figure 6-154   Options to create the local PPRC-SYNC pair

Figure 6-155 on page 288 shows the local PPRC-SYNC pair when it is in the process of initial synchronization (duplex pending).
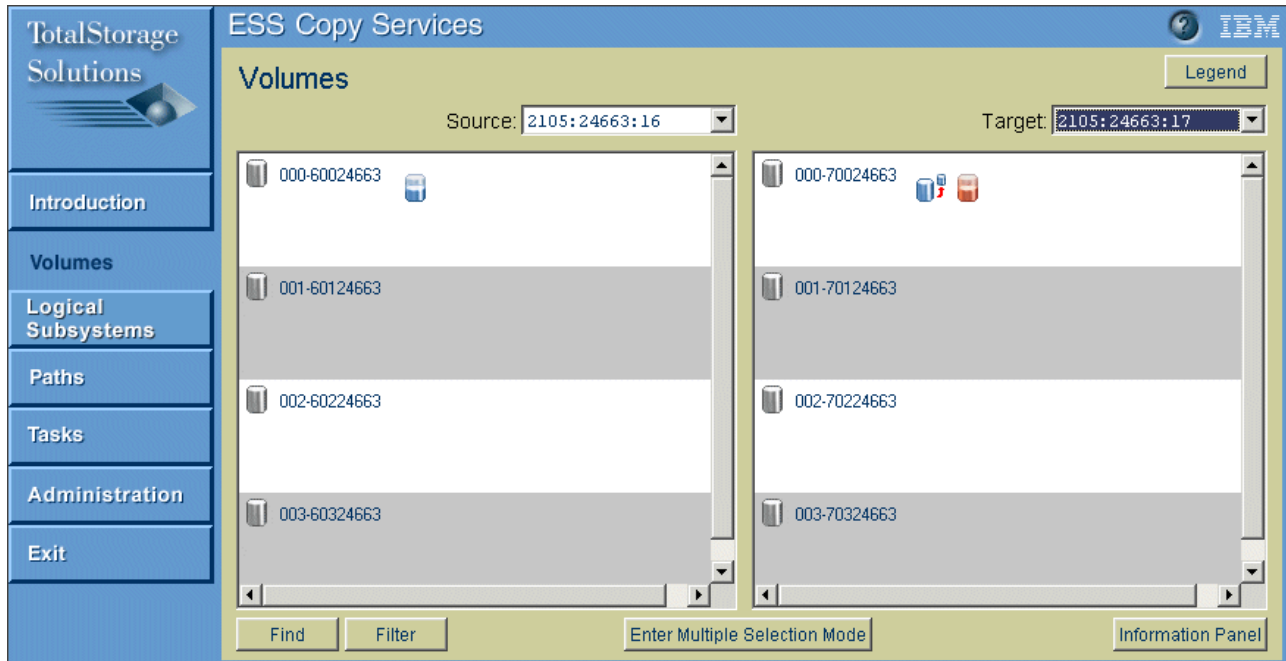
Figure 6-155   Local PPRC-SYNC pair during initial synchronization (duplex pending)

Once the local pair reaches full synchronization (duplex state), then the asynchronous cascading PPRC configuration is set up and running.

PPRC Failover task to site B volumes

With the asynchronous cascading PPRC configuration setup and running, now we can start the fail-over procedure at any moment, so we can stop application updates on site A and resume them on site B volumes. This will leave site A available for any required maintenance while applications run using site B volumes.

For this, we run a PPRC Failover task onto site B volumes. Figure 6-156 on page 289 shows the options selected for this task.
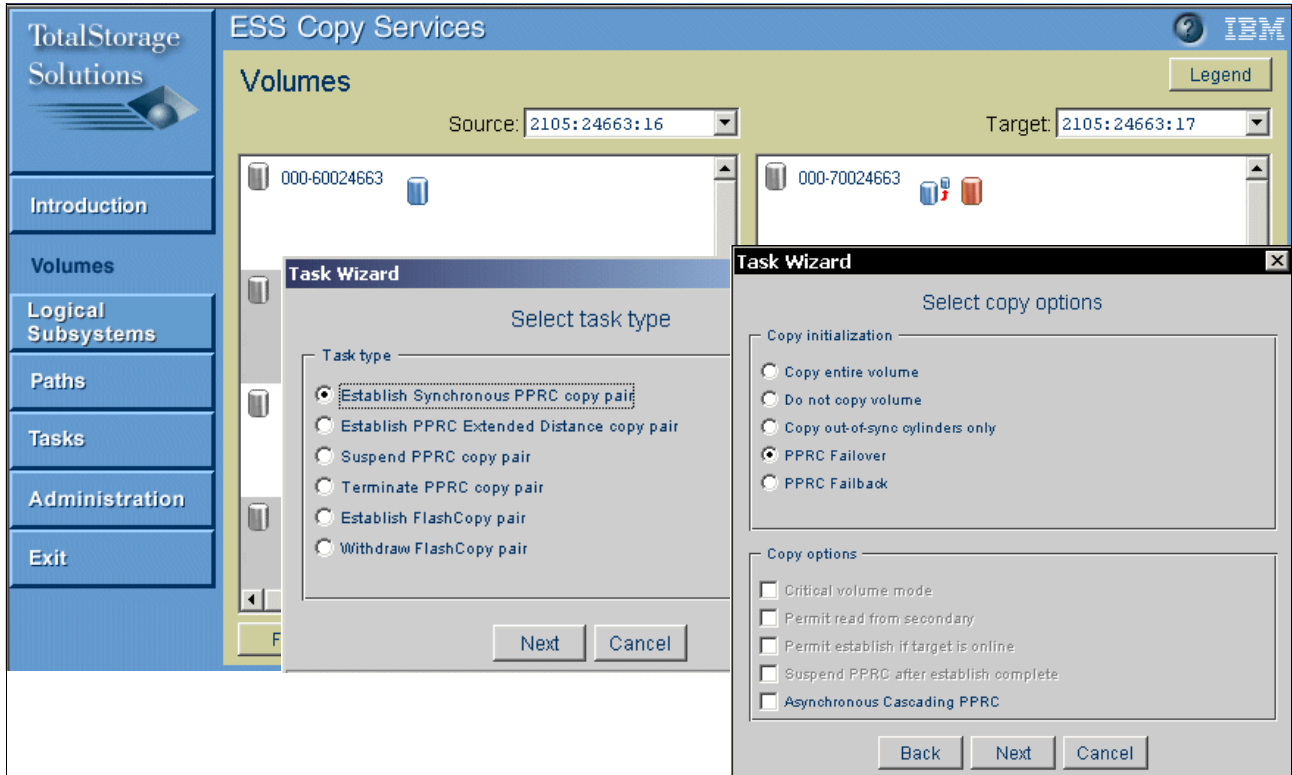
Figure 6-156   Doing a PPRC Failover on site B volumes

Figure 6-157 illustrates the state of the local pair of volumes (site A → site B) once the PPRC Failover is completed.
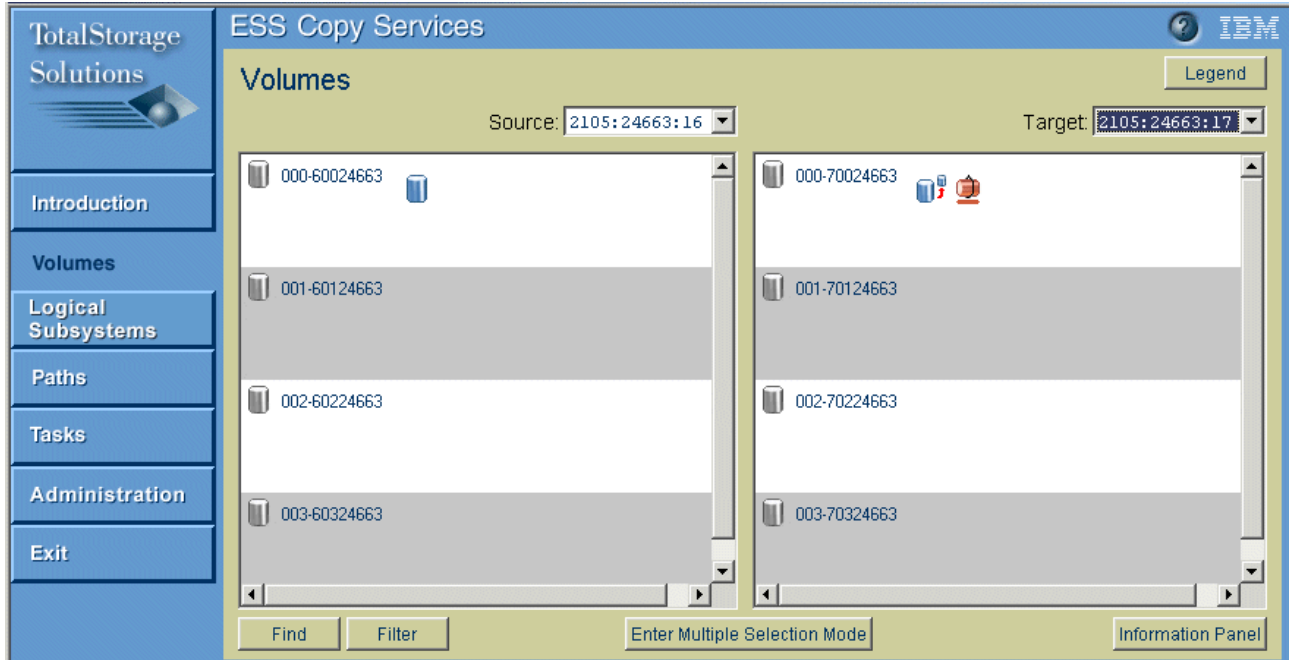
Figure 6-157   PPRC Failover to site B completed

At this point, application I/O can be swapped to site B volumes.

Figure 6-158 shows the Information Panel for the intermediate volume of our example (the volume in LSS17 on ESS-A/B); note the "Secondary Host Writes Enabled" status. This allows a volume that is a PPRC target to accept host I/O.
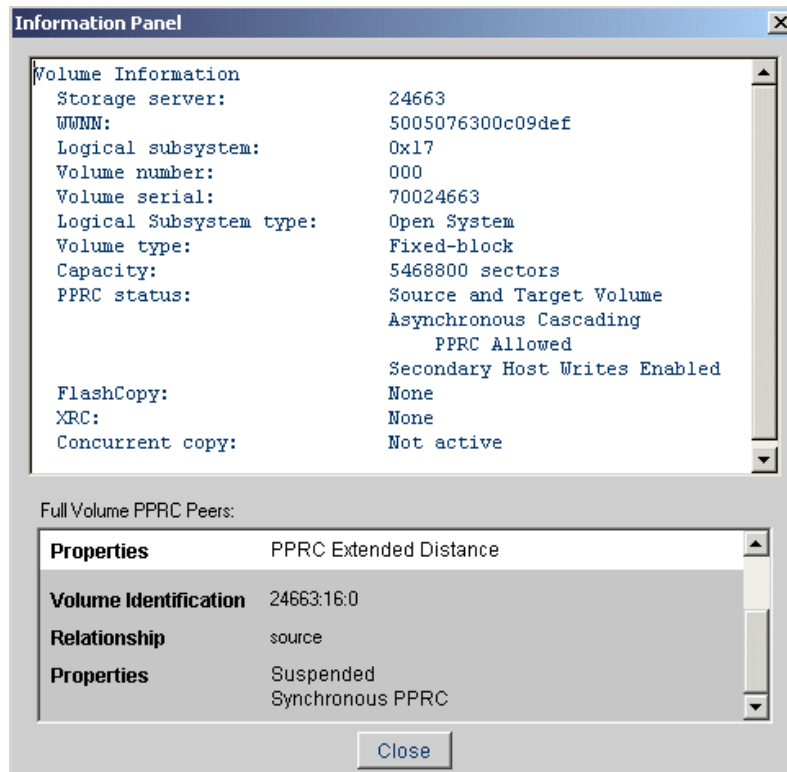
Figure 6-158   Information Panel for site B volume after failover

In these conditions, the production applications can run using site B volumes. Site B volumes are still secondary volumes for the local PPRC-SYNC relationship and at the same time are primary for the remote PPRC-XD relationship. And because the "Secondary Host Writes Enabled" status is on, they can receive application updates, even if they are secondary volumes.

Once the required activities are completed on site A, we can resume application processing back on site A volumes. The following steps discuss the procedure to switch back to site A (failback to site A).

Suspend the B -> C PPRC sessions

This is the first step in moving application I/O back to site A. The B → C PPRC-XD sessions (one session in our example) must be suspended, as we need to establish a B → A PPRC session in order to fail back, and B volumes cannot be the source of two sessions at the same time.
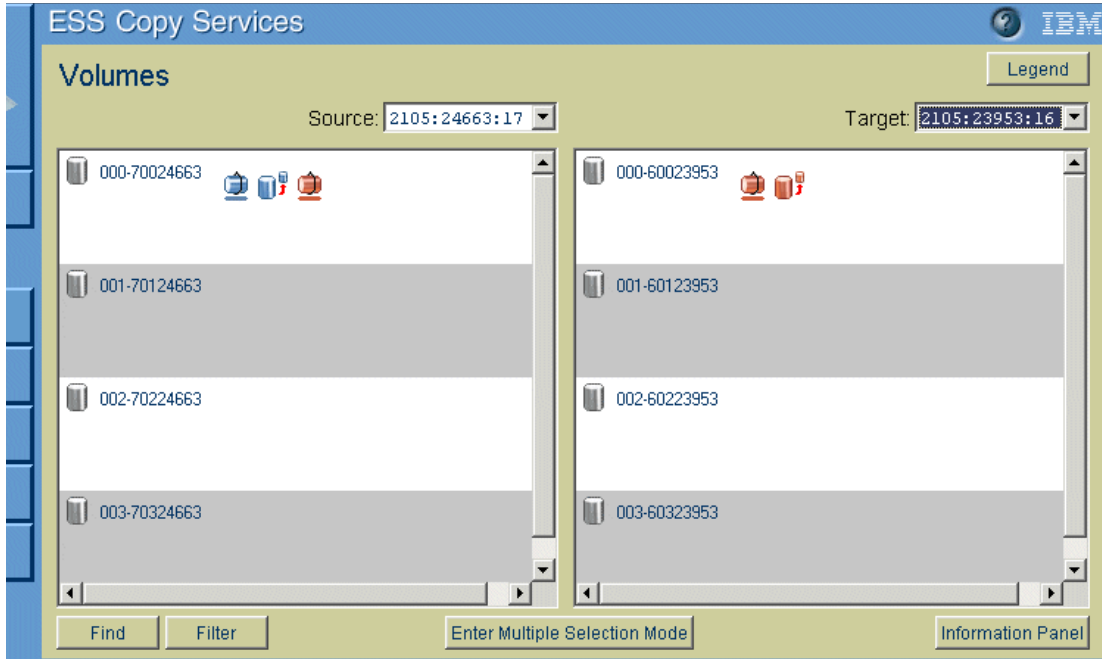
Figure 6-159   Display of remote PPRC-XD pair after suspending

Figure 6-159 shows the status of the PPRC-XD remote (site B <--> site C) volume pair after suspending. The status of the PPRC-SYNC source volume in LSS16 on ESS-A/B (site A, not shown in the Figure 6-159) has not changed.

Establish PPRC paths B -> A

We now establish the site B → site A PPRC paths (one in our example) to allow the site A volumes to be synchronized with the site B volumes. Figure 6-160 shows the path once it has been established, from site B to site A in preparation for the PPRC Failback task.
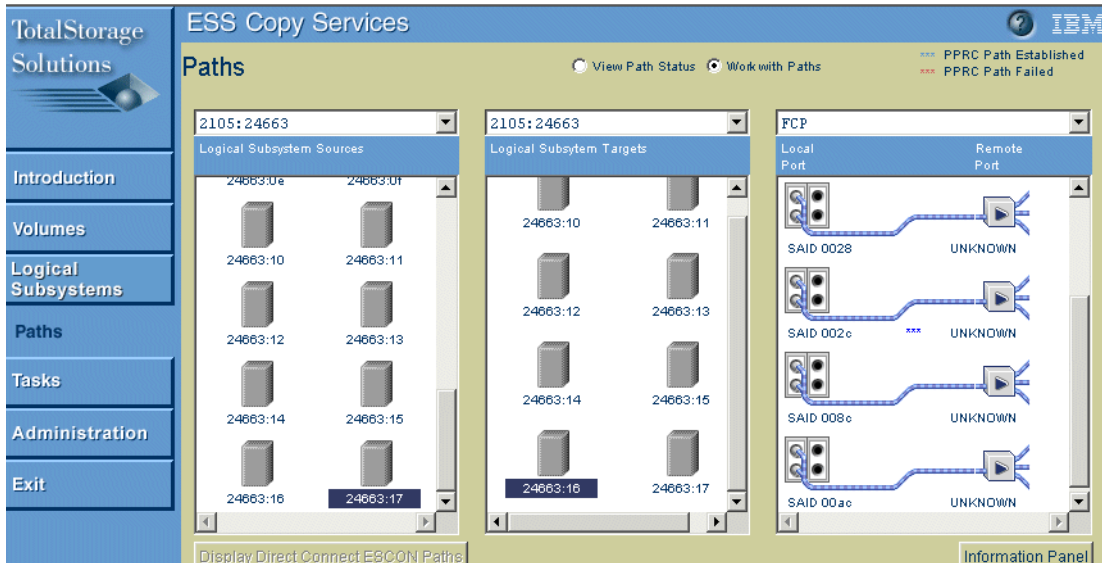
Figure 6-160   Site B to site A path established

PPRC Failback task on site B volumes

The PPRC Failback processing will synchronize the site A volumes with the site B volumes (one volume pair in our example). It will also set the "trusted resync" flag to allow a later re-synchronization of site C volumes with site B volumes, as these pairs currently remain suspended.

Figure 6-161 shows (in the background) the state of the local volume pair before the PPRC Failback operation and it also shows (in the foreground) the Task Wizard options selected for the Failback task (note that in Figure 6-161 the site A volume is on the right side of the Volumes panel, and the site B volume is on the left side of the panel).
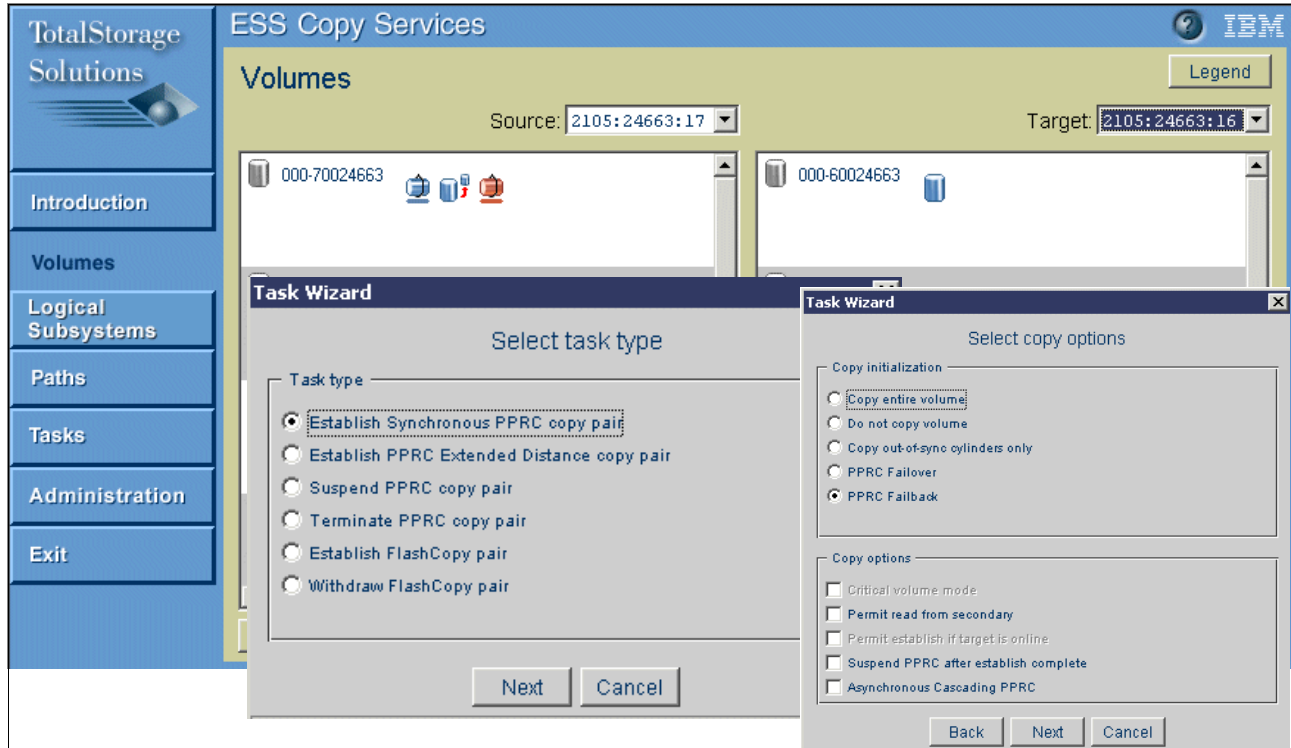
Figure 6-161   PPRC Failback on site B volume: Task options

Figure 6-162 on page 293 shows the status of the local PPRC volume pair (site B volume on the left side of the Volumes panel, site A volume on the right) while the PPRC Failback task is progressing. Notice that the volumes are re-synchronizing (disk icons half shaded). Also notice the "Trusted Primed for Resynchronization" icon on site B volume; this will allow the later re-synchronization of the remote PPRC-XD pair (site B to site C).
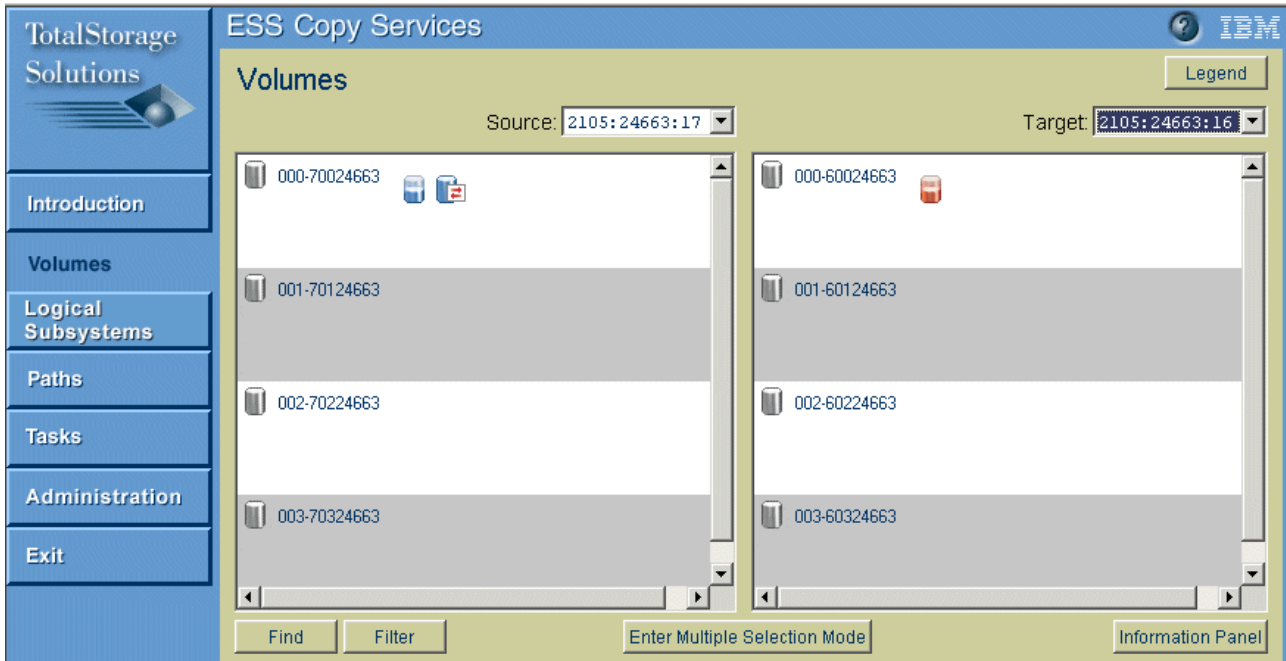
Figure 6-162   Volumes in LSS17 and LSS16 on ESSA while PPRC Failback is proceeding

Figure 6-163 shows the volumes state once the PPRC Failback has completed; the volume pairs are now synchronized (disk icons fully shaded).
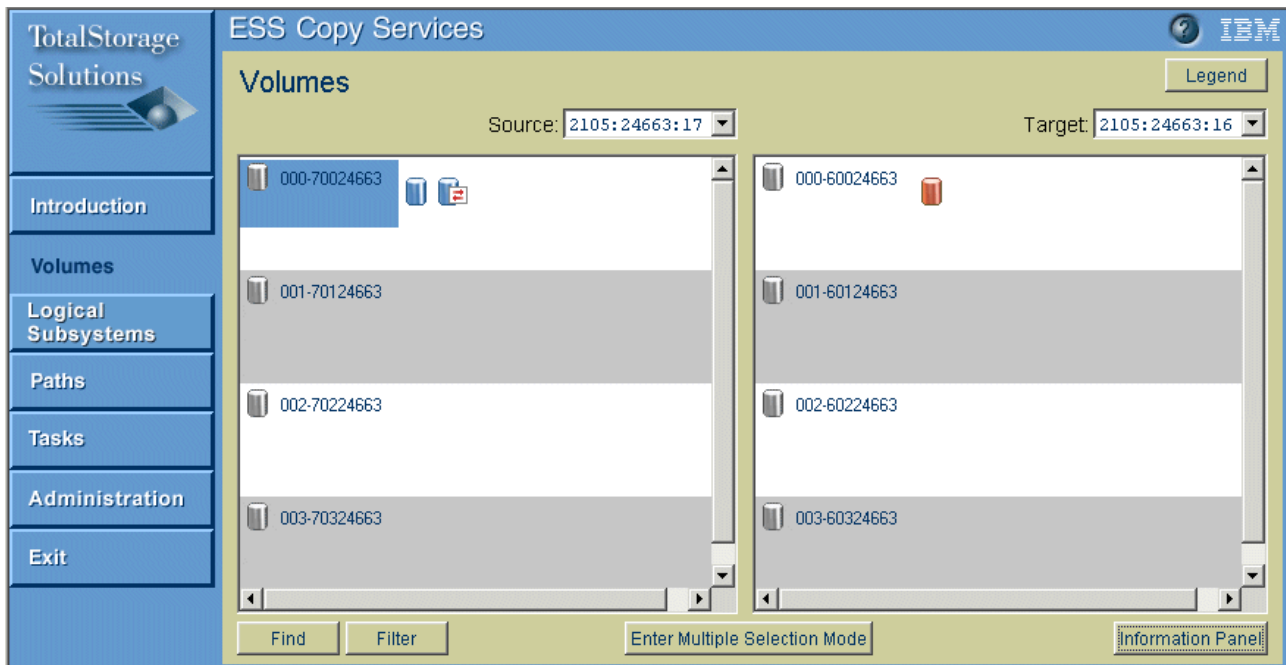
Figure 6-163   Volumes in LSS17 and LSS16 on ESSA after PPRC Failback on B has completed

Figure 6-164 on page 294 shows the information panel for the site B volume (in LSS17 on ESS-A/B); note the "Trusted Primed for Resynchronization" flag.
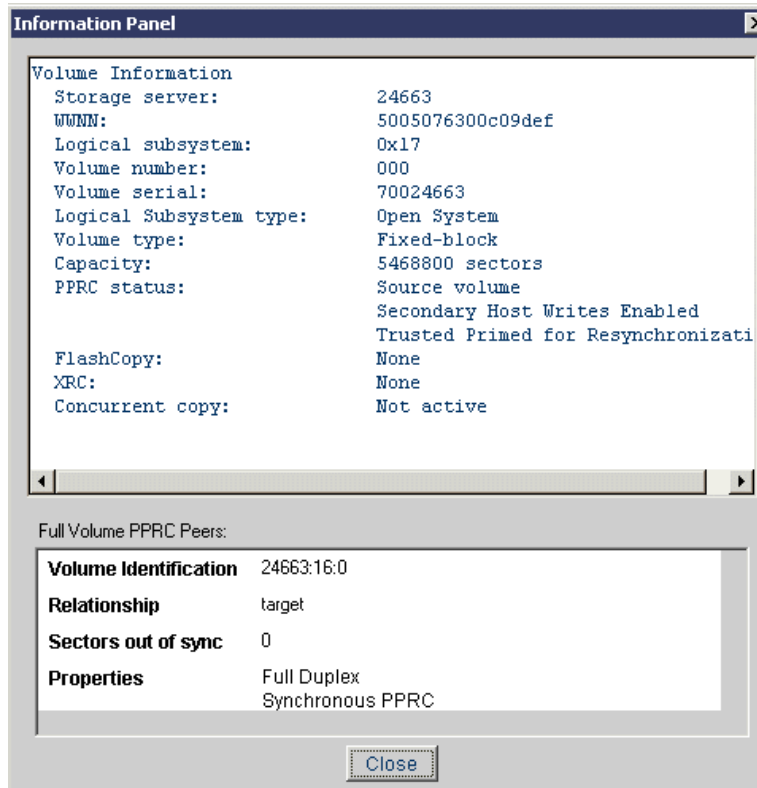
Figure 6-164   Information Panel for site B volume after PPRC Failback

Site A volumes are now synchronized with site B volumes. We are ready to stop application I/O processing on site B and switch back to site A. The following sections describe the steps needed to switch back to site A.

Establish PPRC paths A -> B

We can now establish the PPRC paths from site A → site B (one path in our example). Figure 6-165 on page 295 shows the path successfully established (blue asterisks) from site A LSS 16 to site B LSS 17 (both in the same ESS-A/B).
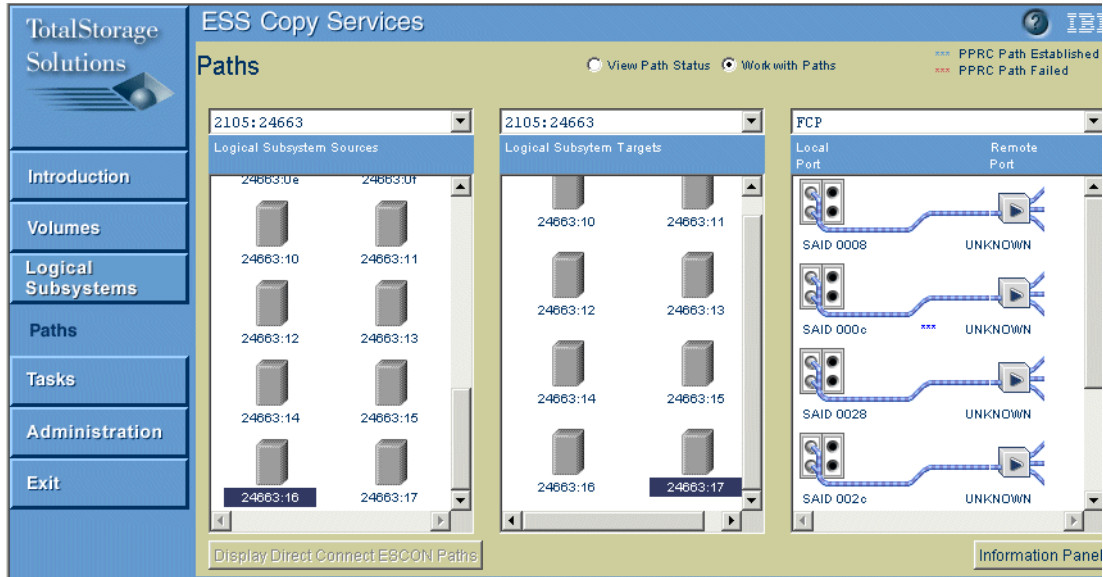
Figure 6-165   Path successfully established from site A LSS16 to site B LSS17

PPRC Failover task to site A volumes

We can execute the PPRC Failover task to the site A volumes (one volume pair in our example). This will terminate the PPRC B $\rightarrow$ A relationship and establish the A $\rightarrow$ B relationship.

Figure 6-166 shows the state of the volume pair before the PPRC Failover task is executed.
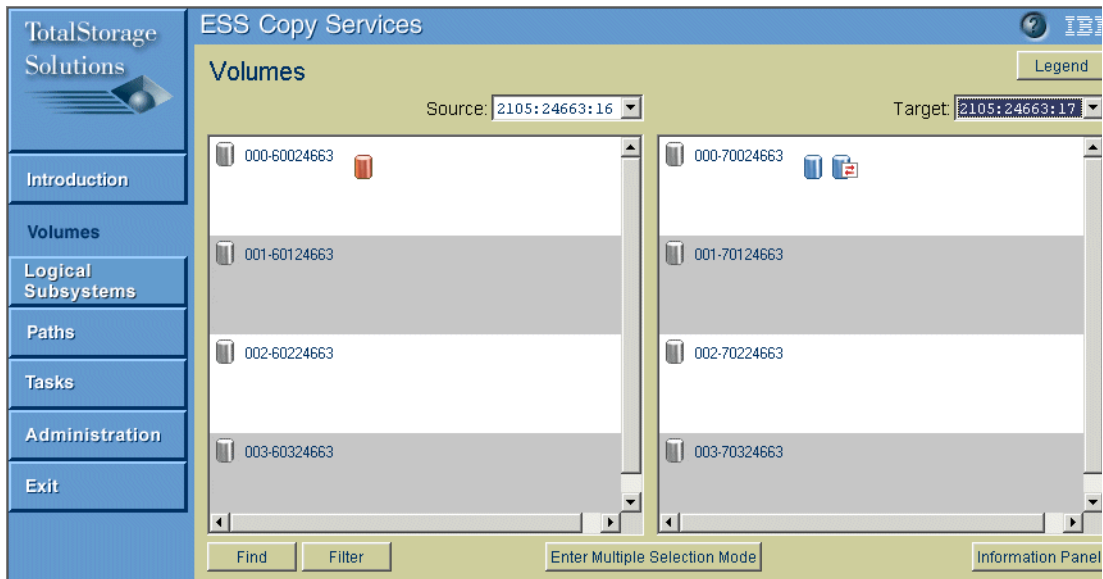
Figure 6-166   Local PPRC volume pair: State before PPRC Failover

Figure 6-167 on page 296 shows the state of the local volume pair after the PPRC Failover task has completed.
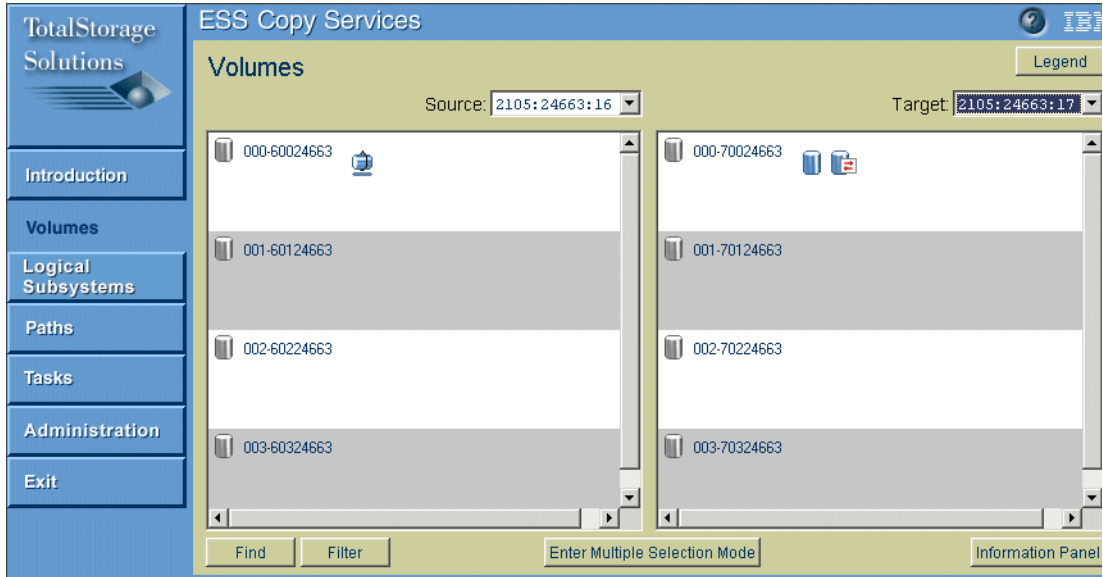
Figure 6-167   Local PPRC volume pair: State after PPRC Failover task on volume A

PPRC Failback to site A volumes

The PPRC Failback task on site A volumes will complete the restore of the A → B PPRC session and reset the "Enable Secondary Host Writes" flag on the B volumes. Figure 6-168 shows the volume pair before the task is initiated (background) and the copy options used for the task.
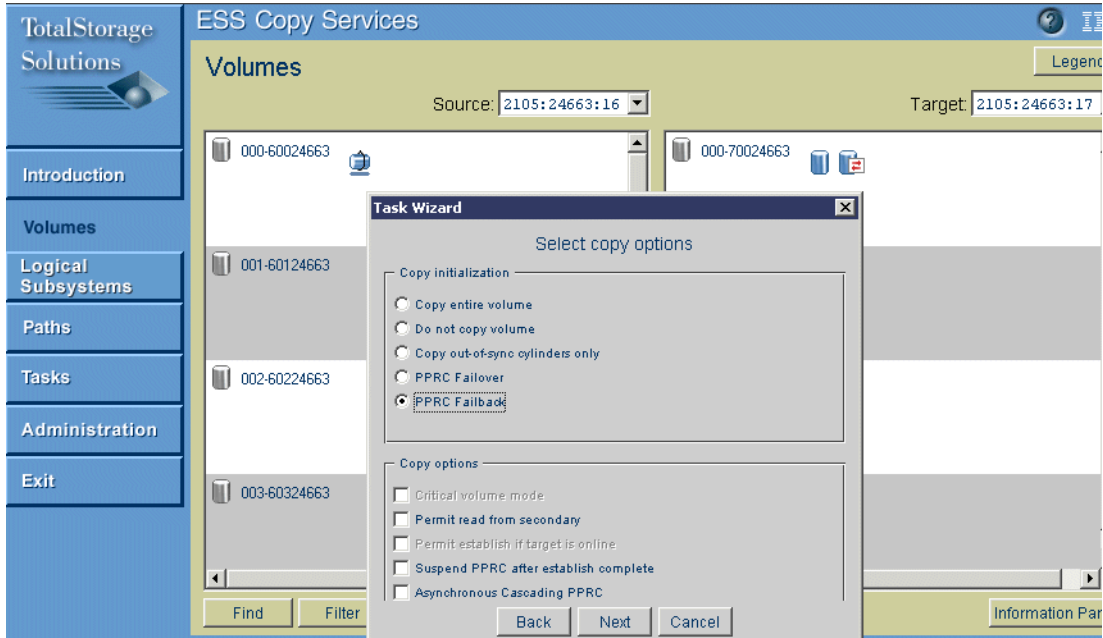
Figure 6-168   PPRC Failback task: Copy options

Figure 6-169 on page 297 shows the state of the volumes after the PPRC Failback task to site A volumes is completed. Site B volume is now the PPRC secondary (red disk icon).
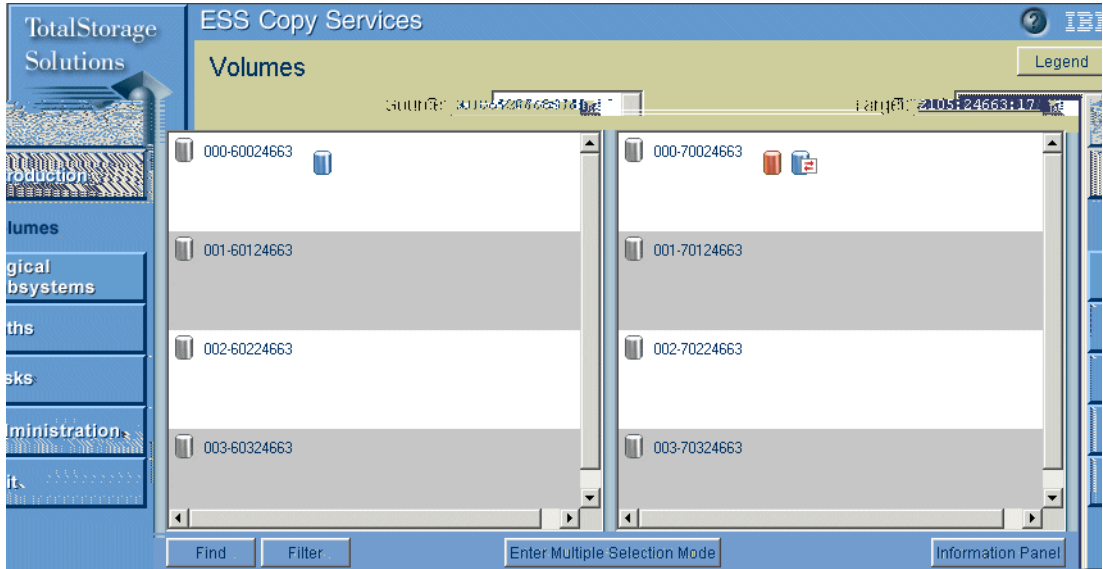
Figure 6-169   PPRC Failback task completed

Figure 6-170 shows the Information Panel for the volume in LSS17 on ESS-A/B (the site B volume). Notice that the "Trusted Primed for Resynchronization" flag is still set, but the "Enable Secondary Host Writes" has been reset.
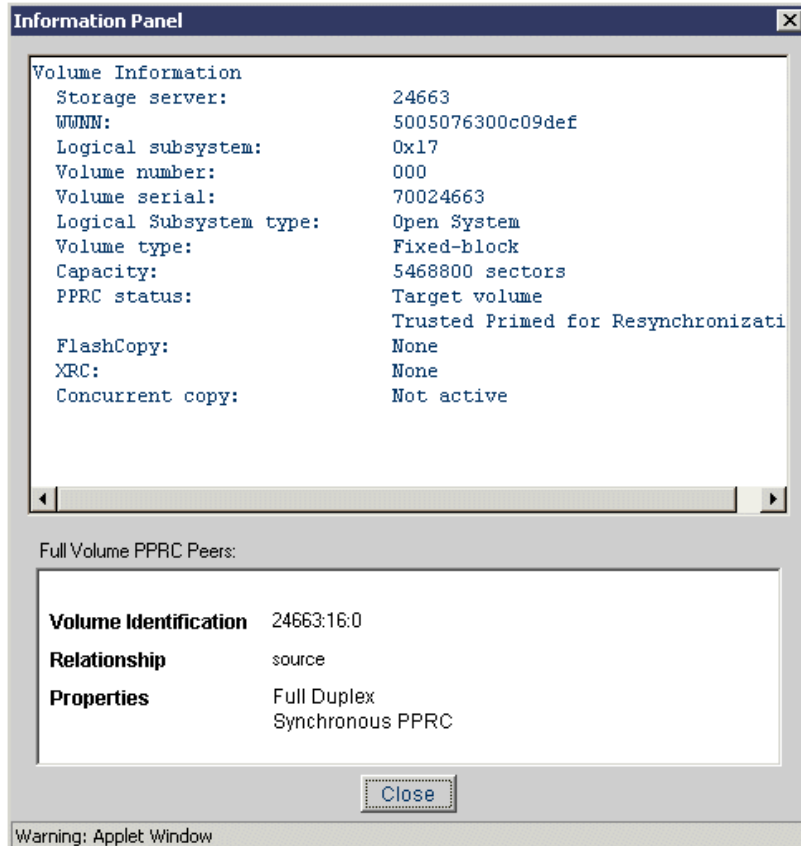
Figure 6-170   Information Panel for the site B volume after Failback to site A

Re-establish B -> C sessions

Now that the site A → site B PPRC-SYNC session has been restored, we can re-establish the site B → site C PPRC-XD pairs using the out-of-sync option. This resynchronization is allowed, as the site B volumes have the "Trusted Primed for Resynchronization" flag set. Figure 6-171 shows the remote volume pair (site B and site C volumes) before the resynchronization task.
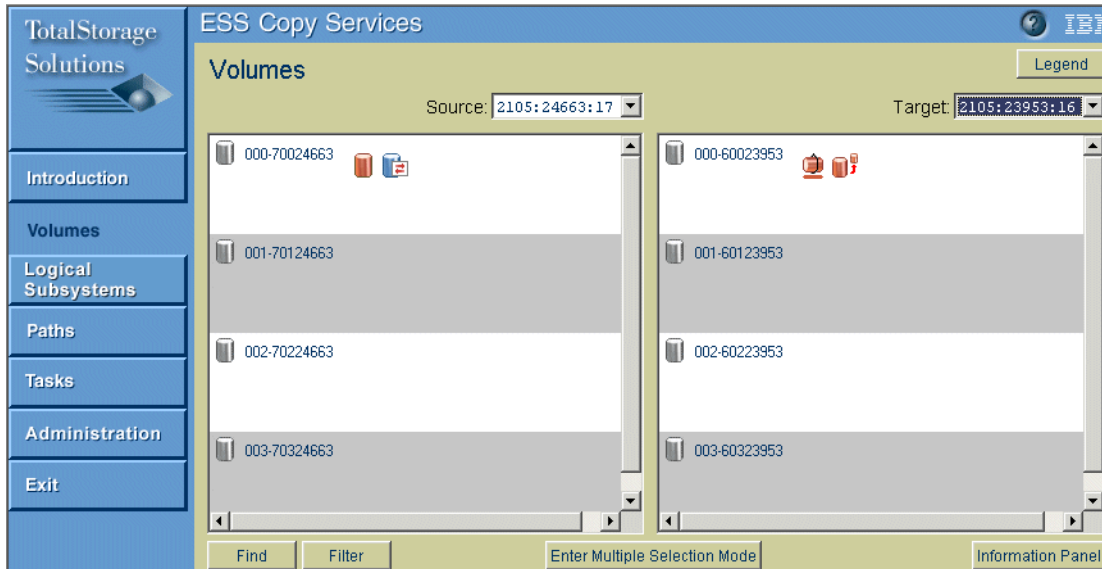
Figure 6-171   Remote volume pair before resynchronization (note trusted primed for resynch)

We re-establish the PPRC-XD relationship between the B and C volumes, using the options shown in Figure 6-172.
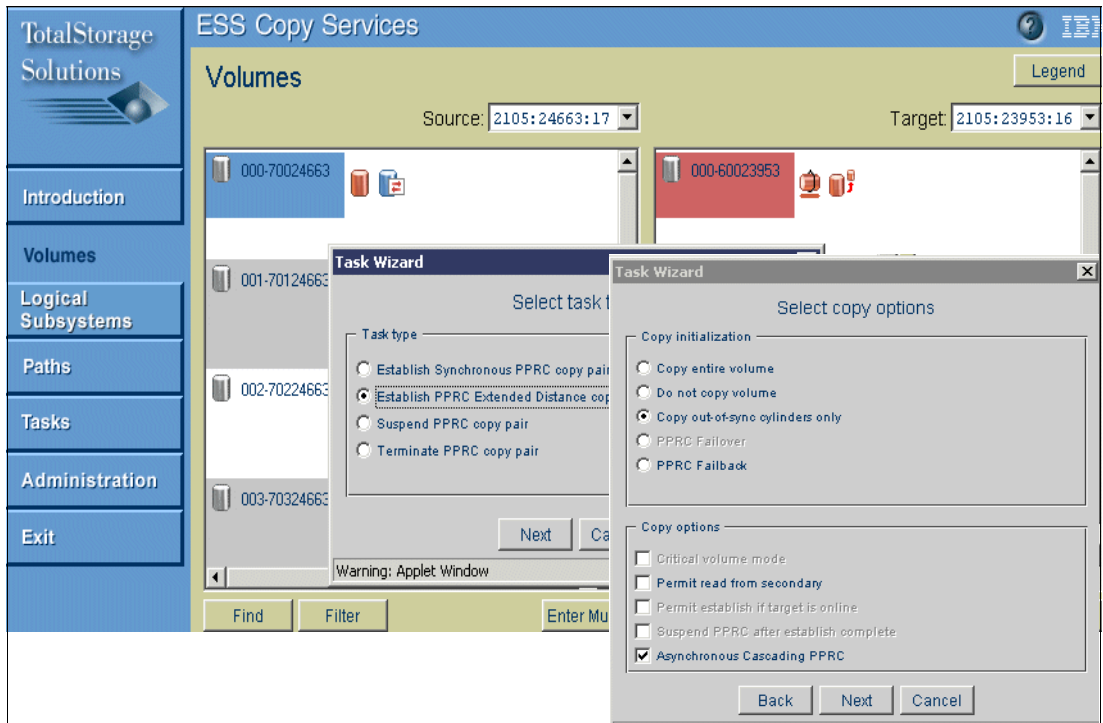
Figure 6-172   Re-synchronization between B and C volumes

Once the resynchronization task completes, the "Trusted Primed for Resynchronization" flag is reset. Figure 6-173 shows the information panel for the site C volume once the re-synchronization has completed.
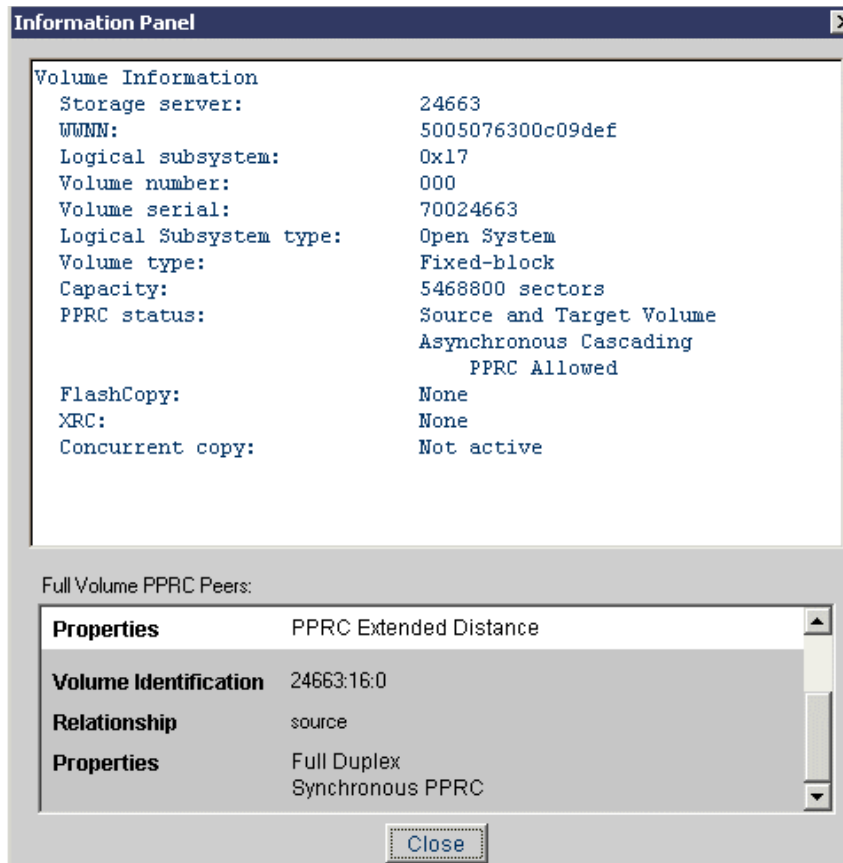
Figure 6-173   Information Panel for the site C volume after resynchronization

The completion of this last step restores the asynchronous cascading PPRC configuration to its fully operative condition.

6.13 Summary of the benefits of the WUI

Starting with LIC Version 2.2.0, the WUI provides an enhanced Copy Services Web User Interface, compared with the WUI provided with LIC versions prior to 2.2.0.

The main differences are:

► **New Tools panel with new options**

   – The user can define the active servers of the Copy Services domain ("Defining the active servers" on page 173).

   – The user can define the clients in the Copy Services domain ("Defining the clients" on page 174).

   – The user can perform a Domain Wide Reset ("Domain Wide Reset" on page 176).

► **New way to start the Copy Services WUI**

Since LIC 2.2.0 or higher provides a new dual-active mode, you can now choose from which server you want to start the WUI (6.4, "Connecting to the Copy Services Web User Interface" on page 176).

► **New volume Panels in the WUI**

New volume presentations with new icons. The volumes are now presented in rows. When the volume is involved in one or more copy services relationships, multiple icons will identify each relationship and the state of the volume (see 6.5, "Volume panel" on page 179).

► **New Option for PPRC pairs**

To benefit from the new option, you must have the PPRC V2 feature installed. The new option is Asynchronous Cascading PPRC (see 6.12.15, "Establishing an Asynchronous Cascading PPRC" on page 273).

► **New Option for FlashCopy pairs**

To benefit from the new option, you must have the FlashCopy V2 feature installed. The new options are:

– Source and target volumes can be in a different LSS in the same ESS.

– Multiple FlashCopy using the same Source volume (6.11.4, "Multiple FlashCopy using a single source volume" on page 214).

– Incremental FlashCopy (6.11.6, "Incremental FlashCopy" on page 216).

– Reversing a FlashCopy (6.11.7, "Reversing a FlashCopy relationship" on page 222).

– Creating FlashCopy consistency group (6.11.8, "Creating a FlashCopy consistency group" on page 224).

– Inband Commands (6.11.11, "Inband FlashCopy" on page 229).

With the availability of LIC level 2.3.0. all the previous benefits apply, plus the ability to configure PPRC using Fibre Channel links and also having double the ESSs in a Copy Services domain.

7

ESS Command Line Interface

The Enterprise Storage Server (ESS) provides two Command Line Interfaces (CLIs) for different host platforms:

► **ESS Copy Services CLI:** The ESS Copy Services CLI allows you to communicate with the ESS Copy Services server from the host's command line. An example would be to automate a task like doing a FlashCopy by invoking the Copy Services command with customized scripts.

► **ESS Storage Management CLI:** The ESS Storage Management CLI allows you to create automated routines to manage the storage and to change the logical configuration of the ESS.

In this chapter, we describe how to use the ESS Copy Services CLI and we describe the most useful commands regarding the management of the ESS Copy Services. We also present the ESS Storage Management CLI.

7.1 ESS Copy Services CLI

The Copy Services Command Line Interface (CLI) is an interface that allows you to communicate with the ESS Copy Services server from the host's command line.

You can use the Copy Services CLI to perform several tasks:

► Execute one or more ESS Copy Services tasks that you have defined using the ESS Copy Services Web-based interface.

► Display the mapping of a host disk name to a 2105 volume serial number.

► Create, update, or remove a list view of volumes associated with a particular host in the ESS Copy Services server; you can see the same view in the Volumes panel of the ESS Copy Services Web-based interface.

► Query the status of one or more volumes.

► Access the ESS Copy Services server tasks, as defined and saved within the ESS Copy Services Web-based interface, and determine whether the volumes defined in a given task have completed the initial PPRC synchronization.

► Determine whether you can successfully connect to the ESS Copy Services server.

The ESS CLI is available for selected open servers:

► Compaq Tru64 UNIX
► Compaq OpenVMS Alpha
► HP-UNIX
► IBM AIX
► IBM NUMA-Q
► Novell
► Red Hat or SUSE Linux
► SUN Solaris
► Windows NT, Windows 2000, and Windows 2000 Datacenter

For the latest information about supported operating system versions, refer to the server page documentation in the interoperability matrix, which can be found at:

http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Detailed and complete information on the ESS Copy Services Command Line Interface is provided in the *IBM TotalStorage Enterprise Storage Server Command-Line Interface User's Guide*, SC26-7494.

7.1.1 Requirements for Copy Services CLI

The Copy Services Command Line Interface is Java based, and therefore the Java JDK needs to be installed on each host system from which you want to issue the commands.

The use of the Command Line Interface requires at a minimum Java JDK level 1.1.8. It is also supported with Java level 1.3. For current recommendations, refer to the *IBM TotalStorage Enterprise Storage Server Command-Line Interfaces User's Guide*, SC26-7494.

On AIX, Windows, SUN Solaris, Linux (Red Hat or SUSE), and OpenVMS, the CLI installation process automatically installs a JRE on your host system. If you do not accept this option, you must install the required Java on your host system.

The host system does not necessarily need to be connected to storage assigned on one of the host ports of the ESS. The only requirement is that the server from which you want to

invoke the commands is connected to the ESS that is defined as the primary Copy Services server via a local area network (LAN). However, commands that relate the ESS volume serial number to the host's physical volume name or number (for example, the hdisk number in the case of AIX), such as the `rsPrimeserver` and `rsList2105s` commands, will only work on a host system that is physically connected to the ESS storage.

Because the copy commands need to communicate with the Copy Services server, you have to identify the Copy Services server before using any of the commands. For instructions on how to do this, refer to 5.3.2, "Location of ESS Copy Services server" on page 108.

Optionally, it is possible to authorize the usage of the Copy Services commands by specifying a user and its password when invoking the commands. The administration of these users is done from the Administration panel of the Copy Services Web User Interface (see 5.8.3, "Managing the CLI user ID and password for an open-systems host" on page 129). In this case, the user authorization for the Copy Services CLI needs to be enabled, so you have to create new users for that purpose. The accounts already created for the ESS Specialist to administrate the ESS and its storage cannot be used for the Copy Services commands.

7.1.2 Installing Copy Services CLI

The Command Line Interface program is located on the Host Attach CD, which is delivered with your ESS. To install the CLI, you must be logged on to the respective platform with the highest administrative privilege (root, administrator, or supervisor).

For specific operating system installation information, refer to the *IBM TotalStorage Enterprise Storage Server Command-Line Interfaces User's Guide*, SC26-7494.

7.1.3 General command guidelines

In general, all flags can be specified in any order on the command line. The specification of duplicate command flags is not allowed.

The entire interface is case sensitive. For example, if you specify a task name for `rsExecuteTask`, the case is significant. Use the same upper/lower case combination as defined in the Web User Interface. Also, on UNIX systems, the file names are case sensitive.

Place the user name, password, task name, and host name parameters inside quotation marks when specified on an OpenVMS host.

If an argument value contains any embedded blanks, the value must be surrounded by quotes, either single or double quotes, depending on the operating system being used, for example:

```
-a "c:\data files\esscli\ess access file.dat"
```

7.1.4 Copy Services CLI Command description

The Copy Services CLI provides six commands you can use to monitor and manage predefined ESS Copy Services tasks. In the following section, we describe these commands and show their name variations.

Table 7-1 on page 304 provides an overview of the platform specific notations.

The name of the Copy Services CLI commands are the same for all the supported platforms. Only the CLI command for Novell 4.x differs slightly from those. For example, use `rsExTask` instead of `rsExecuteTask` on a Novell 4.x system.

Table 7-1   Platform specific notations

Commands	File ext.	Flag	Host systems	Example

`rsExecuteTask`
`rsList2105s`
`rsPrimeServer`
`rsQuery`
`rsQueryComplete`
`rsTestConnection`
	.sh	-	AIX, Tru64, Sun, NUMA-Q, Linux, and HP-UX	`rsExecuteTask.sh -a securityfilepath -s primaryserver`

		-	OpenVMS and Alpha	`rsExecuteTask -u "username "-p "password "`
	.exe	/	Win NT, Win 2000, Win 2000, and Datacenter	`rsExecuteTask.exe /a securityfilepath /s primaryserver`
	.nlm	/	Novell 5.x and Novell 6.x	`rsExecuteTask.nlm /a securityfilepath /s primaryserver`

`rsExTask`
`rsL2105s`
`rsPSrv`
`rsQ`
`rsQComp`
`rsTConn`
		/	Novell 4.x	`rsExTask.nlm /a securityfilepath /s primaryserver`

rsExecuteTask (rsExTask) command

The `rsExecuteTask` command executes one or more ESS Copy Services tasks, as defined and saved using the ESS Copy Services Web-based interface, and it waits for the given task to complete execution.

Successful completion of the task means that the copy pair relationship has been established, suspended, or terminated.

Use the `rsQueryComplete` command to determine if the PPRC or FlashCopy pairs have completed copying.

Syntax of rsExecuteTask command

The syntax for the `rsExecuteTask` command is:

```
rsExecuteTask.* [-v] {-u username -p password |-a securityfilepath} -s primaryserver
[-b backupserver] tasknames
```

In this command, the asterisk (*) represents the file extension used on the particular operating system. For further information, refer to Table 7-1.

Flags and parameters for rsExecuteTask command

You can specify one or more of the following flags and parameters when you issue the `rsExecuteTask` command:

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter.

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p	*Password*
	Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.
a	*securityfilepath*
	Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the u and p flags and associated parameters are not used.
	The format of the security path file is as follows: *username password*.
s	*primaryserver*
	Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.
b	*backupserver*
	Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.
taskNames	Specifies the name of one or more previously saved ESS Copy Services tasks. Separate the task names with a space when you specify more than one task.

Example

Example 7-1 shows the output you might see when you invoke the `rsExecuteTask` command on an AIX host system.

Example 7-1   Output of rsExecuteTask

```
# ./rsExecuteTask.sh -v -u yourusername -p yourpassword -s primaryservername es_pair12
   rsExecuteTask: Got task manager reference
   rsExecuteTask: *************Finding the tasks****************
   rsExecuteTask: Task es_pair_12 found by task manager
   rsExecuteTask: *************Scheduling the tasks****************
   rsExecuteTask: Task es_pair_12 scheduled with copy services server
   rsExecuteTask: *************Monitoring the tasks****************
   rsExecuteTask: Waiting on server...
   rsExecuteTask: Task es_pair_12 completed successfully
   rsExecuteTask: Command successful
```

rsList2105s (rsL2105s) command

The `rsList2105s` command displays the mapping of a host disk name to a 2105 volume serial number.

If you have the IBM Subsystem Device Driver (SDD) running on host systems other than Linux, NUMA-Q, Tru64, and OpenVMS, and if you have changed the configuration for the

host, such as adding or removing a volume, you must restart the host so that the SDD can recognize, add, or remove the paths.

If you use the `rsList2105s` command on an OpenVMS Alpha host system, which is a member of an OpenVMS cluster, the output does not display information about the following devices:

► ESS volumes to which the host system has only MSCP paths

► ESS volumes to which the host system uses only MSCP paths at this time, even though it has both MSCP and direct paths

Syntax of rsList2105s command

The syntax for the `rsList2105s` command is simply:

```
rsList2105s.*
```

In this command, the asterisk (*) represents the file extension used on the particular operating system. For further information, refer to Table 7-1 on page 304.

Flags and parameters for the rsList2105s command

The `rsList2105s` command does not have any flags or parameters.

Example

If the SDD is installed, you might see output similar to Example 7-2 when you invoke the `rsList2105s` command on an AIX host system.

Example 7-2   Output of rsList2105s if SDD is installed

```
#./rsList2105s.sh
   VpathName      Serial      VolumeNames
   -----------    ------      -----------
   vpath10        40EFC102    hdisk14
   vpath11        40FFC102    hdisk15
   vpath13        40CFC102    hdisk17
   vpath14        40DFC102    hdisk18
```

If SDD is not installed, you might see output similar to Example 7-3 when you invoke the `rsList2105s` command on an AIX host system.

Example 7-3   Output of rsList2105s if SDD is not installed

```
#./rsList2105s.sh
   disk name       2105 serial number
   ---------       ------------------
   hdisk14           40EFC102
   hdisk15           40FFC102
   hdisk17           40CFC102
   hdisk18           40DFC102
```

rsPrimeServer (rsPSrv) command

The `rsPrimeServer` command creates, updates, or removes a list view of volumes associated with a particular host in the ESS Copy Services server. You can see the same view in the **Volumes** panel of the ESS Copy Services Web User Interface.

This command is useful when you use the ESS Copy Services **Volumes** panel to perform FlashCopy, PPRC, or both, because it shows the host names for all the volumes that are connected to the selected host in one view.

If the SDD is installed on your host system and you change the configuration for the host (such as adding or removing volumes), you must perform the following steps:

1. Restart your host system so that the SDD can recognize, add, or remove the paths.

2. Run the **rsPrimeServer** command to generate the current mapping for the host volumes.

If the SDD is not installed on your host system and you change the configuration for the host (such as adding or removing volumes), you must run the **rsPrimeServer** command to generate the current mapping for the host volumes.

If you use the **rsPrimeServer** command on an OpenVMS Alpha host system with MSCP-only access to some ESS volumes, these volumes are not shown in the host-specific volume list view in the ESS Copy Services server.

Syntax of rsPrimeServer command
The syntax for the **rsPrimeServer** command is:

```
rsPrimeServer.* [-v] {-u username -p password |-a securityfilepath} -d hostname -s
primaryserver [-b backupserver]
```

In this command, the * represents the file extension used on the particular operating system. For further information, refer to Table 7-1 on page 304.

Flags and parameters for rsPrimeServer command
You can specify one or more of the following flags and parameters when you issue the **rsPrimeServer** command:

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter.

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p   *Password*

Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

a   *securityfilepath*

Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password-protection for the host system and if the u and p flags and associated parameters are not used.

The format of the security path file is as follows: *username password*.

s   *primaryserver*

Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the

primaryserver parameter, you must identify and configure an ESS Copy Services server.

b   *backupserver*

Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.

d   *hostname*

Specifies the IP address or host name to be removed. The host name must exactly match one of the listing entries on the **Volumes** panels of the ESS Copy Services Web-based interface. The d flag with the hostname parameter is required when you are removing the host name.

Example

Example 7-4 shows the output that you see when you invoke the **rsPrimeServer** command to add a host name on an AIX host system.

Example 7-4   Output of rsPrimeServer to add a host name

```
# ./rsPrimeServer.sh -v -u yourusername -p yourpassword -s primaryservername
   rsPrimeServer: Local host is abc.def.ghi.com
   rsPrimeServer: Got task manager reference
   rsPrimeServer: Host abc.def.ghi.com registered
   rsPrimeServer: Configuration change requested made to primaryserver
   rsPrimeServer: Command successful
```

Example 7-5 shows the output that you might see when you invoke the **rsPrimeServer** command to delete a host name on an AIX host system.

Example 7-5   Output of rsPrimeServer command to delete a host name

```
# ./rsPrimeServer.sh -v -u yourusername -p yourpassword -d abc.def.ghi.com
-s primaryservername
   rsPrimeServer: host to unregister is abc.def.ghi.com
   rsPrimeServer: Got task manager reference
   rsPrimeServer: Host abc.def.ghi.com unregistered
   rsPrimeServer: Configuration change requested made to primaryserver
   rsPrimeServer: Command successful
```

rsQuery (rsQ) command

The **rsQuery** command queries the status of one or more volumes. You can invoke the **rsQuery** command for the vpath name, the host volume name, or the volume serial number under various conditions.

If the SDD is installed, use the **rsQuery** command for the vpath name or volume serial number. If the SDD is not installed, use the **rsQuery** command for the host volume name or volume serial number.

The **rsQuery** command fails on an OpenVMS host system if you specify the -m flag for a volume to which the host system has MSCP-only access. In other words, if your OpenVMS host system only has MSCP access to a volume, specify the volume serial number when using the **rsQuery** command.

Syntax of rsQuery command

The syntax for the rsQuery command is:

```
rsQuery.* [-v]-m {-u username -p password | -a securityfilepath} {-q volumename | -f
filename} -s primaryserver [-b backupserver]
```

In this command, the asterisk (*) represents the file extension used on the particular operating system. For further information, refer to Table 7-1 on page 304.

Flags and parameters for rsQuery command

You can specify one or more of the following flags and parameters when you issue the `rsQuery` command:

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter

m   Maps host volume names to 2105 volume serial numbers. This flag allows host volume names to be used as a volume flag. Use the m flag if you use the rsQuery command for the vpath name of host volume name. Do not use the m flag if you use the rsQuery command for a serial number.

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p   *Password*

Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

a   *securityfilepath*

Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password-protection for the host system and if the u and p flags and associated parameters are not used.

The format of the security path file is as follows: *username password*.

q   *volumename*

Specifies the vpath name, host volume name, or volume serial number. Use the q flag with the volume parameter if the flag with the filename parameter is not used.

Use the m flag if you use the `rsQuery` command for the vpath name or the host volume name. Do not use the m flag if you use the `rsQuery` command for a serial number.

f   *filename*

Specifies the name of a file that contains vpath names, host volume names, or volume serial numbers, all of which are to be queried. Use

the f flag parameter with the filename parameter if the q flag with the volume parameter is not used.

The format of the filename file can be either:

sourceVolume1   *targetVolume1*

sourceVolume2   *targetVolume2*

....

sourceVolumeN   *targetVolumeN*

or

volume1

volume2

volume3

....

volumeN

s   *primaryserver*

Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.

b   *backupserver*

Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.

Example

Example 7-6 shows the output that you might see when you invoke the **rsQuery** command to query a serial number on an AIX host system.

Example 7-6   Output of rsQuery

```
#./rsQuery.sh -v -u yourusername -v yourpassword  -q G0622532 -s primaryserverna
   ************************Volume Information************************
   Volume G0622532 found on 22532:00 as volume number 0x06
   PPRC State=source, Type=synchronous, Status=copy_pending
   PPRCPendingTracks=00017323
   FlashCopy_state=none, Size=1.03_GB
   PPRC Peer=G7A22532
   ******************************************************************
   rsQuery: Command successful
```

The sample output contains the following fields and values:

► **PPRC State**: Specifies the current state of a volume in a PPRC relationship with one of the following values:

   – simplex
   – source
   – target
   – unknown

► **Type**: Specifies one of the following types of a volume PPRC relationship:

  – synchronous
  – extended distance

► **Status**: Specifies the current status of a volume in a PPRC relationship with one of the following values:

  – copy_pending
  – suspended
  – fullcopy
  – none
  – unknown

  If the volume status is copy_pending, the `rsQuery` command also reports the status of PPRCPending sectors for fixed blocks or PPRC PendingTracks for CKD volumes.

► **FlashCopy_state**: Specifies the current state of a volume in a FlashCopy relationship with one of the following values:

  – source
  – target
  – none
  – unknown

rsQueryComplete (rsQComp) command

The `rsQueryComplete` command accepts ESS Copy Services server tasks as defined and saved with the ESS Copy Services Web-based interface and determines whether all volumes defined in the given tasks have completed the initial PPRC synchronization or FlashCopy background copy.

When the `rsQueryComplete` command is executed on grouped tasks, the output displays the first subtask until it reaches the completion threshold, then it displays the next subtask until it reaches the completion threshold. It repeats this process until all the subtasks are queried.

Syntax of rsQueryComplete command

The syntax for the `rsQueryComplete` command is:

```
rsQueryComplete.* [-v] {-u username -p password |-a securityfilepath} [-m minutes]
[-t threshold] -s primaryserver [-b backupserver] tasknames
```

In this command, the asterisk (*) represents the file extension used on the particular operating system. For further information, refer to Table 7-1 on page 304.

Flags and parameters for rsQueryComplete command

You can specify one or more of the following flags and parameters when you issue the `rsQueryComplete` command:

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter

u   *Username*

              Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p	*Password*
	Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.
a	*securityfilepath*
	Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password-protection for the host system and if the u and p flags and associated parameters are not used.
	The format of the security path file is as follows: *username password*.
m	*minutes*
	Specifies the number of minute and second intervals between queries for the PPRC or FlashCopy synchronization completion status. The m flag and the minutes parameter are optional. The default is one minute.
	The format of this parameter is as follows: *mm:ss*.
	The *mm* stands for the number of minutes and ss for the number of seconds before the command issues a new query.
t	*threshold*
	Specifies the threshold percentage that defines the query completion status. The default is 100%. The t flag and the threshold parameter are optional.
s	*primaryserver*
	Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.
b	*backupserver*
	Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.
taskNames	Specifies the name of one or more previously saved ESS Copy Services tasks. Separate the task names with a space when you specify more than one task.

Example

Example 7-7 on page 313 shows the output that you might see when you invoke the `rsQueryComplete` command and specify the minutes (m) parameter with a value of 0:10 and the threshold (t) parameter with a value of 80% on an AIX host system.

Example 7-7   Output of rsQueryComplete

```
#./rsQueryComplete.sh -v -u yourusername -p yourpassword -m 0:10 -t 80
-s primaryservername es_pair_12
   rsQueryComplete: Got task manager reference
   rsQueryComplete: --------- Task Name: es_pair_12 ---------
   rsQueryComplete: Task es_pair_12 found by TaskManager
   rsQueryComplete: PPRC Type = Synchronous.  Threshold = 80%
   rsQueryComplete: waiting 10 seconds...
   rsQueryComplete: Sampling volumes...
   rsQueryComplete: Percentage complete = 17
   rsQueryComplete: The tracks remaining to be copied = 994
   rsQueryComplete: waiting 10 seconds...
   rsQueryComplete: Sampling volumes...
   rsQueryComplete: Percentage complete = 32
   rsQueryComplete: The tracks remaining to be copied = 234
   rsQueryComplete: waiting 10 seconds...
   rsQueryComplete: Sampling volumes...
   rsQueryComplete: Percentage complete = 82
   rsQueryComplete: The tracks remaining to be copied = 0
   rsQueryComplete: ---------------------------------------------
   rsQueryComplete: Command successful
```

The sample output contains one of the following fields and values:

► **PPRC Type**: Specifies one of the following types of PPRC tasks:

– Synchronous

– Extended Distance

► **FlashCopy**: Specifies one of the following types of FlashCopy tasks:

– Background Copy

– No Background Copy

rsTestConnection (rsTCConn) command

The `rsTestConnection` command determines whether you can successfully connect to the ESS Copy Services server.

Syntax of rsTestConnection command

The syntax for the `rsTestConnection` command is:

`rsTestConnection.* [-v] -s servername`

In this command, the asterisk (*) represents the file extension used on the particular operating system. For further information, refer to Table 7-1 on page 304.

Flags and parameters for rsTestConnection command

You can specify one or more of the following flags and parameters when you issue the `rsTestConnection` command:

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter.

s   *servername*

Specifies the IP address or the complete host name of the ESS Copy Services server to which you want to test the connection.

Example

Example 7-8 shows the output that you might see when you invoke the `rsTestConnection` command on an AIX host system.

Example 7-8   Output or rsTestConnection on an AIX host

```
#./rsTestConnection.sh -v -s primaryservername
    rsWebTest: Using yourhostname as server name
    rsWebTest: rsVSServer reference obtained successfully
    rsWebTest: rsVSServer reference narrowed successfully
    rsWebTest: HeartBeat to the server was successful.
    rsWebTest: Command successful
```

7.2 ESS Storage Management CLI

The Storage Management CLI provides an alternate method to perform ESS logical configuration and storage management functions. With the Storage Management CLI (ESSCLI), routine configuration and management tasks can now be automated through their incorporation into scripts and applications, helping to simplify ESS administration.

The ESS CLI supports the following functions:

► **Asset Management:** Obtain information about ESS volumes, I/O ports, volume spaces, disk groups, and connected hosts.

► **LUN Masking:** Expose/unexpose volumes to Fibre Channel based initiators and obtain volume access information, including a list of volumes not exposed to any initiator.

► **Space Management:** Query available free space and create new volumes.

► **Volume Identification:** Assign a user-specified label to a volume.

► **Host Connections:** Define, undefine, and modify host connections and obtain host type information, including the host port's world-wide name.

► **Audit Log:** Obtain a log of configuration activity by user ID.

► **Parallel Access Volumes (PAV):** List and create new PAVs.

► **Copy Services:** Query functions to view tasks in the Copy Services task repository, obtain PPRC path status, and determine completion of the FlashCopy background copy.

The Storage Management CLI is available for the following operating systems:

► Hewlett-Packard-UX
► IBM AIX
► Red Hat Linux
► Sun Solaris
► Windows 2000

For the latest information about supported operating system versions, refer to the server page documentation in the inter-operability matrix, which can be found at:

http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Detailed and complete information on the ESS Copy Services Command Line Interface is provided in the *IBM TotalStorage Enterprise Storage Server Command-Line Interface User's Guide*, SC26-7494.

7.2.1 Requirements for the Storage Management CLI

The Storage Management CLI is implemented using the Java programming language. It uses the Java Secure Socket Extension (JSSE) and as such, requires a minimum level of the Java Runtime Environment (JRE 1.3.1).

On AIX, Windows, Sun Solaris, and Red Hat Linux, the ESSCLI installation process automatically installs a JRE on your host system. If you do not accept this option, you must install the required Java on your host system.

The host system does not necessarily need to be connected to storage assigned on one of the host ports of the ESS. The only requirement is that the server from which you want to invoke the commands is connected to the ESS to be managed via a local area network (LAN).

Because the copy commands need to communicate with the Copy Services server, you have to identify the Copy Services server before using any of the commands. For instructions on how to do this, refer to 5.3.2, "Location of ESS Copy Services server" on page 108.

7.2.2 Installing the Storage Management CLI

The Storage Management CLI is located on the Host Attach CD, which is delivered with your ESS. To install the ESSCLI, you must be logged on the respective platform with the highest administrative privilege (root, administrator, or supervisor).

For specific operating system installation information, refer to the *IBM TotalStorage Enterprise Storage Server Command-Line Interfaces User's Guide*, SC26-7494.

7.2.3 General command guidelines

In general, all command flags, as well as command actions and command categories, can be specified in any order on the command line. The same applies to the arguments; the positioning of the name-value pairs is not relevant to the operation of Storage Management CLI. However, the specification of duplicate command flags and duplicate argument names is not allowed.

The entire interface is case insensitive, so any combination of upper and lower case characters will be accepted by the Storage Management CLI for any of the command flags, argument names, as well as command action and category. For command parameters and argument values, the case may be significant depending on the function being performed. For example, file names are case sensitive in UNIX, of course, and the values of host names and volume labels are case sensitive in the ESS configuration.

For command parameters and argument values that allow a list of values to be specified, the list must be comma-separated, without any intervening blanks or other white-space. For example:

```
ports = A0,80,04
```

If a command parameter or argument value contains any embedded blanks, the value must be surrounded by quotes, either single or double quotes depending on the operating system being used. For example:

```
"c:\data files\esscli\ess access file.dat"
"ess=2105.12345 host='Jean Luc' profile=linux"
```

7.2.4 Command description

The Storage Management CLI provides a single `esscli` command with multiple flags, parameters, and arguments you can specify.

This section defines the Storage Management CLI commands managing ESS Copy Services tasks by monitoring and querying the tasks in the Copy Services task repository. This section also presents the Storage Management CLI command use to monitor and query the Paths defined for Copy Services. These commands described in this chapter are:

```
esscli list task
esscli show task
esscli list PPRCPaths
```

For further information about configurations and storage management functions, refer to the *IBM TotalStorage Enterprise Storage Server Command-Line Interfaces User's Guide*, SC26-7494.

esscli list task

This command generates a report that includes all defined tasks in the Copy Services Task repository.

Syntax of the esscli list task command

The syntax for the `esscli list task` command is:

```
esscli list Task {-a securityfilepath | -p password -u userid} -s primaryserver
[-b backupserver] [-v] [-nohdr]
```

Flags and parameters for esscli list task command

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p   *Password*

Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

a   *securityfilepath*

Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password-protection for the host system and if the u and p flags and associated parameters are not used.

The format of the security path file is as follows: *username password*.

s	*primaryserver*
	Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.
b	*backupserver*
	Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.
nohdr	Omits all header lines from the output.

Example

You might see output similar to Example 7-9 when you invoke the `esscli list task` command on a host system.

Example 7-9   Output of the esscli list task command

```
#esscli list task -v -u yourusername -p yourpassword -s primaryservername

Sun Aug 11 02:23:49 PST 2002 IBM ESSCLI 2.1.0.0
TaskName                        Type                         Status
-------------------------------------------------------------------------------
EstPath1                        PPRCEstablishPaths           NotRunning
RemovePaths                     PPRCRemovePaths              Successful
CharliesTask1                   PPRCEstablishPair            NotRunning
CharliesTask2                   PPRCSuspendPair              Failed
CharliesTask3                   PPRCTerminatePair            NotRunning
CharliesTask4                   PPRCFreezeGroup              Successful
CharliesTask5                   PPRCConsistencyCreated       Successful
MaggiesTask                     FCWithdraw                   Successful
RicksTask                       FCEstablish                  Successful
aGroupedTask                    Group                        Successful
```

The output contains the following fields and values:

► **TaskName**: The name of the task within the Copy Services task repository.

► **Type**: The task type, with one of the following values:

– FCEstablish
FlashCopy Establish Task

– FCWithdraw
FlashCopy Withdraw Task

– Group
Group Task

– PPRCconsistencyCreated
PPRC Consistency Group Created Task

– PPRCEstablishPair
PPRC Establish Pair Task

– PPRCEstablishPath
PPRC Establish Path Task

– PPRCFreezeGroup
PPRC Consistency Group Freeze Task

- PPRCRemovePath
  PPRC Remove Path Task

- PPRCSuspendPair
  PPRC Suspend Pair Task

- PPRCTerminatePair
  PPRC Terminate Pair Task

▶ **Status**: Last known status of Copy Services task.

esscli show task command

This command returns all available information about one or all tasks. It lists attributes of a task, including source and target volumes. If TaskName is specified, only information for this task is provided. If omitted, all tasks in the Copy Services Repository are included.

Syntax for esscli show task

The syntax for the `esscli show task` command is:

```
esscli show task {-a securityfilepath | -p password -u userid} -s primaryserver
[-b backupserver] [-v] [-nohdr] [-d "name=TaskName"]
```

Flags and parameters for the esscli show task command

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p   *Password*

Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

a   *securityfilepath*

Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the u and p flags and associated parameters are not used.

The format of the security path file is as follows: *username password*.

s   *primaryserver*

Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.

b   *backupserver*

Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter

is optional. Note that you can use the Backup server only after the primary server is down.

nohdr   Omits all header lines from the output.

d   Additional argument to specify the task name. If a task name is specified, only information for this task is provided. The syntax for this parameter is: -d "name=*Taskname*"

Example
Example 7-10 shows the output that you might see when you invoke an `esscli show task` command.

Example 7-10   Output of the esscli show tasks command

```
#esscli show task -v -u yourusername -p yourpassword -s primaryservername -d
"name=MaggiesTask"

Sun Aug 11 02:23:49 PST 2002 IBM ESSCLI 2.1.0.0
TaskName=MaggiesTask
TaskType=FCEstablish
Options=TargetOnline,ExtendedDistance,ReadFromSecondary
SourceServer=2105.FA123
TargetServer=2105.FA123
SourceVol       TargetVol
--------------------------------------------------------------------------------
1645            1690
1646            1691
1647            692
```

The sample output contains the following fields and values:

▶ **TaskName**: The name of the task within the Copy Services task repository.

▶ **Type**: The task type with one of the following values:

– FCEstablish
FlashCopy Establish Task

– FCWithdraw
FlashCopy Withdraw Task

– Group
Group Task

– PPRCconsistencyCreated
PPRC consistency Group Created Task

– PPRCEstablishPair
PPRC Establish Pair Task

– PPRCEstablishPath
PPRC Establish Path Task

– PPRCFreezeGroup
PPRC Consistency Group Freeze Task

– PPRCRemovePath
PPRC Remove Path Task

– PPRCSuspendPair
PPRC Suspend Pair Task

– PPRCTerminatePair
PPRC Terminate Pair Task

► **Options**: The task option with one or more of the following:

– AcceleratedDestageMode
Accelerate the destage of tracks on the source FlashCopy volume.

– CopyAllTracks
During PPRC establish copy all tracks from source to target.

– CopyOutOfSyncTracks
During PPRC establish copy out of sync tracks from source to target.

– CreateConsistencyGroup
Create a PPRC consistency group.

– CriticalVolumeMode
Prohibit writes to the primary PPRC volume if data cannot be copied to the secondary volume.

– DoNotCopyVol
During PPRC establish, do not perform initial copy from source to target.

– DoNotDestageModifiedData
Do not require data to be destaged to the source FlashCopy volume when a source track is modified.

– FCWithdrawToTarget
FlashCopy Withdraw to Target.

– FlashCopyStartBackgroundCopy
Initiate a background copy on an existing FlashCopy pair.

– ForceRemovalOfPPRCPath
Force the remove of the PPRC Paths even if PPRC pairs exist.

– IssueToPrimaryDevice
Issue the PPRC command to the primary volume.

– IssueToSecondaryDevice
Issue the PPRC command to the secondary volume.

– NoBackgroundCopy
Do not perform a background copy for the FlashCopy pair.

– NoForceEstablish
Do not force the establishment of PPRC Paths if paths already exist.

– PersistentFlashCopy
Persistent FlashCopy Pair establish.

– PPRCExtendedDistance
Establish an Extended Distance PPRC Pair.

– PPRCFailback
Fail back the PPRC pair.

– PPRCFailover
Fail over the PPRC pair.

– ReadFromSecondary
Permit host reading from the secondary volume involved in a PPRC pair.

– SecondaryOnlineOk
Permit the establish of the FlashCopy pair even if the target volume is online to the host.

– SingleDeviceSpecified
Only a single volume was specified to the task (source or target).

– SuspendAfterEstablish
When initial PPRC copy of tracks is complete, suspend the pair.

► **SourceServer**: The ESS ID of the primary storage server. The format of the ESS ID consists of the machine type, followed by a period (.), followed by the machine sequence number.

► **TargetServer**: The ESS ID of the secondary storage server. The format of the ESS ID consists of the machine type, followed by a period (.), followed by the machine sequence number.

► **SourceLSS**: The logical subsystem on the primary storage server.

► **TargetLSS**: The logical subsystem on the secondary storage server.

► **SourceVol**: The volume ID of the source volume. The first two characters are the LSS number (00-1F) and the following digits are the volume number (00-FF).

► **TargetVol**: The volume ID of the target volume. The first two characters are the LSS number (00-1F) and the following digits are the volume number (00-FF).

esscli list PPRCPaths

Use this command to list the path status for each established PPRC path between a source and one or more target logical subsystems.

Syntax for esscli list PPRCpaths

The syntax for esscli list PPRCpaths is:

```
esscli list PPRCpaths {-a securityfilepath | -p password -u userid} -s primaryserver
[-b backupserver] [-v] [-nohdr] [-fmt Format] -d "ess=ESSID [srclss=LssId]"
```

Flags and parameter for esscli list PPRCpaths

v   Displays all responses from the server. This verbose flag is optional and does not require a parameter

u   *Username*

Specifies the user name that is authorized to execute the Copy Services CLI commands. The u flag with the username parameter is required if the administrator for the ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

p   *Password*

Specifies the password of a user name that is authorized to execute the Copy Services CLI commands. The p flag with the password parameter is required if the administrator for ESS Copy Services server has enabled password protection for the host system and if the a flag with the securityfilepath parameter is not used.

a   *securityfilepath*

Specifies the full path to the file that contains the user name and password defined at the ESS Copy Services server for using Copy Services CLI commands on a host system. The a flag with the securityfile parameter is required if the administrator for the ESS Copy Services server has enabled password-protection for the host system and if the u and p flags and associated parameters are not used.

The format of the security path file is as follows: *username password*.

s	*primaryserver*
	Specifies the IP address or the complete host name of the ESS Copy Services primary server. Before you use the s flag with the primaryserver parameter, you must identify and configure an ESS Copy Services server.
b	*backupserver*
	Specifies the IP address or the complete host name of the ESS Copy Services Backup server. The b flag with the backupserver parameter is optional. Note that you can use the Backup server only after the primary server is down.
nohdr	Omits all header lines from the output.
fmt	Use the fmt flag to define the output fields in the desired order. The default *Format* tag is: "srclss,srcport,tgtserver,tgtlss,conn,status".
d	Additional argument to specify the *EssId* and the *LssID* (optional).
	The *EssId* identifies the Enterprise Storage Server (ESS). The format is "Type.Sequence". For example: 2105:FA123.
	The *LssId* identifies the logical subsystem in the ESS. The format consists of two hexadecimal characters, for example, 1F. If specified, `esscli list pprcpaths` will report only those paths where the given LSS is the source.

Example

Example 7-11 shows the output that you might see when you invoke an `esscli list pprcpaths` command.

Example 7-11   Output of esscli list pprcpaths command

```
esscli list pprcpaths -v -u yourusername -p yourpassword -s primaryservername -d
"ess=2105.20288 srclss=16"

Sun Aug 11 02:23:49 PST 2002 IBM ESSCLI 2.1.0.0
SourceServer=2105.20288
TotalPaths=7
SrcLss    SrcPort    TgtServer      TgtLss    Conn      Status
------------------------------------------------------------------------------
16        80         2105.FA123     18        Fabric    01
16        00         2105.FA123     18        Fabric    01
16        80         2105.FA14      20        Fabric    01
16        00         2105.FA145     20        Fabric    01
16        01         2105.FA178     18        P2P       01
16        04         2105.FA178     18        P2P       01
16        05         2105.FA178     18        P2P       02
```

The sample output contains the following fields and values:

► **SourceServer**: The ESS ID of the primary storage server. The format of the ESS ID consists of the machine type, followed by a period (.), followed by the machine sequence number.

► **TotalPaths**: The total number of paths for the source LSS.

► **SrcLss**: The logical subsystem on the primary storage server.

► **SrcPort**: The port ID on the primary storage server. See Appendix D, "System Adapter ID (SAID)" on page 493 for details.

- ► **TgtServer**: The ESS ID of the secondary storage server. The format of the ESS ID consists of the machine type, followed by a period (.), followed by the machine sequence number.

- ► **TgtLss**: The logical subsystem on the secondary storage server.

- ► **Conn**: PPRC Connection Type:

  - – Fabric
    Path connected through switch

  - – P2P
    Direct connect path

- ► **Status**: The detailed status of the path. See Table 7-2.

- ► **TgtPort**: The port ID on the secondary storage server. See Appendix D, "System Adapter ID (SAID)" on page 493 for details.

- ► **SwId**: The switch identifier.

- ► **SwPort**: The outgoing port on the switch.

Table 7-2   Path status table

Status	Description
1	Path established.
2	Initialization failed.
3	Timeout. No reason available.
4	No resources available at primary site for the logical path establishment.
5	No resources available at secondary site for the logical path establishment.
6	Secondary site sequence number or logical subsystem number mismatch.
7	Secondary site SSID mismatch.
8	Path is offline. This is caused by lack of light detection coming from a host, peer, or switch.
9	Establish failed but will try again when conditions improve.
0A	The port at the primary storage server cannot be converted to channel mode because an inbound logical path is already established.
0B	Reserved.
10	Configuration error.
11	Error during path establish.
12	Error during path establish.
13-FF	Reserved.

7.3 ESS Storage Management CLI Version 2

The ESS Storage Management CLI V2 proposes new commands you can use to invoke, perform, and manage various ESS functions from your host system.

We do not propose a full description of those commands in this chapter, since they are not directly linked to the implementation of Copy Services. Nevertheless, we will give a short description of the new commands.

To get a full description of these commands, refer to *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface User's Guide*, SC26-7449.

esscli create volume

The ESS management CLI V2 enhanced this command by enabling you to create also CKD volumes

esscli create/delete volumespace

You can use this command to create or delete configured storage space. For the ESS, a VolumeSpace is a RAID array or a non-RAID (JBOD) disk. When a Disk Group is configured as a RAID array, one VolumeSpace is created. When a Disk Group is configured as non-RAID array, up to eight Volume Spaces are created, one for each disk in the group.

The command will return before the actual Volumespace is formatted. A delete and create command cannot be run immediately in sequence on the same disk group The user should poll the Shark to see if the volume space has finished formatting.

esscli list featurecode

This command lists the active Feature Codes on an ESS.

esscli list/create/set/delete webuseraccount

This new set of commands helps to manage the user accounts on the ESS. You can use the commands to:

► List the active user on the ESS.
► View the settings of the currently defined user accounts that are permitted to access the ESS.
► Create user Accounts to have access to the ESS.
► Modify Web user accounts that have access to the ESS.
► Change user names.
► Change passwords.
► Modify Access levels.
► Modify IP ranges.
► Delete users.

Most of these commands require Administration access level.

esscli show/set perfstats

Use this set of commands to display and modify the current settings of the data collector. These settings include the name of the machine that is to receive the collected statistics, the frequency at which statistics should be collected, the port number where the statistics should be sent, and the status of the data collection.

esscli show/set remotesupport

Display the current settings set for the ESS remote support. With this set of commands you can:

- ► Indicate whether incoming and outgoing calls are enabled.
- ► Modify or set the remote support properties available in the ESS.
- ► Set the call home and remote access properties.
- ► Set service and PE passwords.

esscli show/create/delete/set email

This set of commands will let you manage all the information regarding the e-mail addresses configured for problem notification. You can:

- ► Display the information about e-mail addresses.
- ► Output the e-mail addresses and the corresponding error notification trigger value indicating the types of problems.
- ► Add a new e-mail address to the problem notification configuration.
- ► Remove an e-mail address from the current problem notification configuration.
- ► Modify an existing e-mail address and its corresponding options.
- ► Change the trigger options on an existing e-mail address.
- ► Rename the existing e-mail address with a new address.
- ► Change the global settings of e-mail.
- ► Maximum number of e-mails sent per problem.
- ► Sending test e-mail.

esscli show/create/delete/set pager

You can manage the pager settings configured for problem notification with the CLI commands. With these commands you can:

- ► Display a list of all pagers, with all their related information, and global pager settings.
- ► Add a new pager number to the problem notification configuration.
- ► Remove a pager from the current problem notification configuration.
- ► Modify existing pager information.
- ► Update the global properties of pagers.

esscli show/create/ delete/setsnmp

This set of commands will let you manage all the information regarding the snmp configuration for problem notification. You can:

- ► Add trap addresses to the existing problem notification configuration.
- ► Modify existing SNMP information.
- ► Delete trap addresses from the SNMP problem notification configuration.

esscli list/show/delete/ problem

You can use these commands to manage the active problems on an ESS. With these CLI commands you can:

- ► List all outstanding problems active on the ESS (this will output a table that displays the problem ID, cluster, and description of each active problem).

- ► Get information on specific active Problems.
- ► View the currently active problem log records (when an ESS monitors the operation of its internal components and logical resources, abnormal conditions are entered in the problem log).
- ► Cancel any outstanding problems active on the ESS.

8

Open systems specifics

In this chapter, we describe the basic tasks that need to be performed on the individual host systems when using ESS Copy Services.

We explain how to bring FlashCopy target volumes online to the same host as well as to a second host. The chapter covers various UNIX and WinTel platforms.

8.1 AIX specifics

In this section, we describe the steps needed to use volumes created by the ESS Copy Services on AIX hosts.

8.1.1 AIX and FlashCopy

The FlashCopy functionality from the ESS Copy Services copies the entire contents of a source volume to a target volume. If the source volume is defined to the AIX Logical Volume Manager (LVM), all of its data structures and identifiers are copied to the target volume, as well. This includes the Volume Group Descriptor Area (VGDA), which contains the Physical Volume Identifier (PVID) and Volume Group Identifier (VGID).

For AIX LVM, it is currently not possible to activate a volume group with a physical volume (hdisk) that contains a VGID and a PVID that is already used in a volume group existing on the same server. The restriction still applies even if the hdisk PVID is cleared and reassigned with the two commands listed in Example 8-1.

Example 8-1   Clearing PVIDs

```
chdev -l <hdisk#> -a pv=clear
chdev -l <hdisk#> -a pv=yes
```

Therefore, it is necessary to redefine the volume group information on the FlashCopy target volumes using special procedures or the **recreatevg** command (refer to "AIX recreatevg command" on page 332). This will alter the PVIDs and VGIDs in all the VGDAs of the FlashCopy target volumes, so that there are no conflicts with existing PVIDs and VGIDs on existing volume groups that reside on the source volumes. If you do not redefine the volume group information prior to importing the volume group, then the **importvg** command will fail.

Accessing FlashCopy target volume from another AIX host

The following procedure makes the data of the FlashCopy target volume available to another AIX host that has no prior definitions of the target volume in its configuration database (ODM):

1. The target volume (hdisk) is new to AIX, and therefore the Configuration Manager should be run on the specific SCSI or Fibre Channel adapter:

   ```
   cfgmgr -l <host_bus_adapter>
   ```

2. Find out which of the physical volumes is your FlashCopy target volume:

   ```
   lsdev -Cc disk | grep 2105
   ```

3. Import the target volume group:

   ```
   importvg -y <volume_group_name> <hdisk#>
   ```

4. Vary on the Volume Group (the **importvg** command should varyon the volume group):

   ```
   varyonvg <volume_group_name>
   ```

5. Verify consistency of all file systems on the FlashCopy target volume:

   ```
   fsck -y <filesystem_name>
   ```

6. Mount all the target file systems:

   ```
   mount <filesystem_name>
   ```

The data is now available. You can, for example, back up the data residing on the FlashCopy target volume to a tape device. This procedure can be run once the relationship between the FlashCopy source and target volume is established, even if data is still being copied in the background.

The disks containing the target volumes may have been previously defined to an AIX system, for example, if you periodically create backups using the same set of volumes. In this case, there are two possible scenarios:

► If no volume group, file system, or logical volume structure changes were made, then use Procedure 1 to access the FlashCopy target volumes from the target system

► If some modifications to the structure of the volume group were made, such as changing file system size or the modification of logical volumes (LV), then it is highly *not recommended* to perform the steps described in Procedure 1 to access the FlashCopy target volumes. In this case, it is recommended to use Procedure 2.

Procedure 1

1. Unmount all the source file systems:

   `umount <src_filesystem>`

2. Unmount all the target file systems:

   `umount <tgt_filesystem>`

3. Deactivate the target volume group:

   `varyoffvg <tgt_volume_group_name>`

4. Establish the FlashCopy relationship(s).

5. Mount all the source file systems:

   `mount <src_filesystem>`

6. Activate the target volume group:

   `varyonvg <tgt_volume_group_name>`

7. Perform a file system consistency check on target file systems:

   `fsck -y <tgt_filesystem>`

8. Mount all the target file systems:

   `mount <tgt_filesystem>`

Procedure 2

1. Unmount all the target file systems:

   `umount <tgt_filesystem>`

2. Deactivate the target volume group:

   `varyoffvg <tgt_volume_group_name>`

3. Export the target volume group:

   `exportvg <tgt_volume_group_name>`

4. Delete the target physical volumes (optional; not needed if using the same set of physical volumes):

   `rmdev -dl <hdisk#>`

5. Establish the FlashCopy relationship(s).

6. Then perform tasks as if the volumes were new to the system (if you did not remove the physical volumes from the target systems, then you do not have to run the Configuration Manager or identify the volumes).

Accessing the FlashCopy target volume from the same AIX host

In this section, we describe a method of accessing the FlashCopy target volume on a single AIX host while the source volume is still active on the same server. The procedure is intended to be used as a guide and may not cover all scenarios.

The steps needed to access the FlashCopy target volume depend on the level of AIX used. If your level is AIX Version 4.3.3 maintenance level 05 (APAR IY10456) or higher, then skip over the next section and use the procedure outlined in "AIX recreatevg command" on page 332. Otherwise, use the procedure described in "AIX recreatevg command not used" on page 330.

AIX recreatevg command not used

In our example, we have a volume group called fc_source_vg on hdisk3. In this volume group, there is a logical volume (fc_source_lv) defined with a mounted file system on /fc_source_fs. This volume will be our source for the FlashCopy operation. The target of our FlashCopy volume will be hdisk4, which is currently not in use.

Make sure that the source for your FlashCopy is in a consistent state for the short period of time while establishing the FlashCopy pair. Use the following procedure to access the target volume:

1. If necessary, configure the FlashCopy target LUN to AIX:

   ```
   cfgmgr -l <host adapter>
   ```

   This will make the target LUN known to AIX as an *hdisk*. It is recommended to run the configuration manager against an individual SCSI or Fibre Channel adapter rather than globally.

2. Get the source physical partition size:

   ```
   lsvg fc_source_vg | grep "PP SIZE"
   ```

   In our example, the physical partition size is 16 MB.

3. Bring down applications that access the FlashCopy source and unmount the related file systems for the short period of FlashCopy establishment.

4. Establish the FlashCopy pair.

5. Once the FlashCopy pair is established, mount all file systems and restart the applications. This could be done even if data is still being copied from the source to the target in the background.

6. Clear the PVID from the target hdisk to allow a new volume group to be created:

   ```
   chdev -l hdisk4 -a pv=clear
   ```

7. Create a new volume group on the target volume using the physical partition size of the volume group containing the source logical volume:

   ```
   mkvg -y fc_target_vg -s 16 -f hdisk4
   ```

8. Generate the physical partition map for the AIX logical volume on the source and write it to a file. The output of the `lslv -m` command needs to be reformatted for input to the `mklv` command. This step can be performed while the FlashCopy process is copying data. For example:

   ```
   lslv -m fc_source_lv | awk '/hdisk/ {print $3 ":" $2 }' | sed 's/\(:0*\)/:/' | sed "s/hdisk3/hdisk4/g" > /tmp/lv_map.out
   ```

You should note that the foregoing command *does not work for mirrored file systems*. To generate the map file for a mirrored file system, use the following command. For example:

```
lslv -m fc_source_lv | awk '/hdisk/ {print $3 ":" $2 " " $5 ":" $4}' | sed
's/\(:0*\)/:/' | sed "s/hdisk3:/hdisk4:/g" | sed "s/hdisk13:/hdisk14:/g"> /tmp/map.out
```

9. Determine the number of physical partitions in the logical volume. This value is used in Step 10:

```
cat /tmp/lv_map.out | wc -l
```

10. Create the target logical volume using the physical volume map file that was created from the source logical volume's partition map. Specify the new logical volume name, the map file name, the logical volume type, the volume group, and the number of partitions in the new logical volume:

```
mklv -y fc_target_lv -m /tmp/lv_map.out -t jfs fc_target_vg 50
```

11. Make a new JFS log logical volume. This step must be performed after all logical volumes have been created. Otherwise, creation of a subsequent logical volume may fail with a physical partition number conflict error. Run:

```
mklv -y fc_log -t jfslog fc_target_vg 1 hdisk4
```

12. Format the new logical volume for use as a JFS log:

```
logform /dev/fc_log
```

13. Add a new stanza to /etc/filesystems to include the new file system attributes. Make sure the device entry points to the target logical volume and the log entry points to the new JFS log (see Figure 8-1).

```
/fc_target_fs:
dev             = /dev/fc_target_lv
vfs             = jfs
log             = /dev/fc_log
mount           = false
options         = rw
account         = false
```

Figure 8-1   File system stanza in /etc/filesystems

14. Create a new mount point for the target:

```
mkdir /fc_target_fs
```

15. Check the new file system (target):

```
fsck -V jfs -y /fc_target_fs
```

16. Mount the new file system:

```
mount /fc_target_fs
```

Repeat steps 8 through 16 for each file system on the FlashCopy source volume, with the exception of steps 11 and 12. Of course, you need only one JFS log logical volume. You may wish to automate the steps by creating a shell script that recreates each file system.

Note: If you are using JFS and JFS2 in the same volume group, you will need a special journaled log device for each type of the journaled file system.

Once the file system is mounted, you have access to the data on the FlashCopy target volume.

Keep in mind that the procedure described above will, more than likely, need to be adapted to your configuration and specific objectives. For example, you may have multiple physical volumes in the volume group, mirrored files systems, and so on.

AIX recreatevg command

In "AIX recreatevg command not used" on page 330, you can see that the manual procedure in AIX to overcome the problems of making a FlashCopy target volume available to the same host as the source volume is not trivial.

Copying the source volume's content using FlashCopy causes all of the data structures and identifiers used by AIX's Logical Volume Manager to be duplicated to the target volume. The duplicate definitions (PVID and VGID) in turn cause conflicts within LVM. This problem is solved by using the AIX command `recreatevg`.

The `recreatevg` command is packaged as a PTF for AIX Version 4.3.3 in APAR IY10456 and higher. It is officially available in:

► AIX 4.3.3 Recommended Maintenance Level 05 (RML05) or higher
► AIX 5L Version 5.1
► AIX 5L Version 5.2

The `recreatevg` command overcomes the problem of duplicated LVM data structures and identifiers caused by a disk duplication process such as FlashCopy. It is used to recreate an AIX Volume Group (VG) on a set of target volumes that are copied from a set of source volumes belonging to a specific VG. The command will allocate new physical volume identifiers (PVIDs) for the member disks and a new volume group identifier (VGID) to the volume group. The command also provides options to rename the logical volumes with a prefix you specify, and options to rename labels to specify different mount points for file systems.

Here is the AIX man page synopsis (might be different, depends on the AIX version used):

```
recreatevg [ -y VGname ] [ -p ] [ -f ] [ -Y lv_prefix | -l LvNameFile ] [ -L label_prefix ]
[ -n ] \ PVname...
```

► Description

This command can be used to recreate a VG on a set of disks that are mirrored from a set of disks belonging to a specific VG. This command will allocate new physical volume identifiers (PVID) for the member disks, as the PVIDs will also be duplicated by the disk mirroring. Similarly, other LVM logical members that are duplicated will also be changed to new names with the specified prefixes.

► Flags

    -y   VolumeGroup specifies the volume group name rather than having the name generated automatically. Volume group names must be unique system wide and can range from 1 to 15 characters. The name cannot begin with a prefix already defined in the PdDv class in the Device Configuration database for other devices. The volume group name created is sent to standard output.

    -p   Disables the automatic generation of the new PVIDs. If the -p flag is used, you must ensure that there are no duplicated PVIDs on the system. All the disks that were hardware mirrored must have had their PVIDs changed to an unique value.

    -Y   lv_prefix causes the logical volumes on the VG being recreated renamed with this prefix. The number of characters in the prefix should be such that the total length of the prefix and the logical volume name must be less than or equal to 15 characters. If the

length exceeds 15 characters, the logical volume will be renamed with the default name. The name cannot begin with a prefix already defined in the PdDv class in the Device Configuration Database for other devices, nor be a name already used by another device.

-L   label_prefix causes the labels of logical volumes on the VG being recreated to be changed with this prefix. The user must modify the /etc/filesystems stanza manually if a simple modification of the mount point is not enough to define the stanza uniquely.

-l   LvNameFile entries in the LvNameFile must be in the format LV1:NEWLV1. After **recreatevg**, LV1 will be renamed with NEWLV1. All the logical volumes that are not included in the LvNameFile will be recreated with the default system generated name.

-f   Allows a volume group to be recreated that does not have all disks available.

-n   After **recreatevg**, the volume group is imported but varied off. The default is imported and vary on.

► Notes

  a. To use this command, you must have root user authority.

  b. All the member physical volumes of the volume group must be specified on the command line. The command will fail if the input list does not match with the list compiled from the Volume Group Descriptor Area (VGDA).

  c. If you perform a Copy Services function on one half of a RAID-1 pair to reduce the capacity required for FlashCopy targets or PPRC secondary volumes, then use the -f option to force the creation of the volume group. Otherwise, the VGDA will have PVIDs of volumes that made up the other half of the mirror at the source or primary site.

► Examples:

  a. To recreate a volume group that contains three physical volumes, enter this command:

```
recreatevg hdisk1 hdisk2 hdisk3
```

  The volume group on hdisk1, hdisk2, and hdisk3 is recreated with an automatically generated name, which is displayed.

  b. Run:

```
recreatevg -y testvg hdisk1
```

  The volume group on hdisk1 is recreated with the new name testvg.

  c. Run:

```
recreatevg -Y newlv hdisk14
```

  The volume group on hdisk14 is recreated and all logical volumes in that volume group are recreated and renamed with the prefix newlv.

Accessing FlashCopy target volume using the recreatevg command

In this example, we have a volume group containing two physical volumes (hdisks) and wish to FlashCopy the volumes for the purpose of creating a backup.

The source volume group is fc_source_vg, containing hdisk4 and hdisk5.

The target volume group will be fc_target_vg, containing hdisk8 and hdisk9.

Perform these tasks to run the FlashCopy and make the target volumes available to AIX:

1. Stop all applications that access the FlashCopy source volumes.

2. Unmount all source file systems for the short period of FlashCopy establishment.

3. Establish the FlashCopy pairs with the **No Background Copy** option selected. Use the ESS Copy Services Web user interface (WUI) to establish the pairs or, if you have a task defined, use `rsExecuteTask.sh` in the CLI.

4. Mount all source file systems.

5. Restart applications that access the FlashCopy source volumes.

6. The target volumes, hdisk8 and hdisk9, will now have the same volume group data structures as the source volumes hdisk4 and hdisk5. Clear the PVIDs from the target hdisks to allow a new volume group to be made:

```
chdev -l hdisk8 -a pv=clear
chdev -l hdisk9 -a pv=clear
```

The output from `lspv` shows the result (see Figure 8-2).

```
# lspv
hdisk4          000567992d4c9024      fc_source_vg
hdisk5          000567995abe005e      fc_source_vg
hdisk8          none                  None
hdisk9          none                  None
```

Figure 8-2   lspv after pv=clear

7. Create the target volume group and prefix all file system path names with */backup* and prefix all AIX logical volumes with *bkup*:

```
recreatevg -y fc_target_vg -L /backup -Y bkup hdisk8 hdisk9
```

You must specify the hdisk names of all disk volumes participating in the volume group. The output from `lspv` illustrates the new volume group definition (see Figure 8-3).

```
# lspv
hdisk4          000567992d4c9024      fc_source_vg
hdisk5          000567995abe005e      fc_source_vg
hdisk8          000567995abdf345      fc_target_vg
hdisk9          00056799b2d831b9      fc_target_vg
```

Figure 8-3   Recreated FlashCopy target volumes

An extract from /etc/filesystems shows how `recreatevg` generates a new file system stanza. The file system named /u01 in the source volume group is renamed to /backup/u01 in the target volume group. Also, the directory /backup/u01 is created. Notice also that the logical volume and JFS log logical volume have been renamed. The remainder of the stanza is the same as the stanza for /u01 (see Figure 8-4 on page 338).

```
/backup/u01:
        dev             = /dev/bkupelv001
        vfs             = jfs
        log             = /dev/bkupelvlog001
        mount           = true
        check           = false
        options         = rw
        account         = false
```

Figure 8-4   Target file system stanza

8. Mount the new file systems belonging to the target volume group to make them accessible.

More detail is given on using AIX and FlashCopy as part of your backup strategy in Appendix F, "AIX and FlashCopy for backups" on page 509.

8.1.2 AIX and PPRC

When you have the primary and secondary volumes in a PPRC relationship, it is not possible to read the secondary unless the **Permit read from secondary** option has been selected when establishing the relationship (see 6.12.3, "Establishing a synchronous PPRC pair" on page 239). To be able to read the secondary volumes, they must also be in the full duplex state (in addition to **Permit read from secondary option**). Therefore, if you are configuring the secondary volumes on the target server, it is necessary to terminate the copy pair relationship. Once the volumes are in the simplex state, the secondary volumes can be configured (`cfgmgr`) into the target systems customized device class (CuDv) of the ODM. This will bring in the secondary volumes as hdisks and will contain the same physical volume IDs (PVID) as the primary volumes. Because these volumes are new to the system, there is no conflict with existing PVIDs. The volume group on the secondary volumes containing the logical volume (LV) and file system information can now be imported into the Object Data Manager (ODM) and the /etc/filesystems file using the `importvg` command.

If the PPRC secondary volumes were previously defined on the target AIX system as hdisks or vpaths, but the original volume group was removed from the primary volumes, the old volume group and disk definitions must be removed (`exportvg` and `rmdev`) from the target volumes and redefined (`cfgmgr`) before running `importvg` again to get the new volume group definitions. If this is not done first, `importvg` will import the volume group improperly. The volume group data structures (PVIDs and VGID) in ODM will differ from the data structures in the VGDAs and disk volume super blocks. The file systems will not be accessible.

If the secondary volumes that are already configured on the target AIX server are in a PPRC relationship and you do not have the **Permit read from secondary** option enabled (and the volumes are not in full duplex state), after rebooting the target server, the hdisks will be configured to AIX again. In other words, you will see each PPRC secondary volume twice on the target server. The reason for this situation is as follows: AIX knows that these physical volumes already exist with entries in the Configuration Database (ODM). However, when the configuration manager runs during reboot, it cannot read their PVIDs because, as PPRC targets, they are locked by the ESS Copy Services server. This results in AIX causing the original hdisks to be configured to a Defined state, and new (phantom) hdisks being configured and placed in an Available state. This is an undesirable condition that must be remedied before the secondary volumes can be accessed.

To access the secondary volumes, the phantom hdisks must be removed and the real or original hdisks must be changed from a Defined state to an Available state.

For example, hdisk6 through hdisk9 are assigned to a volume group, evg001. Each of the disk volumes is currently participating as a secondary volume in a PPRC relationship. If the server is rebooted, four new hdisks are configured to AIX. These phantom disks, hdisk13 through hdisk16, appear in the output from **lspv**, as shown in Figure 8-5.

```
# lspv
hdisk6          000567992d4c9024      evg001
hdisk7          000567995abe005e      evg001
hdisk8          000567995abdf345      evg001
hdisk9          00056799b2d831b9      evg001
hdisk13         none                  None
hdisk14         none                  None
hdisk15         none                  None
hdisk16         none                  None
```

Figure 8-5   PPRC phantom disks

When you execute **lsdev -Cc disk**, you can observe that the state of the original PPRC secondary volumes has become Defined during reboot (shown in Figure 8-6).

```
# lsdev -Cc disk
hdisk6  Defined    20-58-01      IBM FC 2105F20
hdisk7  Defined    20-58-01      IBM FC 2105F20
hdisk8  Defined    20-58-01      IBM FC 2105F20
hdisk9  Defined    20-58-01      IBM FC 2105F20
hdisk13 Available  20-58-01      IBM FC 2105F20
hdisk14 Available  20-58-01      IBM FC 2105F20
hdisk15 Available  20-58-01      IBM FC 2105F20
hdisk16 Available  20-58-01      IBM FC 2105F20
```

Figure 8-6   Phantom hdisks

It is important to execute both the **lspv** and **lsdev** commands, so that you can be sure which disks are the phantoms. From the **lspv** output, the phantom disks will have no PVIDs and will not be assigned to a volume group. From the **lsdev** output, the phantom will be in an Available state. The original disks will have PVIDs assigned to a volume group and will be in a Defined state.

To remove the phantom hdisks from the configuration database, run the **rmdev -dl** command on each phantom disk device (see Example 8-2).

Example 8-2   Removing phantom hdisks

```
# for i in 13 14 15 16
do
   rmdev -dl hdisk$i
done
```

Set the original hdisks to an Available state with the **mkdev** command (see Example 8-3).

Example 8-3   Configuring hdisks

```
# for i in 6 7 8 9
do
mkdev -l hdisk$i
```

You can reactivate the volume group, evg001, (`varyonvg`) and mount its file systems.

Making updates to the LVM information

When performing PPRC between primary and secondary volumes, the primary AIX host may create/modify or delete existing LVM information from a volume group. However, because the secondary volume is not accessible when in a PPRC relationship, the LVM information in the secondary AIX host would be out-of-date. Therefore, scheduled periods need to be allotted where write I/Os to the primary PPRC volume can be quiesced and file systems unmounted. At this point, the copy pair relationship can be terminated and the secondary AIX host can perform a learn on the volume group (`importvg -L`).

Once the updates have been imported into the secondary AIX hosts ODM, you can establish the PPRC copy pair again. However, select **Do not copy volume** from the Select copy Options when establishing the PPRC copy pair. As soon as the PPRC pair has been established, immediately suspend the PPRC relationship. Because there was no write I/O to the primary volumes, both primary and secondary are consistent.

Now that the primary volume has been suspended, the file systems can be remounted and write I/O resumed. Once the write I/O has been going for a while, you can re-establish the relationship with the primary and secondary by choosing **Copy out-of-sync cylinders only** (see 6.12.9, "Resynchronizing PPRC copy pairs" on page 256).

If the **Permit read from secondary** option was selected during the PPRC copy pair establish, then it would only be advisable to suspend the primary volume (while in the full duplex state) and then perform the import learn function. Once completed, all that is necessary would be to re-establish the copy pair only by copying the out-of-sync cylinders.

The following example shows two systems, sanf50 and sanh70, where sanf50 has the primary volume vpath5 and sanh70 has the secondary volume vpath16. Both systems have had their ODMs populated with the volume group itsovg from their respective PPRC volumes and, prior to any modifications, both systems ODM have the same time stamp, as shown in Figure 8-7.

```
root@sanf50:/ > getlvodm –T itsovg
3d99d8911542ab68

root@sanh70:/ > getlvodm –T itsovg
3d99d8911542ab68
```

Figure 8-7   Original time stamp

Volumes vpath5 and vpath16 are in the PPRC duplex state, and the volume group itsovg on sanf50 is updated with a new logical volume. The time stamp on the VGDA of the volumes gets updated and so does the ODM on sanf50, but not on sanh70 (see Figure 8-8 on page 338).

```
root@sanf50:/ > lqueryvg -p vpath5 -Tt
Time Stamp:     3d99e5de077392e4
root@sanf50:/ > getlvodm -T itsovg
3d99e5de077392e4

root@sanh70:/ > lqueryvg -p vpath16 -Tt
Time Stamp:     3d99e5de077392e4
root@sanh70:/ > getlvodm -T itsovg
3d99d8911542ab68
```

Figure 8-8   Updated source time stamp

To update the ODM on the secondary server, it is advisable to suspend the PPRC copy pair prior to performing the **importvg -L** command to avoid any conflicts from LVM actions occurring on the primary server. Figure 8-9 shows the updated ODM entry on sanh70.

```
root@sanh70:/ > importvg -L itsovg vpath16
root@sanh70:/ > lqueryvg -p vpath16 -Tt
3d99e5de077392e4
root@sanh70:/ > getlvodm -T itsovg
3d99e5de077392e4
```

Figure 8-9   Update secondary server's ODM

Once the **importvg -L** command has completed, you can re-establish the PPRC copy pairs and copy only the out-of-sync cylinders.

8.1.3 Scripting Copy Services tasks

The following scripts were created and run in the lab. Notice that they are shown as examples only and will need modification to run in your environment.

Automation of consistent backup using PPRC-XD

This script uses the following ESS Copy Services tasks when calling **rsExecuteTask.sh** commands.

The **PPRC_Est_Susp** task will transform the asynchronous PPRC-XD relationship into a synchronous PPRC relationship and suspend the pair once they are synchronized (see Example 8-4).

Example 8-4   PPRC_Est_Susp task

```
Task Information

  ---------------------------------------------------------------
  Task(C)

    Task type: Establish synchronous PPRC pair

    Task options: Copy out-of-sync cylinders only
                  Permit read from secondary
                  Suspend PPRC after establish

    Source:                 Target:
    18767:14                12320:14
```

```
    volume: 002 (40218767)    volume: 000 (40012320)
```

The **FlashCopy_1** task creates a FlashCopy of the PPRC secondary volume (40012320)
with the background copy option, which can be used to back up the data (see Example 8-5).

Example 8-5   FlashCopy_1 task

```
Task Information

  ----------------------------------------------------------------
  Task(FlashCopy_1)

    Task type: FlashCopy establish

    Task options: none

    Source:               Target:
    12320:14              12320:14
    volume: 000 (40012320)    volume: 001 (4012320)
```

The **PPRC-XD_Est_OS** task is used to re-establish the PPRC-XD relationship as it was
before the backup process began by establishing the PPRC-XD relationship and copying
out-of-sync cylinders (see Example 8-6).

Example 8-6   PPRC-XD_Est_OS task

```
Task Information

  ----------------------------------------------------------------
  Task(PPRC-XD_Est_OS)

    Task type: Establish extended distance PPRC pair

    Task options: Copy out-of-sync cylinders only
                  Permit read from secondary


    Source:               Target:
    18767:14              12320:14
    volume: 002 (40218767)    volume: 000 (40012320)
```

The **consbkup.sh** script can be used to perform a consistent backup of the production data
by stopping application IO, changing the PPRC-XD pair into a synchronized PPRC pair,
suspending the pair, running a FlashCopy of the PPRC secondary volume to create tertiary
volume, backing up the tertiary volume, restarting the IO, and then returning the PPRC pair
back to an asynchronous PPRC-XD pair. The script will require modification to work properly.
It is intended to provide the steps that must be performed to accomplish the mentioned goals.
It is just a simple example without any testing of prerequisites and error condition handling
(see Example 8-7).

Example 8-7   consbkup.sh

```
#!/bin/sh
#
#
#Stop all applications and unmount filesystems on source
umount /fs1
umount /fs2
#
```

```
#Synchronize the PPRC-XD pairs and then suspend
rsExecuteTask.sh -u username -p password -s 9.88.9.1 -b 9.88.9.2 PPRC_Est_Susp
#
#FlashCopy the PPRC target
rsExecuteTask.sh -u username -p password -s 9.88.9.1 -b 9.88.9.2 FlashCopy_1
#
#Backup the data
tar -cvf files backupfiles
#
#Restart application and mount filesystems
mount /fs1
mount /fs2
#
#Re-establish the PPRC-XD relationship and copy out-of-sync cylinders
rsExecuteTask.sh -u username -p password -s 9.88.9.1 -b 9.88.9.2 PPRC-XD_Est_OS
```

> **Note:** When calling Copy Services tasks inside shell scripts or at the command line, be aware that the task names are *case-sensitive*.

8.2 Windows NT specifics

In this section, we discuss what tasks are necessary when performing Copy Services operations on volumes owned by Microsoft Windows NT hosts:

8.2.1 Windows NT basic volumes

Windows NT handles disks in a way that is not similar to any other operating system covered in this book. The need to reboot a server in order to scan for new disks and the need to run a GUI-based Disk Administrator in order to manipulate the disks are the main factors that restrict the routine use of PPRC or FlashCopy.

It is possible to automate the actions of the GUI-based Disk Administrator using third-party software to remotely reboot the server and to remotely assign the drive letter from the server that starts the Copy Services task. This was not tested during our project. You can automate the invocation of PPRC and FlashCopy using the ESS Copy Services Command Line Interface (CLI).

If you are going to create an automated script with Windows NT, you need to be very careful about data consistency. It may be that some part of the automation process may run a script on a source server and subsequent actions maybe taken by a script on a target server. Therefore, inter process communication across servers may be required for timing or you may get inconsistent data (not all applications will allow this).

You have two options on how to make a PPRC or FlashCopy target available to the server, *with reboot* or *without reboot*. We recommend that you reboot the server, it is safer because then it is guaranteed that all the registry entries get created. However, using PPRC and FlashCopy without rebooting is faster.

Registering PPRC/FlashCopy volumes to NT

If you are going to reboot the server, you do not have to make the target disks known to Windows NT before you do the PPRC or FlashCopy. However, we recommend that you have them preassigned and registered in the server. The "assign disk and run PPRC/FlashCopy" approach is useful for a non-routine PPRC/FlashCopy, for example, for testing or migration.

For routine purposes, we recommend having target disks already present in the Disk Administrator with partitions created and partition information saved. You can accomplish that by clicking **Start** → **Programs** → **Administrative Tools** → **Disk Administrator**:

1. If the target disk has not been previously seen by the system, Disk Administrator will issue a pop-up message saying "No signature on Disk *X*. Should I write a signature?" where *X* is the number assigned to the newly present disk.

2. Click **OK** to save the signature on the target disk. The Disk Administrator will come up.

3. Click on the disk that is to be used as the PPRC/FlashCopy target (it should be grey and marked as Free Space) and select **Create**.

4. Confirm the partition parameters and click **OK**. The partition appears as Unknown.

5. Click on the newly created partition and select **Commit Changes Now**.

6. Right-click on the partition and select **Assign Drive letter**.

7. Assign a drive letter and click **OK.** Exit Disk Administrator.

After this procedure, the information about the PPRC/FlashCopy target volume is properly stored in the Windows NT registry.

Stopping the target server

Stop the server that will use the target if you want to use the safer method. Also bear in mind that if you were assigning the volume to the host just before performing the PPRC/FlashCopy, you will have to use the volume serial number for the target. You cannot use the `rsPrimeServer` CLI command to insert disk names into the ESS Copy Services Web Interface, as the server is down.

Performing a FlashCopy

Stop all applications using the source volume. Now you must flush the data to the source volume. You can accomplish that by clicking **Start** → **Programs** → **Administrative Tools** → **Disk Administrator**.

1. Click on the disk that is to be used as the PPRC/FlashCopy source (it should have a drive letter assigned and be formatted) and select **Assign Drive letter**.

2. In the pop-up box, select **Do not assign a drive letter** and click **OK**.

   Now the data is flushed to the source and you can start the PPRC/FlashCopy task from the ESS Copy Services Web Interface or from any server CLI.

   Observe the GUI or issue the `rsQuery` command to the volumes to see if the PPRC/FlashCopy task had successfully started. Then you can reassign the drive letter to the source volume.

3. Click on the disk that is a PPRC/FlashCopy source and select **Assign Drive Letter**.

4. Assign a drive letter and click **OK.** Exit Disk Administrator.

You can resume using the source volume.

Starting the target server

After that, you may boot up the target server. In this case, you have just assigned the target volumes to the host that will create the disk entry in the Windows NT registry. To verify that the registry entry is created, do the following:

1. Click **Start** → **Settings** → **Control Panel**.

2. In Control Panel, double-click **SCSI Adapters**.

3. Click the adapter that has the target volume attached.

That opens a list of targets. Verify the list includes the target ID and LUN of the volume you have just made available to the server. If you are using SDD, you will see each disk entry several times, depending on how many paths to a volume you have defined.

You may also run the command `datapath query device` from the SDD command line to check if the PPRC/FlashCopy targets are listed between the volumes. This command will also enable you to check volume serial numbers, and will give you a more understandable overview of the volumes and their paths.

Making PPRC/FlashCopy target available

Log in, start Windows NT Disk Administrator, write a signature if necessary, and assign a drive letter. To do that, click **Start** → **Programs** → **Administrative Tools** → **Disk Administrator**.

1. If the disk has not been previously seen by this system, Disk Administrator will issue a pop-up message saying "No signature on Disk X. Should I write a signature?", where X is the number assigned to the newly present disk.

2. Click **OK** to save the signature on the target disk. The Disk Administrator will come up.

3. Click on the disk that is a PPRC/FlashCopy target (you should see a formatted partition on it) and select **Assign Drive Letter**.

4. If you cannot assign a drive letter, the target is probably corrupt. Try repeating the whole process and consider the scenario that includes reboot.

5. Assign a drive letter and click **OK.** Exit Disk Administrator.

6. From a Windows NT command prompt, run `chkdsk x: /f /r`, where *x* is the letter assigned to the PPRC/FlashCopy target. An option is to run the disk check from the **Properties** of a disk in Windows NT Explorer.

7. After this procedure, the PPRC/FlashCopy target is available to the Windows NT and can be handled like normal disk.

8.2.2 Copy Services with Windows NT volume sets

The following section describes how to perform Copy Services functions with Windows NT volume sets. Basically, the same procedures apply to other types of Windows NT multi-volume logical drives, such as mirrored sets, striped sets, and striped sets with parity. But in practice, because the ESS provides RAID-1+0 and RAID-5, only volume sets are needed (to add more space to existing logical drives).

In Windows NT, information about multi-volume logical drives is stored only in the registry key HKEY_LOCAL_MACHINE\System\Disk (shortly called the Disk key). Because there is no volume information on the disks itself, it is not possible to import a multi-volume logical drive by another host. Instead, the Disk key of that host must be changed. Unfortunately, the Disk key value is in binary format, requiring special tools for manipulations.

You can find information about the Disk key's internal structure at the Web site:

http://www.sysinternals.com/ntw2k/info/diskkey.shtml

Both PPRC and FlashCopy are supported when using normal disks and volume sets. When using either PPRC or FlashCopy with volume sets, because these outboard copy features do not copy the volume set information in the Windows registry, certain limitations exist, and a special procedure is required, as outlined below. After SP6, it is possible to have the FlashCopy source and target volumes accessible by the same server. Prior to SP6, the FlashCopy source and target volumes must be attached to different servers. PPRC primary and secondary volumes must always be attached to different servers.

Procedure for using PPRC and FlashCopy with volume sets

This special procedure is required in order to FlashCopy or PPRC an NT volume set. If the target disks are in the same order as the source disks, *and* the target disks are contiguous (that is, all the disks are next to each other, as viewed by the target machine's Disk Administrator), then simply create an identical volume set on the target machine and reboot prior to performing the FlashCopy. This only has to be done before performing FlashCopy or PPRC for the first time. Subsequent copies should work as expected, provided that the file system is unmounted (the drive letter is unassigned) on the target prior to performing a copy.

If the target disks do not appear contiguous to NT or appear in a different order than on the source machine, then a different procedure must be used. Microsoft's FTEDIT program, available on the NT Resource Kit, is a Microsoft supported tool designed to read and modify the Disk key. Because of the binary structure of the Disk key, using FTEDIT is highly recommended instead of editing the registry directly.

> **Attention:** Incorrect use of FTEDIT could result in loss of access to software RAID arrays. It is recommended that you use Disk Administrator to save your disk configuration before using FTEDIT. In general, most errors made using FTEDIT are recoverable. For more information on how to recover from FTEDIT errors, and on FTEDIT in general, see the Microsoft Knowledge Base articles 131658 and 149927:
>
> http://support.microsoft.com/?kbid=131658
> http://support.microsoft.com/?kbid=149927

In the following sections, we describe the procedure for using FlashCopy and PPRC with FTEDIT.

Preparation

On the target machine:

1. Back up the disk data using Disk Administrator, and registry information using REGEDIT.

2. If the target disks have been previously used, delete all of the target disks in Disk Administrator. Do not just unmount them, but delete all of the partitions on the target disks. Commit the changes.

3. In the Control Panel, double-click **Devices** and make sure that Ftdisk is started, and is set to start on boot. Ftdisk is the driver used by NT to identify and access fault tolerant drives and other multi-volume logical drives, such as volume sets. If there are any such drives in use on the system, Ftdisk will be started and set to start on boot. If it is not started, one way to get it started is to create a multi-volume logical drive on a couple of spare disks. This will require a reboot.

On the source machine:

Obtain the order in which the disks were added to the volume set. One way of doing this is to use a freeware utility called `diskkey.exe`, available from:

http://www.sysinternals.com

This utility is not supported by IBM and is known to report disk numbering and other information that is different than what Disk Administrator reports. However, the order in which the disks are included in the volume set will be correct, and the correct ordering of the disks is the information required to create a duplicate volume set on the target server.

Map the disks on the source machine to the disks on the target machine. For example, determine that Disk6 on the source is FlashCopied to Disk9 on the target. One way to do this is to look at the tasks on the ESS. If the FlashCopy task is set up to FlashCopy from volume 10017608 to volume 10117608, for example, then use `rsList2105s` from the Copy Services CLI or the `datapath` command from SDD to show that 10017608 is Disk6 on the source machine and 10117608 is Disk9 on the target machine. Map all of the disks in the volume set to their counterpart on the target.

Performing the PPRC/FlashCopy

On the target machine:

1. Run the FlashCopy establish or PPRC terminate task(s).

2. Start Disk Administrator. If it asks you to write a signature on any of the disks, click **No** (except in the special cases described below in the note). Once Disk Administrator is up, commit the changes (this is very important), and close Disk Administrator.

> **Note:** Disk Administrator will ask to write a signature when the FlashCopy is performed to the same machine, because it detects a duplicate disk signature (the source and target volumes have the same disk signature) and needs to write a new one. It is safe to do this, but be sure that you are writing the signature to the FlashCopy target disk. If a signature is written to the wrong disk, it may cause data corruption.
>
> When FlashCopying to a different machine, usually the disk signature on the target machine's disks are different than the FlashCopy source disks' signature, so Disk Administrator does not need to write a new signature to the target disks to use it. It is unlikely, but possible, that by coincidence the disk signature of one of the source disks is the same as one of the disks on the target machine. In this case, you will have to write a signature on the target disk before using it. Again, it is safe to do this, but be sure that you are writing the signature to the right disk.
>
> **Attention:** Writing a new disk signature to the source disk may cause severe problems. For example, a cluster disk will not become online after a change of its signature.

3. Start FTEDIT by choosing **Start → Resource Kit 4.0 → Disk Tools → Fault Tolerance Editor**. Read the warning and click **OK**. There are two panes in the FTEDIT window. In the left-hand pane is a list of the disks in the system. In the right-hand pane is the list of partitions on that disk. You must add the disks to the volume set in the right order. Use the results of diskkey.exe to determine the order in which the disks were added on the source volume set.

> **Attention:** If there are active PPRC target volumes on the target, then the disk numbering used in FTEDIT may differ from the disk numbering used in Disk Administrator. The PPRC target volumes are not seen by FTEDIT, and so are not included in the disk numbering scheme. Adjust your disk choices accordingly.

4. Click **Make FT set** in the lower left-hand corner. When it asks you what kind of set you would like, choose **Volume set** and click **OK**.

5. Click the first target disk in the left-hand pane. The list of partitions on that disk should appear in the right-hand pane. Choose the partition that contains the volume set on that disk (usually Partition 1).

6. Double-click **Partition 1** in the right-hand pane. This will add this disk/partition to the volume set in order.

7. Repeat steps 5 and 6 for the rest of the disks. If you make a mistake, you can cancel and start from scratch. The disks *must* be added in the right order.

8. When all of the disks have been added, choose **Save FT set** at the bottom.

9. Click **Save Changes to System** in the Edit menu. Close FTEDIT.

10. Reboot.

11. When NT restarts, start Disk Administrator. The target disks should be yellow now, indicating that they are in a volume set. Assign a drive letter and commit the changes. If the drives are not usable at this point, then the disks were probably added in the wrong order.

As long as the disk configuration does not change on the source or target, FlashCopy should work as expected. If the disk configuration is changed in any way, such as adding an additional disk to the volume set or rearranging the disks, then you will have to perform this procedure again.

8.3 Windows 2000 and Copy Services

Windows 2000 handles its disks differently than does Windows NT. Windows 2000 incorporates a stripped-down version of the VERITAS Volume Manager, called the Logical Disk Manager (LDM).

With the LDM, you are able to create logical partitions, perform disk mounts, and create dynamic volumes. There are five types of dynamic volumes: simple, spanned, mirrored, striped, and RAID-5.

On Windows NT, the information relating to the disks was stored in the Windows NT registry. With Windows 2000, this information is stored on the disk drive itself in a partition called the LDM database, which is kept on the last few tracks of the disk. Each volume has its own 128-bit Globally Unique Identifier (GUID) and belongs to a disk group. This is similar to the concept of Physical Volume Identifier (PVID) and Volume Group in AIX. As the LDM is stored on the physical drive itself, with Windows 2000 it is possible to move disk drives between different computers.

Copy Services limitations with Windows 2000

Having the drive information stored on the disk itself imposes some limitations when using Copy Services functionality on a Windows 2000 system:

► The source and target volumes must be of the same physical size. Normally the target volume can be bigger than the source volume; with Windows 2000, this is not the case, for two reasons:

  a. The LDM database holds information relating to the size of the volume. As this is copied from the source to the target, if the target volume is a different size from the source, then the database information will be incorrect, and the host system will return an exception.

  b. The LDM database is stored at the end of the volume. The copy process is a track-by-track copy, unless the target is an identical size to the source, the database will not be at the end of the target volume.

► It is not possible to have the source and target FlashCopy volume on the same Windows 2000 System when they were created as Windows 2000 dynamic volumes. The reason is that each dynamic volume has to have its own 128-bit GUID. As its name implies, the GUID must be unique on one system. When you perform FlashCopy, the GUID gets copied as well, so this means that if you tried to mount the source and target volume on the same host system, you would have two volumes with exactly the same GUID. This is not allowed, and you will not be able to mount the target volume.

8.3.1 Copy Services with Windows 2000 volumes

Basic disks are the same as the NT disks with the same restrictions. Dynamic disks are supported for both PPRC and FlashCopy and the primary/source and secondary/target volumes must be attached to different servers. We also support the use of Spanned Volumes with PPRC and FlashCopy on Dynamic Disks with the same attachment restriction.

Mounting a Copy Services target volume

In order to see target volumes on a second Windows 2000 host, you have to do the following:

1. Perform the PPRC/FlashCopy function onto the target volume. Ensure that when using PPRC that the primary and secondary volumes were in duplex mode, and write I/O was ceased prior to terminating the copy pair relationship.

2. Reboot the host machine on which you wish to mount the Copy Services target volume.

3. Right-click **Open Computer Management**, and then click **Disk Management**.

4. Find the disk that is associated with your volume. There are two "panes" for each disk; the left one should read **Dynamic** and **Foreign**. It is likely that no drive letter will be associated with that volume.

5. Right-click that pane, and select **Import Foreign Disks**. Select **OK**, then **OK** again. The volume now has a drive letter assigned to it, and is of Simple Layout and Dynamic Type. You can read/write to that volume.

> **Tip:** Disable the **Fast-indexing** option on the source disk, otherwise operations to that volume get cached to speed up disk access. However, this means that data is not flushed from memory and the target disk may have copies of files/folders that were deleted from the source system.

When performing subsequent PPRC/FlashCopies to the target volume, it is not necessary to perform a reboot because the target volume is still known to the target system. However, in order to detect any changes to the contents of the target volume, you should remove the drive

letter from the target volume before doing the FlashCopy. Then, after carrying out the FlashCopy, you restore the drive letter in order for the host it is mounted on to be able to read/write to it.

There is a Windows utility, Diskpart, that enables you to script these operations so that FlashCopy can be carried out as part of an automated backup procedure. Diskpart can be found at the Microsoft download site http://www.microsoft.com/downloads and the search on the key word Diskpart. A description of Diskpart commands can be found at the Web site:

http://www.microsoft.com/technet/prodtechnol/winxppro/proddocs/DiskPart.asp

Extending simple volumes

The Copy Services source may initially be a single *simple* volume. However, as requirements change on the application server, the logical volume may be extended over two or more volumes. However, you should not independently extend the target volumes, but let Windows 2000 detect the correct sequence of the extended volumes during the import process. For this reason, the target volumes should be reverted back to *basic* disks prior to the initial FlashCopy after the source has been extended. The target server should also be rebooted for disk manager to pick up the new volumes.

After the reboot, the volumes will be recognized as *foreign* disks, and you can proceed to import them. Reboot of the target system on subsequent FlashCopy is not necessary until the source volume has been further extended. When performing subsequent PPRC/FlashCopy to the target volume, it is not necessary to perform a reboot because the target volume is still known to the target system. However, in order to detect any changes to the contents of the target volume, you should remove the drive letter from the target volume before doing the FlashCopy. Then, after carrying out the FlashCopy, you restore the drive letter in order for the host it is mounted on to be able to read/write to it.

There is a Windows utility, Diskpart, that enables you to script these operations so that FlashCopy can be carried out as part of an automated backup procedure. Diskpart can be found at the Microsoft download site http://www.microsoft.com/downloads and the search on the key word Diskpart.

The Diskpart tool also provides a way to extend an existing partition into free space at the end of the same logical drive. A description of this procedure can be found in the Microsoft Knowledge Base, article 304736:

http://support.microsoft.com/?kbid=304736

Enlarging Extended/Spanned volumes

When you have extended *or* spanned disks, the logical drive may in time grow to include more of the initial volume (extended disk) or include additional volumes. When this occurs, it is necessary, as before, to remove the target volume group information and revert the target volumes back to basic disks. On the initial FlashCopy, it is necessary to reboot the target server to configure the additional disks, and then import all the foreign disks that are part of the volume group.

When performing subsequent PPRC/FlashCopy to the target volume, it is not necessary to perform a reboot, because the target volume is still known to the target system. However, in order to detect any changes to the contents of the target volume, you should remove the drive letter from the target volume before doing the FlashCopy. Then, after carrying out the FlashCopy, you restore the drive letter in order for the host it is mounted on to be able to read/write to it.

Again we refer to the Windows 2000 utility Diskpart, which enables you to script these operations so that FlashCopy can be carried out as part of an automated backup procedure.

Diskpart can be found at the Microsoft download site http://www.microsoft.com/downloads; Search for the key word "Diskpart".

8.3.2 PPRC and Windows 2000 spanned volumes

Here is the procedure we followed when carrying out the PPRC of a Windows 2000 spanned volume set from server A to server B:

1. On the source server A, we created a Windows spanned volume set of multiple dynamic disks.

2. We rebooted the target server B and we imported multiple target disks and wrote a disk signature on each as basic disks

3. Established PPRC between the source and target volumes.

4. Once the source and target volumes were synchronized, we terminated PPRC.

5. We rebooted the target host B.

6. We started Disk Manager; the PPRC target volumes were seen as Foreign Dynamic Disks.

7. The disks were imported into the target host and were seen as a spanned volume.

In order to demonstrate failback to the original setup, we carried out the following tasks:

1. We removed the original paths and re-established them in the reverse direction from B to A.

2. We removed the spanned volume drive letter from the original source (the spanned volume on server A).

3. We established PPRC from B to A and wrote some data onto the spanned volume.

4. PPRC was terminated.

5. We restored the drive letter to the spanned volume on server A.

6. The contents of the spanned volume could now be read from server A.

There is a Windows utility, Diskpart, that enables you to script these operations so that FlashCopy can be carried out as part of an automated backup procedure. Diskpart can be found at the Microsoft download site http://www.microsoft.com/downloads; search for the key word "Diskpart".

8.3.3 ESSSync utility: Available by RPQ

The ESSSync utility is available from IBM by RPQ (request for price quotation) only. This utility flushes Windows buffers and enable you to access FlashCopy target volumes without the need to unassign then reassign drives letters. Here is the procedure for using it:

1. Run ESSSync against the source and target volumes to flush all buffers (both volumes are mounted on their respective hosts).

2. Perform the FlashCopy.

3. Run ESSSync against the target volume.

4. Access the data on the target volumes.

ESSSync is run from the command line and can be incorporated into your backup scripts.

8.3.4 Relating ESS volumes to Windows drive letters

When logical volumes are created in the ESS, they are assigned a serial number. These numbers are shown when you look at the volumes using the ESS Copy Services GUI. This can be seen in Figure 8-10.
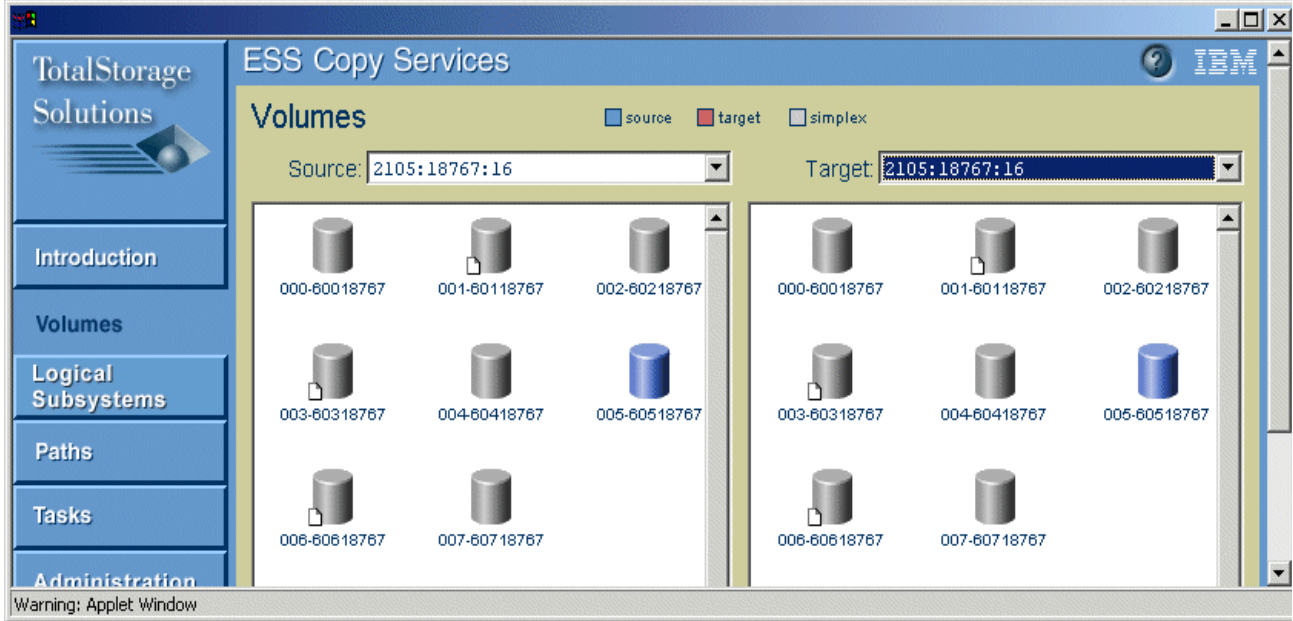
Figure 8-10 ESS logical volume serial numbers

Disk administrator will see the drives as shown in Figure 8-11, that is, as Disk 1, drive letter F and Disk 2, drive letter G.
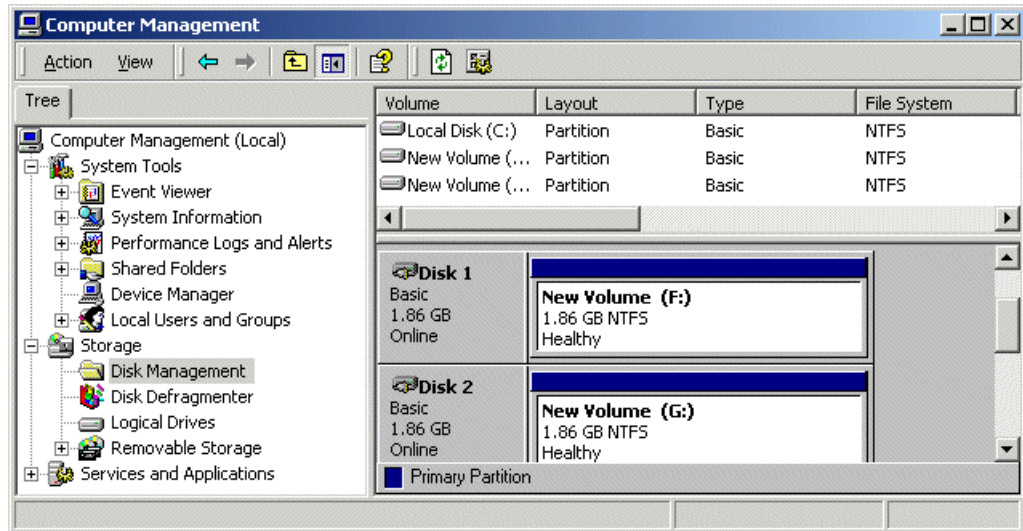
Figure 8-11 View of disks as seen by Disk Administrator

If you have the CLI installed and you run the `rsList2105s` command, then you get the output shown in Figure 8-12 on page 350. This shows the relationship between Disk 1 and Disk 2 and the ESS serial numbers. Note that Disk 3 and Disk 4 are the views of Disk 1 and Disk 2 seen via the second Fibre Channel path. These views of the disk are not seen by applications.
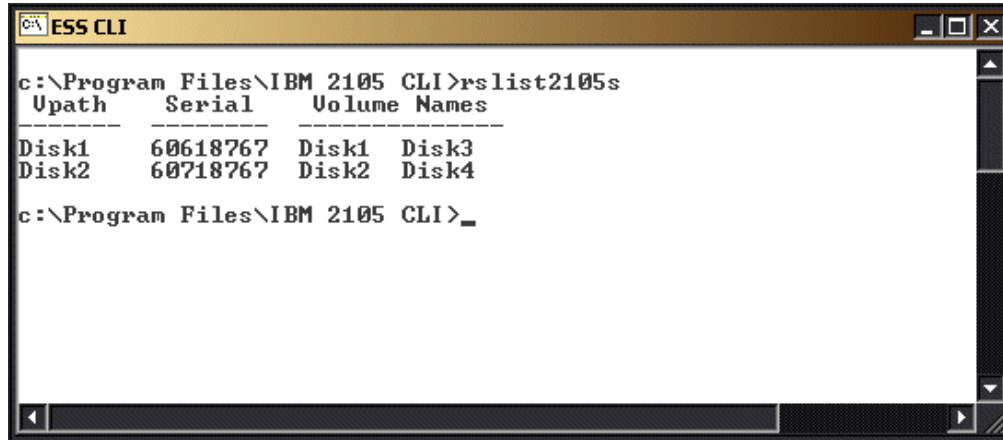
Figure 8-12 Output from rsList2105s

Also, if you have the CLI installed on the host that is accessing the volumes, then you can run the `rsPrimeServer` command. This will then show the Windows volume letters that have been assigned to the ESS volumes when you view the **Volumes** panel in the Copy Services GUI (see Figure 8-13).
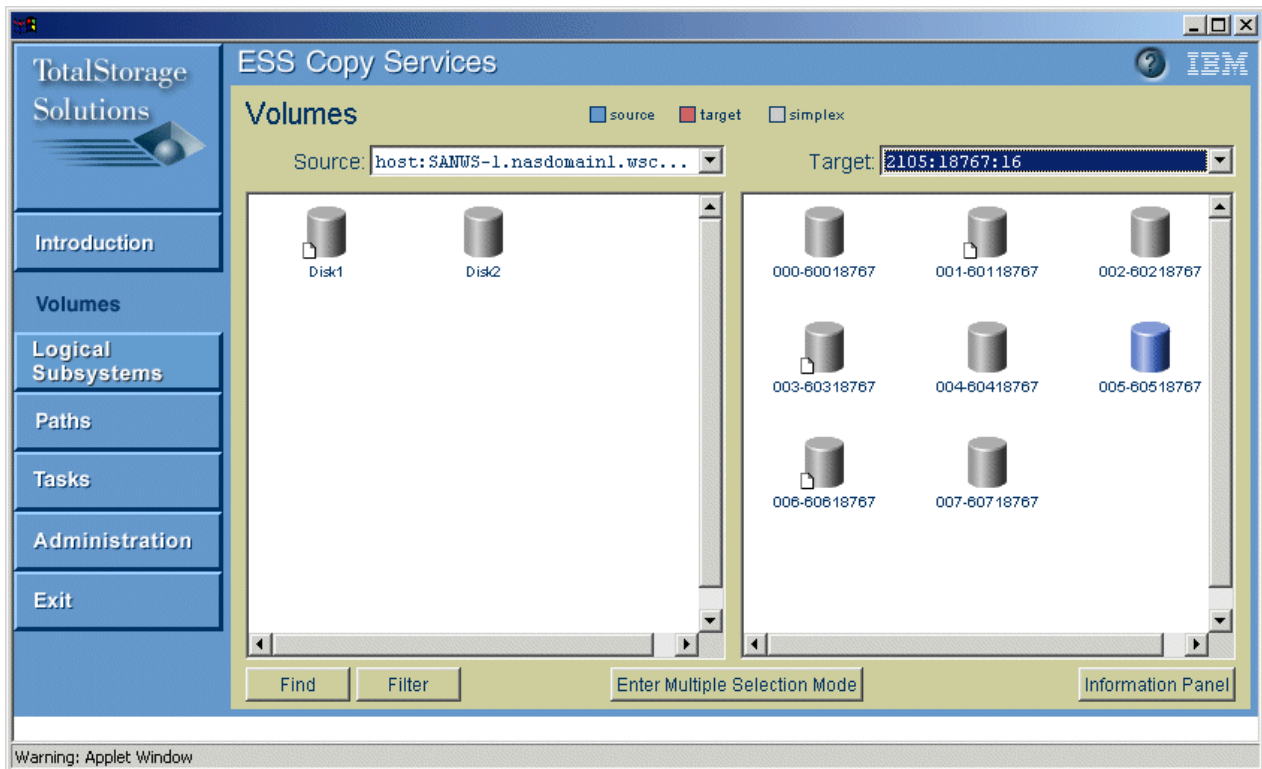
Figure 8-13 View after rsPrimeServer has been run

The IBM SDD software can be used to carry out query commands to find the condition of multipath devices attached to a Windows 2000 server. The two main commands used are `datapath query adapter` and `datapath query device`. This is shown in Figure 8-14 on page 351.

```
Subsystem Device Driver Management                                          _ □ ×
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query device

Total Devices : 2

DEV#:    0  DEVICE NAME: Disk1 Part0  TYPE: 2105F20   SERIAL: 60618767
============================================================================
Path#            Adapter/Hard Disk      State     Mode     Select     Errors
    0     Scsi Port2 Bus1/Disk1 Part0    OPEN    NORMAL       273          0
    1     Scsi Port2 Bus1/Disk3 Part0    OPEN    NORMAL       321          0

DEV#:    1  DEVICE NAME: Disk2 Part0  TYPE: 2105F20   SERIAL: 60718767
============================================================================
Path#            Adapter/Hard Disk      State     Mode     Select     Errors
    0     Scsi Port2 Bus1/Disk2 Part0    OPEN    NORMAL       415          0
    1     Scsi Port2 Bus1/Disk4 Part0    OPEN    NORMAL       407          0

C:\Program Files\IBM Corp\Subsystem Device Driver>
```

Figure 8-14   Datapath query device command output

The **datapath query adapter** command will return a list of adapters, as illustrated in
Example 8-8. The output should include information about all the installed adapters. In this
example, one SCSI adapter is installed on the Windows 2000 host server.

Example 8-8   The datapath query adapter command

```
+--------------------------------------------------------------------------------+
|Active Adapters :1                                                              |
|                                                                                |
|Adpt#      Adapter Name   State     Mode     Select     Errors  Paths  Active   |
|   0  Scsi Port1 Bus0   NORMAL    ACTIVE     4057          0      8      8      |
+--------------------------------------------------------------------------------+
```

8.4 SUN Solaris and Copy Services

In the following section, we describe the actions that should be taken to perform Copy
Services functions and mount a target volume on a SUN Solaris server.

Making a Copy Services target volume available to the same server or to another server is
possible.

You can use the Copy Services CLI for automation and create scripts to automate your
procedures. We recommend that you predefine the tasks to be run and test them thoroughly.
Also, prepare your target mount point.

8.4.1 Copy Services without a volume manager

In this section, we describe how to access Copy Services volumes under SUN Solaris without
volume manager software. Native commands are used to show how it is possible to access
the target volume after the Copy Services function has completed.

The shell script to be run before the application that will use FlashCopy target should include
the operations shown in Example 8-9 on page 352.

Example 8-9   Backup preparation tasks

```
#quiesce an application
# insert the quiescing script here
#unmounting the source
umount /source
#start FlashCopy task
rsExecuteTask.sh -s CopyServicesServer EstablishTaskName
#check if FlashCopy task is established
rsQuery.sh -f VolumeList -s CopyServicesServer
#if yes, you can mount the source back
mount /source
#and resume the application
# insert the resuming script here
#check the target for consistency
fsck -y /dev/rdsk/cXtYdZsN
#if OK mount it
mount /dev/dsk/cXtYdZsN /target
```

The shell script to be run after backup is shown in Example 8-10.

Example 8-10   Post backup tasks

```
#unmount the target
umount /target
#terminate the FlashCopy pair if Do not perform background copy was used
rsExecuteTask.sh -s CopyServicesServer WithdrawTaskName
```

The foregoing steps can be performed for PPRC as well as FlashCopy by substituting the FlashCopy task with a PPRC terminate copy pair task on duplex volumes. Once the operation on the secondary volume has completed, then, rather than execute the FlashCopy withdraw, you can perform a PPRC establish copy pair task. This operation can be performed again once the volumes have become duplex again.

In the following example, we describe a method of accessing the FlashCopy target on a single SUN Solaris system. The FlashCopy source is active on that server at the same time. This example can also be applied to a PPRC secondary volume on a target SUN Solaris server.

In our example, there is a file system named /source on the source volume c1t6d0s2 of the FlashCopy pair. The target of the FlashCopy will be c1t6d1s2.

You can display all available ESS LUNs using the **rsList2105s.sh** command of the Copy Services CLI. This can be seen in Figure 8-15.

```
./rsList2105s.sh
disk name       2105 serial number
---------       ------------------
 c1t6d0            500FCA24
 c1t6d1            501FCA24
 c1t6d2            502FCA24
 c1t6d3            503FCA24
 c1t6d4            504FCA24
 c1t6d5            505FCA24
```

Figure 8-15   rsList2105s.sh on SUN Solaris

Make sure that the PPRC/FlashCopy source is in a consistent state for the short period during the establishment of the FlashCopy pair or prior to terminating the PPRC relationship. Use the following procedure to access the target volume:

1. Bring down applications that access the PPRC/FlashCopy primary/source. Unmount the related file systems for the short period of PPRC/FlashCopy termination/establishment:

   ```
   umount /source
   ```

2. Using the Copy Services Web Interface or the Command Line Interface, terminate the PPRC relationship or establish FlashCopy.

3. Once the PPRC/FlashCopy pair is terminated/established, mount all file systems and restart the applications. This could be done even if data is still copied from the source to the target in the background when using FlashCopy. Run:

   ```
   mount /dev/dsk/c1t6d0s2 /source
   ```

4. Check the consistency of the file system on the target volume:

   ```
   fsck -y /dev/rdsk/c1t6d1s2
   ```

5. Create a mount point for the target file system and mount the file system:

   ```
   mkdir /target
   mount /dev/dsk/c1t6d1s2 /target
   ```

Now the data on the target could be accessed. The **mount -v** output in Figure 8-16 shows that the source and target are active at the same time in the case of a FlashCopy between source and target volume on the same system.

```
/dev/dsk/c1t6d0s2 on /source type ufs read/write/setuid/largefiles...
/dev/dsk/c1t6d1s2 on /target type ufs read/write/setuid/largefiles...
```

Figure 8-16   Output from mount -v on Solaris

8.4.2 Copy Services with a VERITAS Volume Manager

In the following section, we describe how to perform FlashCopy and PPRC on SUN Solaris systems with VERITAS Volume Manager (VxVM) support.

FlashCopy with VERITAS Volume Manager

Here we explain how to simultaneously mount ESS FlashCopy source and target volumes to the same host without exporting the source volumes when using VERITAS Volume Manager.

It is assumed that the sources are constantly mounted to the SUN host, the FlashCopy is performed, and the goal is to mount the copy without unmounting the source or rebooting.

After the target volumes have been assigned, it is necessary to reboot the SUN server using **reboot -- -r** or, if a reboot is not immediately possible, then issue **drvconfig**, **disks**, and then **devlinks**. However, a reboot is recommended for guaranteed results.

It is also assumed that the appropriate actions in order to use the target volumes with the host have already taken place (that is, **devfsadm**, **vxdctl enable**...).

The following procedure refers to these names:

► mydg: As the name of the diskgroup that is being created.

► da_name: As the disk name shown under the DISK column in the **vxdisk list** output.

- ► last_daname: As the disk is known to VxVM as, as shown under DEVICE column in **vxdisk list** output.This is the output of vxdisk list on SUN Solaris.

Use the following procedure to mount the targets to the same host:

1. Determine which disks have a copy of the disk group configuration in their private region. The following command will list the *log disk* disks:

   ```
   # vxdg list <disk group>
   ```

2. Determine the location of the private region (tag 15) on the disks (normally partition 3):

   ```
   # prtvtoc /dev/rdsk/c#t#d#s2
   ```

   or use the following command to get the partition number for the private region:

   ```
   # vxdisk list c#t#d#s2 | grep priv
   ```

3. Dump the private region:

   ```
   # /usr/lib/vxvm/diag.d/vxprivutil dumpconfig /dev/rdsk/c#t#d#s3 > dg.dump
   ```

4. Create a script to initialize the disk group:

   ```
   # cat dg.dump | vxprint -D - -d -F "vxdg -g <mydg> adddisk %name=%last_da_name" > dg.sh
   ```

5. Edit the file dg.sh and change the first line to:

   ```
   # vxdg init <mydg>   <daname>=<last_daname>
   ```

6. Make the file dg.sh executable:

   ```
   # chmod 755 dh.sh
   ```

7. Create a file that can be used to rebuild the VM config:

   ```
   # cat dg.dump | vxprint -D - -hvpsm > dg.maker
   ```

8. Initialize the disk group by executing dg.sh:

   ```
   # ./dh.sh
   ```

9. If this results in the error `Disk is already in use by another system`, then the private region on each disk that is to be added to the disk group will need to be initialized. This can be done with the following command:

   ```
   # vxdisksetup -i <da_name>
   ```

10. Rebuild the VM configuration:

   ```
   # vxmake -g <mydg> -d dg.maker
   ```

11. Start the volumes

   ```
   # vxvol -g <mydg> start <volume>
   ```

PPRC with VERITAS Volume Manager

In the previous section, we described how to perform a FlashCopy and mount the source and target file system on the same server. Here we describe the steps necessary to mount a PPRC secondary volume onto a server that does not have sight of the primary volume.

It assumes that the PPRC copy pair has been terminated prior to carrying out the procedure.

After the secondary volumes have been assigned, it is necessary to reboot the SUN server using **reboot -- -r** or, if a reboot is not immediately possible, then issue **drvconfig**, **disks** and then **devlinks**. However, a reboot is recommended for guaranteed results.

Use the following procedure to mount the secondary volumes to another host:

1. Scan devices in the operating system device tree:

   ```
   vxdisk scandisks
   ```

2. List all known disk groups on the system:

   `vxdisk -o alldgs list`

3. Import the PPRC disk group information:

   `vxdg -C import <disk_group_name>`

4. Check the status of volumes in all disk groups:

   `vxprint -Ath`

5. Bring the disk group online:

   `vxvol -g <disk_group_name> startall`

   or

   `vxrecover -g <disk_group_name> -sb`

6. Perform a consistency check on the file systems in the disk group:

   `fsck -V vxfs /dev/vx/dsk/<disk_group_name>/<volume_name>`

7. Mount the file system for use:

   `mount -V vxfs /dev/vx/dsk/<disk_group_name>/<volume_name> /<mount_point>`

Once you have finished with the PPRC secondary volume, it is recommended that you perform the following tasks:

1. Unmount the file systems in the disk group:

   `umount /<mount_point>`

2. Offline the volumes in the disk group:

   `vxvol -g <disk_group_name> stopall`

3. Export disk group information from the system:

   `vxdg deport <disk_group_name>`

> **Tip:** If you FlashCopy or PPRC only one half of a RAID-1 mirror, it will be necessary to force the import of the disk group because not all of the disks are available. Therefore, it is necessary to issue the following command:
>
> `vxdg -f import <disk_group>`
>
> However, be aware that this may cause disk group inconsistencies.

8.5 HP-UX and Copy Services

The following section describes how it is possible to access a source and target Copy Services volume on the same HP server.

8.5.1 HP-UX with FlashCopy

The following procedure must be followed to permit access to the FlashCopy source and destination simultaneously on an HP-UX host. It could be used to make an additional copy of a development database for testing or to permit concurrent development, to create a database copy for data mining that will be accessed from the same server as the OLTP data, or to create a Point-In-Time copy of a database for archiving to tape from the same server. This procedure must be repeated each time you perform a FlashCopy and want to use the target physical volume on the same host where the FlashCopy source volumes are present in the Logical Volume Manager configuration.

You can use the Copy Services CLI for automation and create scripts to automate your procedures. If you are preparing scripts, you must also prepare your FlashCopy tasks and test them.

Target preparation

In order to prepare the target system, carry out the following steps:

1. Vary off the source volume groups:

   ```
   vgchange -a n /dev/<source_vg_name>
   ```

2. If you did not use the default Logical Volume Names (lvo*ln*n) when they were created, create a map file from your source volume group using the **vgexport** command:

   ```
   vgexport -m <map file name> -p /dev/<source_vg_name>
   ```

   Note: This map file needs to be "ftp'd" to the target host.

3. If the target volume group exists, remove it using the **vgexport** command. The target volumes cannot be members of a volume group when the **vgimport** command is run:

   ```
   vgexport -m /dev/null /dev/<target_vg_name>
   ```

4. Shut down or quiesce any applications that are accessing the FlashCopy source.

FlashCopy execution

To execute the procedure, you must carry out the following steps:

1. Unmount all file systems in the source volume group.

2. Perform the FlashCopy.

3. Ensure that the FlashCopy task has completed using the ESS Specialist Web Interface or the Command Line Interface command **rsQuery.sh**:

   ```
   ./rsQuery.sh -u <CLI_User> -p <CLI_Password> -f disklist -s\
   <primary_copy_services_server>
   ```

4. Mount all the file systems in the source volume group.

5. When the FlashCopy is finished, change the Volume Group ID on each ESS Volume in the FlashCopy target. The volume ID for each volume in the FlashCopy target volume group must be modified on the same command line. Failure to do this will result in a mismatch of Volume Group IDs within the volume group. The only way to resolve this issue is to perform the FlashCopy again and reassign the Volume Group Ids using the same command line:

   ```
   vgchgid -f </dev/rdsk/c#t#d#_1>...</dev/rdsk/c#t#d#_n>
   ```

   Note: This step is not needed if another host is used to access the target devices.

6. Create the Volume Group for the FlashCopy target:

   ```
   mkdir /dev/<target_vg_name>
   mknod /dev/<target_vg_name>/group c <lvm_major_no> <next_available_minor_no>
   ```

   Use the **lsdev -C lvm** command to determine what the major device number should be for Logical Volume Manager objects. To determine the next available minor number, examine the minor number of the group file in each volume group directory using the **ls -l** command.

7. Import the FlashCopy target volume(s) into the newly created volume group using the vgimport command:

```
vgimport -m <map file name> -v /dev/<target_vg_name>
</dev/dsk/c#t#d#_1>...</dev/dsk/c#t#d#_n>
```

8. Activate the new volume group:

```
vgchange -a y /dev/<target_vg_name>
```

9. Perform a full file system check on the logical volumes in the target volume group. This is necessary in order to apply any changes in the JFS intent log to the file system and mark the file system as *clean*.

```
fsck -F vxfs -o full -y /dev/<target_vg_name>/<logical volume name>
```

10. If the logical volume contains a VxFS file system, mount the target logical volume(s) on the server:

```
mount -F vxfs /dev/<target_vg_name>/<logical volume name><mount point>
```

The example in Figure 8-17 on page 359 shows how to import a FlashCopy target on the same system as the source but with default logical volume naming scheme.

Once access to the FlashCopy target volume is no longer required, unmount the file systems and vary off the volume group.

```
vgchange -a n /dev/<target_vg_name>
```

If no changes are made to the source volume group prior to the subsequent FlashCopy, then all that is needed is to vary on the volume group and perform a full file system consistency check, as shown in steps 8 to 10.

8.5.2 HP-UX with PPRC

When using PPRC with HP-UX, it is similar to using FlashCopy, apart from the fact that the volume group should be unique to the target server, so there should be no need to perform the `vgchgid` command to change the physical volume to volume group association. Here is the procedure to bring secondary volumes online to PPRC target HP-UX hosts:

1. Quiesce the source HP-UX application to cease any updates to the primary volumes.

2. If you are using PPRC-XD, allow the copy pair volumes to go into duplex state using the catch-up operation or by leaving the volumes to become synchronized.

3. Terminate the PPRC copy pair relationship, using the Copy Services Web User Interface or by using the CLI.

4. Rescan for hardware configuration changes using the `ioscan -fnC disk` command. Check that the disks are *CLAIMED* using `ioscan -funC disk`. The reason for doing this is that the volume group may have been extended to include more physical volumes.

5. Create the Volume Group for the PPRC secondary. Use the `lsdev -C lvm` command to determine what the major device number should be for Logical Volume Manager objects. To determine the next available minor number, examine the minor number of the group file in each volume group directory using the `ls -l` command.

6. Import the PPRC secondary volume(s) into the newly created volume group using the `vgimport` command.

7. Activate the new volume group.

8. Perform a full file system check on the logical volumes in the target volume group. This is necessary in order to apply any changes in the JFS intent log to the file system and mark the file system as *clean*.

9. If the logical volume contains a VxFS file system, mount the target logical volume(s) on the server.

If changes are made to the source volume group, they should be reflected in the /etc/lvmtab of the target server. Therefore, it is recommended that periodic updates be made to make the lvmtab on both source and target machines consistent. As with the AIX `importvg`, there are two alternatives:

► Using the **Permit read from secondary** option:

   a. If you are using PPRC-XD, issue `go-to-sync` to allow the volumes to go to duplex state.

   b. Once the volumes are in duplex state, suspend the primary volume so that no updates are reflected on the secondary volumes.

   c. Export the source volume group information into a map file.

   d. Export the old volume group definitions from the target host.

   e. Run an `ioscan` to identify any new volumes that have been assigned to the target hosts due to expansion of the source volume group.

   f. Import the target volume group definition using the map file generated from the source host.

   g. Re-establish the PPRC relationship, only copying cylinders out-of-sync.

► Using NOCOPY PPRC establish:

   a. Quiesce all write I/O to the primary volumes, and unmount the source file systems.

   b. If you are using PPRC-XD, issue the `go-to-sync` command to all the secondary volume to catch up.

   c. Once the volumes are in duplex state, terminate the PPRC relationship.

   d. Export the source volume group information into a map file.

   e. Export the old volume group definitions from the target host.

   f. Run an `ioscan` to identify any new volumes that have been assigned to the target hosts due to expansion of the source volume group.

   g. Import the target volume group definition using the map file generated from the source host.

   h. Establish the PPRC relationship with the NOCOPY option, so that only the primary and secondary have a copy pair relationship without updates.

   i. Immediately suspend the primary volumes.

   j. Mount the file systems and start the application at source.

   k. Some time later, re-establish PPRC, only copying the out-of-sync cylinders.

Refer to the code sample in Figure 8-17 on page 359, which shows the import of a FlashCopy target.

```
# umount /test
# ./rsExecuteTask.sh -u storwatch -p specialist -s sls6c1 HPFlash
# ./rsQuery.sh -u storwatch -p specialist -f disklist -s sls6c1
16-Oct-02 3:42:16 PM rsClientImpl: Successfully connected to Name Server, port = 1703
16-Oct-02 3:42:17 PM rsClientImpl: Server identity = IBM2105 Copy Services/1.1.0
16-Oct-02 3:42:18 PM rsClientImpl: Received rsVSServer reference successfully
16-Oct-02 3:42:19 PM rsClientImpl: rsClientImpl registered successfully
*************************Volume Information*************************
Volume 20014744 found on 14744:12 as volume number 000
        State=simplex, status=not_suspended, FlashCopy_state=source, Size=20.0_GB
*******************************************************************
************************Volume Information************************
Volume 20114744 found on 14744:12 as volume number 001
        State=simplex, status=not_suspended, FlashCopy_state=target, Size=20.0_GB
*******************************************************************
# mount /test
# mkdir /dev/vgibm
# mknod /dev/vgibm/group c 64 0x040000
# vgimport -v /dev/vgibm /dev/dsk/c0t6d2
Beginning the import process on Volume Group "/dev/vgibm".
Logical volume "/dev/vgibm/lvol1" has been successfully created
with lv number 1.
Volume group "/dev/vgibm" has been successfully created.
Warning: A backup of this volume group may not exist on this machine.
Please remember to take a backup using the vgcfgbackup command after activating
the volume group.
# vgchange -a y /dev/vgibm
Activated volume group
Volume group "/dev/vgibm" has been successfully changed.
# vgcfgbackup /dev/vgibm
Volume Group configuration for /dev/vgibm has been saved in /etc/lvmconf/vgibm.c
onf
# fsck -F vxfs -p -y -o full /dev/vgibm/lvol1
/dev/vgibm/lvol1:log replay in progress
/dev/vgibm/lvol1:pass0 - checking structural files
/dev/vgibm/lvol1:pass1 - checking inode sanity and blocks
/dev/vgibm/lvol1:pass2 - checking directory linkage
/dev/vgibm/lvol1:pass3 - checking reference counts
/dev/vgibm/lvol1:pass4 - checking resource maps
/dev/vgibm/lvol1:OK to clear log? (ynq)y
/dev/vgibm/lvol1:set state to CLEAN? (ynq)y
# mount /dev/vgibm/lvol1 /ibm
```

Figure 8-17   Import of a FlashCopy target

8.6 OpenVMS with Copy Services

In the following section, we discuss the necessary tasks to perform when bringing a Copy Services target volume online on an HP OpenVMS server.

8.6.1 OpenVMS with FlashCopy

Under OpenVMS, volumes have labels associated to them. The volume label acts as a unique volume identification cluster-wide. When you perform a FlashCopy, all the disk structures get copied from the source, so when the target is brought online, it will have the same information as the source and therefore the same volume label. This causes a volume

label conflict when the FlashCopy target should be mounted at the same host or in a host belonging to the same OpenVMS cluster.

There are two ways of overcoming this problem:

► Dismount the source volume and mount the target volume, therefore not having two volumes with the same volume label mounted at the same time.

► Mount the target volume for private access only by a single process.

   If the target volume is needed only for some specific operation (like creating a backup), then this can be done in that process. If the target is needed for global access, then you change the target volume label, dismount it, and mount it again for global access.

> **Note:** Care must be taken on the sizing and placement of volumes in a SCSI attached environment, because OpenVMS supports eight LUN-IDs per SCSI target, giving a maximum of 120 volumes per adapter. In a Fibre Channel attached environment, the User Defined ID (UDID) must be unique throughout the OpenVMS cluster.

The next procedure shows how you can perform a FlashCopy on a source volume and the steps required for changing the volume label on the target and mounting it for global access:

1. Once the target volumes have been assigned using the ESS StoreWatch Specialist, it is necessary to configure them to the system using the following command:

   ```
   mcr sysman io autoconfigure/log
   ```

   > **Note:** You need the OPER privilege to run the SYSMAN utility. Additionally, you must have CMKRNL and SYSLCK privileges to use the **mcr sysman io autoconfigure** command.

2. Identify whether the new volumes assigned to be the FlashCopy targets are visible to the system, as shown in Figure 8-18, using the following command:

   ```
   show devices dk
   ```

```
$ show device dk

Device         Device      Error   Volume   Free       Trans Mnt
Name           Status      Count   Label    Blocks     Count Cnt
ES40$DKA0:     Mounted       0     ES4073   14988114     1    1
ES40$DKA100:   Mounted       0     ES40_73  13801842   312    1
ES40$DKA200:   Online        0
ES40$DKB200:   Mounted      43     E200     19144495     1    1
ES40$DKB207:   Mounted      43     E207     19142802     1    1
ES40$DKB600:   Mounted      43     E600     16341522     1    1
ES40$DKB601:   Online       43
```

Figure 8-18   Output showing attached SCSI disks

The **show device** command shows you devices of a particular type. In this case, it is a disk device with a SCSI (dk) attachment. If you wanted to see disk devices with Fibre Channel attachment, then dg would be the device specifier (see Figure 8-19 on page 361).

The best way of relating the attached devices to volumes configured in the ESS is to use the Copy Services CLI command **rsList2105s**. However, if you are using SCSI attached, you can translate the physical device name, for example, ES40$DKB601:

– The letters before the $ sign (ES40) are the OpenVMS host name.

– The first two letters after the $ sign (DK) specify the device and connection type as SCSI disk.

– The third letter (B) denotes the SCSI adapter placement.

- The final three digits indicate the SCSI address, where the first digit is the SCSI target ID and the last two digits are the LUN ID. The volumes used in this example were ES60$DKB600 (source) and ES40$DKB601 (target).

The device name for disks with Fibre Channel attachment has the structure 1DGAnnnnn (see Figure 8-19):

- The $1 is a so-called storage allocation classes, replacing the OpenVMS host name.

- The first two letters after the $ sign (DG) specify the device and connection type as a Fibre Channel disk.

- The third letter is always A for any FC disk.

- The final digits (up to five) are the *User-Defined ID* (UDID), which is stored as a label in the ESS. (This label is an ESS-internal identification string, different from the operating system's volume label.)

```
VMS1$ show device dg
Device              Device          Error     Volume      Free  Trans Mnt
 Name               Status          Count      Label     Blocks Count Cnt
$1$DGA2001:    (VMS1)  Mounted          0  QUORUM     1952886      1   1
$1$DGA2002:    (VMS1)  Mounted          0  V73_COMMON 1067115    357   1
$1$DGA2003:    (VMS1)  Online           0
$1$DGA2004:    (VMS1)  Online           0
$1$DGA2005:    (VMS1)  Online           0

VMS1$
```

Figure 8-19   Output showing attached FC disks

An ESS volume label that represents a number in the range 0–32767 is presented as UDID to OpenVMS hosts. Figure 8-20 on page 362 shows the volume panel with the ESS volume labels 2001, 2002, and 2003.
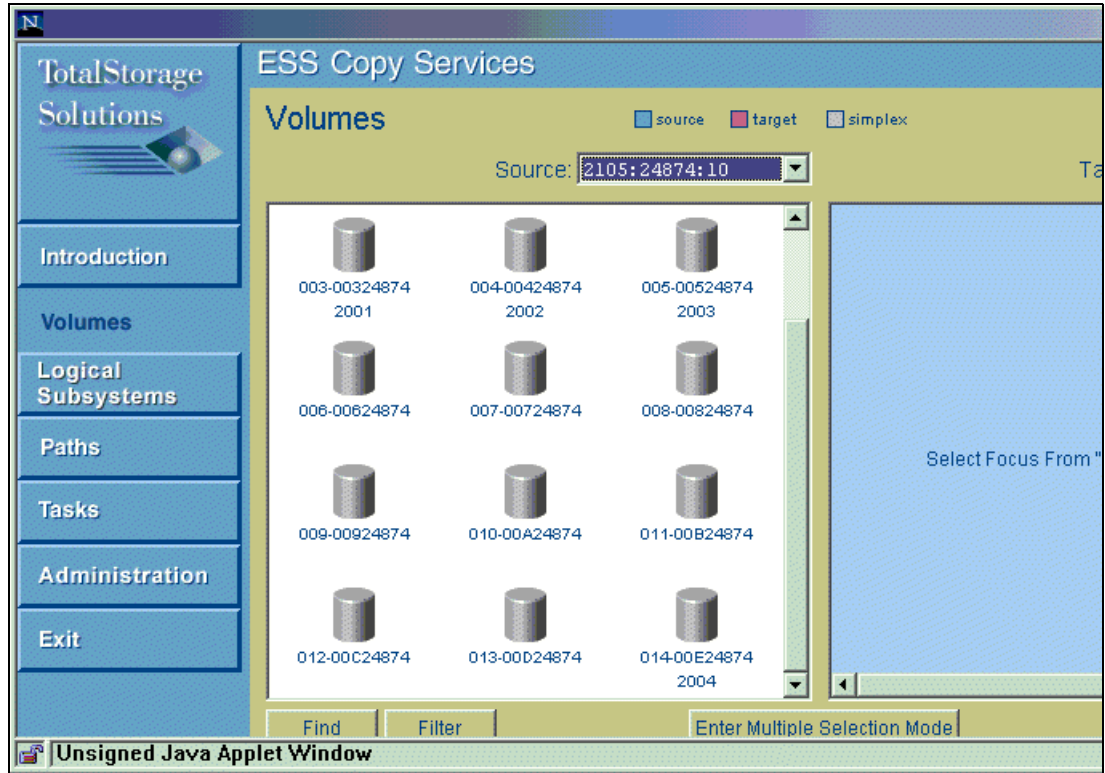
Figure 8-20   OpenVMS volumes with UDID labels

3. Perform the FlashCopy operation using the Copy Services Web User Interface (Figure 8-21 on page 363) or via the CLI. Please note the quotation marks around user name, password, and task name. These items are case-sensitive in the ESS, but the OpenVMS command interpreter would convert them to uppercase without quotation marks. (Currently, there is only the Copy Services CLI available for OpenVMS.) Run:

```
rsExecuteTask -u "storwatch" -p "specialist" -s sls6c1 "OpenVMSFlash"
```
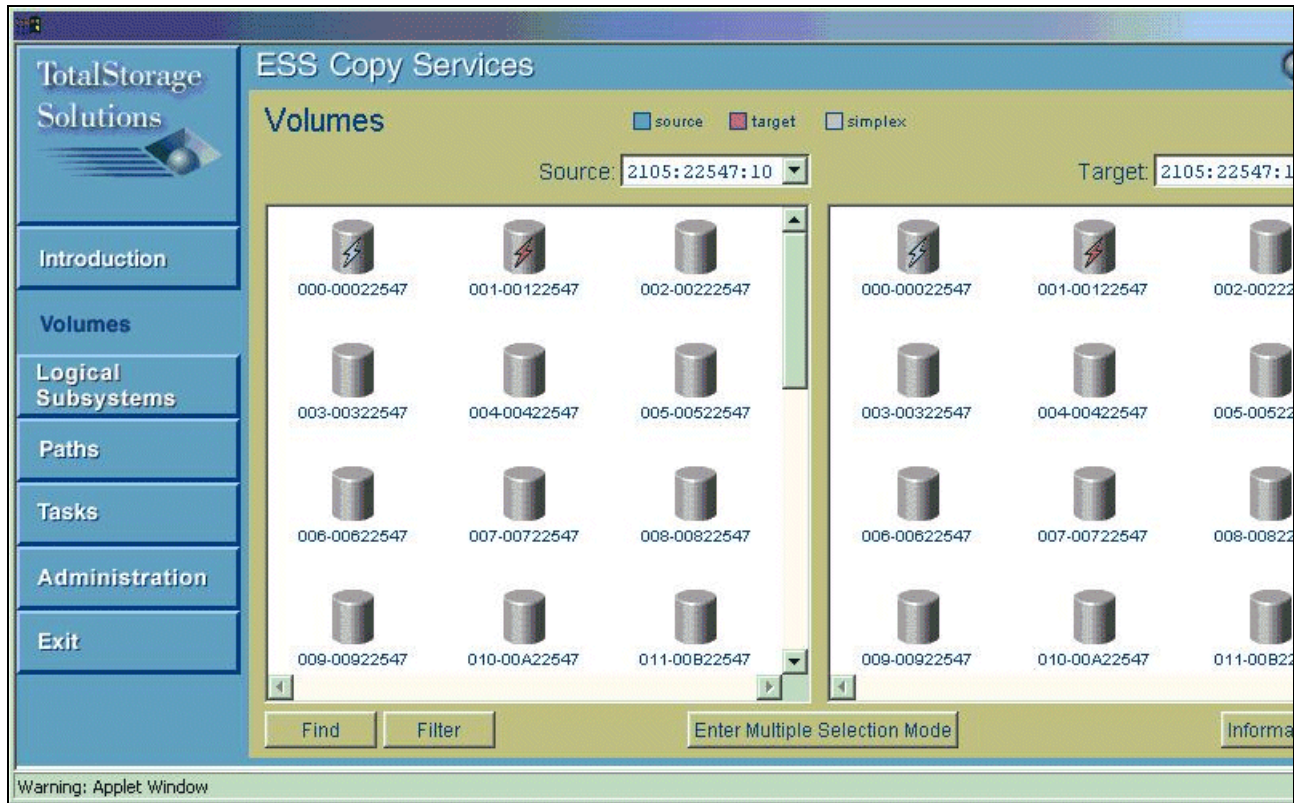
Figure 8-21   Established source and target

Now that the FlashCopy volumes are established, it is possible to access the target volume in read-only mode. This is shown in Figure 8-22. It is not strictly necessary to mount the target volume as read-only at the first time, but we recommend that as "good practice" to look at the volume before doing anything. The important thing is that this mount operation was a private one (not system-wide or cluster-wide), making the volume accessible from only that process that issued the **mount** command.

```
$ mount /nowrite dkb601: e600
%MOUNT-I-MOUNTED, E600 mounted on _ES40$DKB601:
$ dir dkb601:[000000]

Directory DKB601:[000000]

000000.DIR;1         BACKUP.SYS;1        BADBLK.SYS;1        BADLOG.SYS;1
BITMAP.SYS;1         CONTIN.SYS;1        CORIMG.SYS;1        FLASH_TEST.TXT;1
INDEXF.SYS;1         REPEAT_COPY_00000433_1.TMP;1
REPEAT_COPY_00000437_1.TMP;1             REPEAT_COPY_00000438_1.TMP;1
SECURITY.SYS;1       VOLSET.SYS;1

Total of 14 files.
```

Figure 8-22   Checking consistency of FlashCopy target volume

Although we mount the FlashCopy target volume, we must refer to it with the same OpenVMS volume label as assigned to the source volume. Figure 8-23 on page 364 shows that although the FlashCopy target is mounted, it has a writelock initiated. The keyword alloc indicates that the device is allocated for private access by a single process.

```
$ show device dk

Device                 Device          Error   Volume    Free        Trans Mnt
 Name                  Status          Count   Label     Blocks      Count Cnt
ES40$DKA0:             Mounted             0   ES4073    14988114      1   1
ES40$DKA100:          Mounted             0   ES40_73   13801842    312   1
ES40$DKA200:          Online              0
ES40$DKB200:          Mounted            43   E200      19144495      1   1
ES40$DKB207:          Mounted            43   E207      19142802      1   1
ES40$DKB600:          Mounted            43   E600      16341488      1   1
ES40$DKB601:          Mounted alloc      43   E600      15221698      1   1
                      WriteLocked

$
```

Figure 8-23   Status of FlashCopy mount

Now that the FlashCopy target contents have been verified, it is necessary to dismount the target volume and remount it for write access, again for the use by a specific process, in order to not conflict with the source that is mounted for global access.

We notice that the message, %MOUNT-I-REBUILD, volume was improperly dismounted; rebuild in progress in Figure 8-24 comes from the nonzero mount counter on the target volume: OpenVMS manages a mount counter in every volume's home block. Because this number was greater than zero on the source volume when creating the copy, it is also greater than zero on the target volume. Therefore, when mounting the target volume with write access for the first time, OpenVMS suspects an incorrect dismount, so it checks and repairs the volume. (This is suppressed when mounting with the /nowrite qualifier, because this check and repair always requires write access.)

```
$ dismount dk601:
$ mount/override=id dkb601:
%MOUNT-I-MOUNTED, E600 mounted on _ES40$DKB601:
%MOUNT-I-REBUILD, volume was improperly dismounted; rebuild in progress
$ show device dk

Device                 Device          Error   Volume    Free        Trans Mnt
 Name                  Status          Count   Label     Blocks      Count Cnt
 ES40$DKA0:            Mounted             0   ES4073    14988114      1   1
 ES40$DKA100:         Mounted             0   ES40_73   13801842    312   1
 ES40$DKA200:         Online              0
 ES40$DKB200:         Mounted            43   E200      19144495      1   1
 ES40$DKB207:         Mounted            43   E207      19142802      1   1
 ES40$DKB600:         Mounted            43   E600      16341488      1   1
 ES40$DKB601:         Mounted alloc      43   E600      15221698      1   1
 $
```

Figure 8-24   Mount FlashCopy target for write

4. The target still retains the volume label that was copied over from the FlashCopy source, however, because the volume is allocated for exclusive use by the ID, there is no conflict. Since the volume is mounted in exclusive read/write mode, we can change the volume label to what we desire. This is shown in Figure 8-25 on page 365.

```
$ set volume/label=e601 dkb601:
$ show device dk

Device              Device          Error   Volume    Free       Trans Mnt
 Name               Status          Count   Label     Blocks     Count Cnt
ES40$DKA0:          Mounted             0   ES4073    14988114    1    1
ES40$DKA100:        Mounted             0   ES40_73   13801842   312   1
ES40$DKA200:        Online              0
ES40$DKB200:        Mounted            43   E200      19144495    1    1
ES40$DKB207:        Mounted            43   E207      19142802    1    1
ES40$DKB600:        Mounted            43   E600      16341488    1    1
ES40$DKB601:        Mounted alloc      43   E601      15221698    1    1
$
```

Figure 8-25   Change FlashCopy target volume label

5. Since the target volume has now a unique label, it is possible to change the access to system-wide or cluster-wide so that other processes can gain access to the data on the volume (see Figure 8-26).

```
$ dismount dk601:
$ mount/system dkb601: e601
%MOUNT-I-MOUNTED, E601 mounted on _ES40$DKB601:
$ show device dk

Device              Device          Error   Volume    Free       Trans Mnt
 Name               Status          Count   Label     Blocks     Count Cnt
ES40$DKA0:          Mounted             0   ES4073    14988114    1    1
ES40$DKA100:        Mounted             0   ES40_73   13801842   312   1
ES40$DKA200:        Online              0
ES40$DKB200:        Mounted            43   E200      19144495    1    1
ES40$DKB207:        Mounted            43   E207      19142802    1    1
ES40$DKB600:        Mounted            43   E600      16341488    1    1
ES40$DKB601:        Mounted            43   E601      15221698    1    1
$
```

Figure 8-26   Providing system-wide access to FlashCopy target

Note: You need the SYSPRV and SYSNAM privileges to mount a volume for system-wide or cluster-wide access.

Subsequent FlashCopies to the target volume will require the target to be dismounted, whether it be the same or another host system.

8.6.2 OpenVMS and PPRC

When we assume that the PPRC secondary volume is mounted at a different host that is not in the same OpenVMS cluster as the host with the primary volume, then there is no volume label conflict. Therefore, there is no need to change the volume label prior to mounting it for system use.

The following procedural steps are almost exactly the same as those of the FlashCopy; however, certain steps have been excluded due to the assumption that the primary and the secondary volume will not be on the same server/cluster:

1. Quiesce write I/O to the PPRC primary volume.

2. Ensure that the PPRC copy pair is in a duplex state. If you are using PPRC-XD, perform the catch-up operation (see Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53).

3. Terminate the PPRC copy pair relationship, using the Copy Services Web Interface or the CLI.

4. If the secondary volume is not recognized by the other host, use the `mcr sysman io autoconfigure` command to configure the volume.

5. Verify that the secondary volume is recognized by the target system using the `show dev dk` (or `dg` for Fibre Channel) command.

6. Mount the secondary volume for system-wide access using the `mount/system` command (or for cluster-wide access using the `mount/cluster` command).

> **Important:** Prior to re-establishing the PPRC copy pair, ensure that you `dismount` the PPRC secondary volume from the target system. Do not mount a PPRC secondary volume while it is established in a copy pair relationship, as this will cause a system crash. This warning applies also if the **Permit read from secondary** option has been activated.
>
> To prohibit mount operations (which might be performed even by non-privileged users), you should set the secondary volume as considered unavailable using the command `set device/noavailable`.

8.7 Tru64 and Copy Services

In this section, we describe how to perform Copy Services functions such as FlashCopy and PPRC on an HP Tru64 system. We explain how to perform Copy Services functions on volumes in a clustered environment and how to deal with *Persistent Reserves* placed on the target volumes. We also discuss how to handle Advance Filesystem (AdvFS) file systems so Copy Services target volumes containing backup data can be mounted under the same domain.

8.7.1 Tru64 and FlashCopy

The following section describes the procedure to perform FlashCopy on a Tru64 5.x system consisting of UNIX Filesystem (UFS) and Advance Filesystem (AdvFS) file systems. We explain how to mount the source and target FlashCopy volumes containing UFS file systems on a Tru64 5.x server.

In a Tru64 5.x clustered environment, ownership of shared SCSI devices is controlled by placing Persistent Reserves. The symptom of this is that initiation of Copy Services tasks on target volumes fail. Therefore, prior to initiating any Copy Services tasks on target volumes, it is necessary to check for reserves placed on target volumes and deciding whether to release them prior to instigating a copy function.

The procedure uses the essvol script shown in Figure 8-27 on page 367 to identify assigned ESS volumes.

```
echo Extracting ESS volume information...
for ID in `hwmgr -view dev -cat disk | egrep -e 2105F20  -e 2105800 | awk '{ print $1}'`
do
echo; echo ESS vol, H/W ID $ID
hwmgr -get attrib -id $ID | awk '/phys_loc//dev_base//capacity//serial/'
done
```

Figure 8-27   essvol script

UFS Filesystem on Tru64 and FlashCopy

The following procedure describes the steps necessary to perform a FlashCopy to a target volume and mount it on the same server as the source volume.

1. Check whether the FlashCopy target is assigned to the Tru64 host (Figure 8-28):

    ```
    ./essvol
    ```

```
<alpha1> essvol
ESS vol, H/W ID 176:
  phys_location = bus-9-targ-0-lun-0
  dev_base_name = dsk43
  capacity = 3906304
  serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-1042-1380

ESS vol, H/W ID 225:
  phys_location = bus-9-targ-0-lun-7
  dev_base_name = dsk47
  capacity = 3906304
  serial_number = SCSI-WWID:01000010:6000-1fe1-4942-4d20-0000-0000-10a2-1380
```

Figure 8-28   Output of essvol script

2. Use the serial_number field to identify the volumes being used from the host side. The last eight digits denote the volume being used, and the breakdown for 1042-1380 is as follows:

    –   104 denotes the volume 04 on LSS 11.
    –   2-1380 denotes the ESS serial number of 21380.

3. Quiesce all write I/O and unmount file systems from source volumes

4. Ensure that the target volume has no file systems mounted.

5. Check and clear the persistent reserve on target volume if in a clustered environment (see 8.7.3, "Managing Tru64 Persistent Reserve" on page 370).

6. Perform FlashCopy from dsk43 (ESS volume 104) to dsk47 (ESS volume 10a).

7. Remount the file systems in the FlashCopy source volume.

8. Once the FlashCopy has established, perform a consistency check on the target volume file systems:

    ```
    fsck /dev/disk/dsk47c
    ```

9. Mount the file systems on the target volume:

    ```
    mount /dev/disk/dsk47c /ess
    ```

This procedure is the same if you are mounting the FlashCopy target onto a different Tru64 server from the FlashCopy target.

AdvFS Filesystem on Tru64 and FlashCopy

The following procedure describes the steps necessary to mount a FlashCopy target AdvFS volume to the same or another server in the same cluster. To mount the FlashCopy target, it is necessary to manually create the AdvFS domain/fileset and link it to the FlashCopy target volume.

For our example (shown in Figure 8-29), we use an AdvFS volume consisting of three ESS disks. The FlashCopy source AdvFS domain is called $DOM1$ with a fileset name of advol1 and consists of ESS disks: dsk58, dsk59, and dsk60. The FlashCopy target disks are dsk62, dsk63, and dsk64, and will belong to AdvFS domain DOM2.

1. Ensure that the FlashCopy source and target volumes are identical in size.

2. Check the contents of AdvFS domain DOM1.

```
#cd /etc/fdmns
# ls -l DOM1
lrwxr-xr-x  1 root     system        16 Sep 23 11:23 dsk58c -> /dev/disk/dsk58c
lrwxr-xr-x  1 root     system        16 Sep 23 11:24 dsk59c -> /dev/disk/dsk59c
lrwxr-xr-x  1 root      system        16 Sep 23 11:25 dsk60c -> /dev/disk/dsk60c
```

Figure 8-29   Contents of AdvFS domain DOM1

3. Manually create the AdvFS domain used for the FlashCopy target volumes (Figure 8-30).

```
# mkdir /etc/fdmns/DOM2
# cd /etc/fdmns/DOM2
# pwd
/etc/fdmns/DOM2
# ln -s /dev/disk/dsk62c
# ln -s /dev/disk/dsk63c
# ln -s /dev/disk/dsk64c
# ls -al
drwxr-xr-x  2 root     system      8192 Sep 23 14:04 .
drwxr-xr-x 13 root     system      8192 Sep 23 14:04 ..
lrwxrwxrwx  1 root     system        16 Sep 23 14:04 dsk62c -> /dev/disk/dsk62c
lrwxrwxrwx  1 root     system        16 Sep 23 14:04 dsk63c -> /dev/disk/dsk63c
lrwxrwxrwx  1 root      system        16 Sep 2314:04 dsk64c ->/dev/disk/dsk64c
```

Figure 8-30   Manual AdvFS domain creation

4. Unmount the file systems associated to the FlashCopy source volumes.

5. Check and clear the persistent reserve on the target volume if in a clustered environment (see 8.7.3, "Managing Tru64 Persistent Reserve" on page 370)

6. Perform FlashCopy between the source (dsk58, dsk59 and dsk60) and target (dsk62, dsk63 and dsk64) volumes.

7. Once the FlashCopy has been established, you can proceed to remount the file systems on the FlashCopy source volumes or mount the file systems on the target volumes:

   ```
   mount -t advfs DOM2#advol1 /ess/advol2
   ```

   File systems on both source and target cannot be mounted simultaneously on the same system or cluster. This because both the source and target would have the same AdvFS domain ID. The domain ID can be seen using the **advfsstat** command (Figure 8-31 on page 369).

```
# advfsstat DOM1

Domain -1032772981.728915- Stats -

Deref 9          Refhit     9        Refhitwait 0          R_ahead      0
Unpin 0          Pinhit     0        Pinhitwait 0          Pinreads     0
Lazy  0          Log        0        Blocking   0          Clean        0
                 Ubchit     9        Unconsol   0          ConsolAbort 0
                 UnpinMeta 0         UnpinFtx   0          UnpinData   0
                 DerefMeta 9         DerefFtx   9          DerefData   0

Disk      Reads      Writes      Rglobs    AveRglob    Wglobs    AveWglob
----      -----      ------      ------    --------    ------    --------
   1          0           0           0           0         0           0
```

Figure 8-31   Identifying AdvFS domain ID

Note: Do not try to mount a FlashCopy target volume containing the AdvFS file system on the same system or one within the same cluster, because the domain ID contained on the volume will be identical to that on the source.

8.7.2 Tru64 and PPRC

In this section, we discuss the steps necessary to perform PPRC remote volume replication from a primary volume assigned to a Tru64 5.x system and how to mount the secondary on a subsequent system. This section, like the previous one, covers both the UFS and AdvFS file systems under Tru64.

UFS Filesystem on Tru64 and PPRC

Performing volume replication using PPRC on volumes containing UFS file systems is the same as with FlashCopy, ensuring that you quiesce applications and unmount file systems prior to performing split mirror operations for remote vaulting or the termination of pairs when bringing the secondary volumes online for operation. The reason for quiescing all activity on the primary volumes is to maintain consistency, as opposed to any restrictions on the primary volumes.

Therefore, the process for performing volume replication using PPRC with UFS would be as follows:

1. Establish PPRC paths between the LSSs that contain primary and secondary volumes, taking care that there can only be a maximum of eight paths between each LSS pair. Optionally, choose to have a consistency group relationship between the LSS pairs.

2. Establish the primary and secondary copy pair relationship. Choose the method of synchronization, whether synchronous or asynchronous.

3. Wait for the copy pairs to change to duplex state, indicating that data on the secondary volume is the same as on the primary.

   If you instigated a PPRC-XD relationship, then it is necessary to either perform a catch-up operation using `go-to-sync` or to halt all updates to the primary volume so that in time all the contents of the primary volume are reflected on the secondary.

4. Now that the volumes are in synchronous state, stop the application and terminate the pairs.

5. If the volumes are assigned but not recognized by the system kernel, then use:

```
hwmgr -scan scsi
hwmgr -show scsi
```

6. Perform a consistency check on the secondary volume file systems:

```
fsck /dev/disk/dsk47c
```

7. Mount the file systems residing on the secondary volume:

```
mount /dev/disk/dsk47c /ess
```

AdvFS Filesystem on Tru64 and PPRC

The same rules apply with PPRC as with FlashCopy when using AdvFS file systems. It is not possible to mount the secondary volume onto a server that is part of the same cluster as the server on which the primary volume resides. This is due to the Domain/Fileset concepts in the AdvFS file system, which uses a unique domain ID that resides with the volume data structures and since the secondary is a volume level replication of the source, mounting that onto the same server/cluster will cause conflicts.

The procedure for performing PPRC with an AdvFS file system is done as with FlashCopy:

1. Ensure that the primary and secondary volumes are identical in size.

2. Establish the paths between the primary and secondary LSSs, making sure to limit the number of paths to eight per LSS.

3. Check and clear any persistent reservations on the secondary volumes.

4. Establish primary and secondary volume PPRC relationships in synchronous or asynchronous (PPRC-XD) copy mode.

5. Wait for the volumes to go into duplex (synchronized) mode. If you are using PPRC-XD, issue `go-to-sync` and wait for the secondary to catch-up. Monitor the progress of the volumes using the Copy Services Web User Interface or the `rsQuery.sh` command.

6. Once the volumes are in sync, stop all write I/O or quiesce the application and then terminate the copy pairs.

7. If the volumes are assigned from the ESS but not recognized by the system, issue:

```
hwmgr -scan scsi
hwmgr -show scsi or ./essvol
```

8. Manually create the AdvFS domain used for the PPRC secondary volumes under the /etc/fdmns directory and create symbolic links to the secondary volume block special file under /dev/disk.

9. If the secondary server is not part of the same cluster as the primary, it is possible to have both the PPRC primary and secondary volumes mounted at the same time:

```
mount -t advfs DOM2#advol1 /ess/advol2
```

8.7.3 Managing Tru64 Persistent Reserve

In the clustered environment, Tru64 5.x hosts place a Persistent Reserve (PR) on a volume as soon as the cluster can access it. Persistent Reserve means that no other system except the cluster members themselves are able to access the volume.

> **Important:** FlashCopy and PPRC will fail if the target volume contains a Persistent Reserve.

Concept of Persistent Reserve

Within the cluster environment, each cluster member creates a key entry in the Persistent Reserve (PR) table that is held by the disk. This table can be viewed with the scu utility:

```
# scu -f /dev/rdisk/dskxxc show keys
```

One cluster member holds the reservation. Its' active reservation key can also be viewed with the scu utility:

```
# scu -f /dev/rdisk/dskxxc show reservations
```

A Tru64 cluster sets the reservation in such a way that all cluster members have read and write access to the volume and all other hosts have read access only. All cluster members can also modify the PR state (for example clear the PR) using *their own* reservation key.

For example, we have a Host A that uses the key 0x10001 and Host B with the key 0x10002. Check the active reservation using the **scu** command, as seen in Figure 8-32. Host A is the owner of the existing reservation.

```
# scu -f /dev/rdisk/dsk45c show reservations
Persistent Reservation Header:

                  Generation Value: 4
                 Additional Length: 16

Reservation Descriptors:

                   Reservation Key: 0x10001
            Scope-Specific Address: 0
                  Reservation Type: 0x5 (Write Exclusive Registrants Only)
                 Reservation Scope: 0 (LU - full logical unit)
                     Extent Length: 0
```

Figure 8-32   Checking for Persistent Reserve

Scenario 1

To remove the Persistent Reserve (PR) from the volume by using the host who owns it, Host A in this example, issue the command:

```
# scu -f /dev/rdisk/dsk45c pres clear key 0x10001
```

Scenario 2

To remove the Persistent Reserve (PR) from a different host (Host B) than the one who owns it (Host A) (for example, after the owning Host A crashed), see the command in Figure 8-33.

```
# scu -f /dev/rdisk/dsk45c pres clear key 0x10001
scu: 'persistent reserve out' failed on device '2105800' at nexus [2/3/2]
     SCSI Status = SCSI_STAT_RESERVATION_CONFLICT (0x18) - Target reservation conflict
```

Figure 8-33   Trying to remove the PR from another system in the key list

The clear PR failed, because Host A used the key 0x10001; we are trying to remove the PR from Host B, which has to use its own key. To find out the valid key of Host B, we need to check the key table, as in Figure 8-34 on page 372. Host A uses the key 0x10001, so the key for Host B on this volume is 0x10002.

```
# scu -f /dev/rdisk/dsk45c show keys

Persistent Reservation Header:

                 Generation Value: 24
                Additional Length: 48

Reservation Key List:

                        Key Entry 0: 0x10002
                        Key Entry 1: 0x10002
                        Key Entry 2: 0x10002
                        Key Entry 3: 0x10002
                        Key Entry 4: 0x10001
                        Key Entry 5: 0x10001

#
```

Figure 8-34   Checking the key table

If we try the same command again with the valid key, the PR will be cleared:

```
# scu -f /dev/rdisk/dsk45c pres clear key 0x10002
```

Scenario 3

In this scenario, we want to remove the PR from a different host that has no entry in the key table. If the host has no key entry in the PR table, it can not change the PR state of the volume, as seen in Figure 8-35.

```
# scu -f /dev/rdisk/dsk45c show reservation
scu: Unable to open device '/dev/rdisk/dsk45c', EIO (5) - I/O error
# scu -f /dev/rdisk/dsk45c show keys
scu: Unable to open device '/dev/rdisk/dsk45c', EIO (5) - I/O error
```

Figure 8-35   Trying to remove the PR from a system not in the key list

You must remove the PR using a system already listed in the key table.

9

IBM AS/400 and iSeries

In this chapter, we provide information on the positioning of external storage (ESS) and the Copy Services functions in the iSeries context. The iSeries already has a very powerful and reliable integrated disk and various alternate techniques for providing functions similar to the ESS Copy Services. Therefore, users must decide which combination of internal/external disk and related functions is the best fit for their requirements.

For detailed technical information on these topics, refer to the redbook *IBM @server iSeries in Storage Area Networks: Implementing Fibre Channel Disk and Tape with iSeries*, SG24-6220, which is available at:

http://www.redbooks.ibm.com

9.1 Overview

The iSeries and its predecessors, including the AS/400, are systems that were designed for commercial processing workloads. A key premise of their design was ease-of-use and simplicity of management in order to lower the total cost of ownership. The originators had tremendous foresight in designing an architecture that was decades ahead of its time. As a result, iSeries has been "first in the marketplace" time and again as new innovations have come to the technological world. Key examples include the first commercially available RAID-5 disk and seamless transition to a full 64-bit architecture during the early to mid-1990s, and storage virtualization and autonomic computing that have been inherent in the architecture since inception.

Because of the advanced architecture of the iSeries, external storage is positioned very differently on this platform as compared to others. Understanding these differences is key to selecting the most appropriate disk storage and availability options for an iSeries environment.

This chapter will describe the iSeries Storage architecture, and then explain how an external disk fits within that architecture. It will then explain and contrast the various disaster recovery and high availability options that are available on the iSeries, some of which can run on any type of disk, and some that require the ESS disk specifically. Finally, this chapter will review the key V2 enhancements to ESS Copy Services and explain which ones are a fit on iSeries.

9.2 The iSeries storage architecture

Most platforms have storage administrators who spend their days managing the storage on their systems. Some of the tasks they may do in a typical day include:

► Monitor a series of distributed disks individually to ensure they do not become too full. Add new disks as needed and shuffle applications around to use those new disks. Discard smaller disks and install larger disks where needed.

► Monitor and analyze those same disks, watching for performance *hot spots*, where too many I/Os are being directed at one set of arms, therefore causing a bottleneck. Shuffle applications around to reduce those hot spots.

► Add new technology disks to the system, adjust applications as needed, and shuffle applications so the most performance sensitive applications are on the newest disk.

By comparison, the iSeries and its predecessor, the AS/400, have had system-managed storage since inception. Rather than doing all the above tasks, iSeries customers typically just wait for a system message telling them that they have hit their disk-full threshold. They then order a new disk and attach it to the system. By eliminating the storage management workload, iSeries reduces the total cost of ownership of their system. The functions that make this possible are:

► **Storage Virtualization**: The iSeries has what is called a Technology Independent Machine Interface (TIMI) that hides the details of the underlying hardware from the applications. For a disk, this means that the applications are unaware of the number or type of disks that are installed. iSeries sees the disk as one big pool. As a matter of fact, the iSeries sees the main storage and disk as one big address space that is called *single-level storage*. The applications just read/write from the address space, and the TIMI figures out the details. The underlying disk and main storage (and other hardware) can be completely changed or replaced without impact to the applications.

- **Scatter-Loading**: iSeries disk has always been scatter-loaded, not just across all arms in an array, but across all disks in the pool. The pool is called an Auxiliary Storage Pool (ASP) on the iSeries. This scatter-loading is key to the system-managed storage because:
  - The disks fill evenly, so it is no longer necessary to monitor each drive to know when it is approaching capacity.
  - The I/Os are spread evenly across all disks, therefore avoiding hot spots and performance bottlenecks.
  - The OS/400 can use parallel I/O so that many arms all work together to retrieve or write a given file, therefore giving good performance.
  - A new disk can be added without having to juggle other applications around to distribute free space or I/Os. The customer just hooks up the disk and uses the `Add and Balance` command to merge the disk into the ASP, and migrate its share of capacity and I/O onto the new disk.
- **Integrated Read and Write Cache**: The iSeries has a built-in function called Expert Cache that uses the main storage as a very intelligent read cache. Expert Cache watches how the applications are performing, then allocates more or less main storage to each, depending whether the application is getting a benefit from the read cache (in other words, sequential reads) or not (in other words, random reads). The iSeries also has large amounts of write cache on the adapters that are used for integrated disk.
- **Multi-User system**: iSeries users share applications and databases on the system, rather than keeping a separate copy for each user. This can reduce the overall amount of disk storage required.

These pieces of the iSeries architecture allow the iSeries to offer many of the functions typically associated with external storage, but they offer it for *all* disks attached to the system, including the integrated disk.

9.3 iSeries storage terminology

There are a number of iSeries storage items that are important when understanding the iSeries availability functions or setting up the ESS Copy Services functions. These items will be described here and then referred to later in this chapter.

9.3.1 Auxiliary Storage Pools (ASPs)

On the iSeries, disks are included in Auxiliary Storage Pools (ASPs). Data is scatter-loaded across all disks in a given ASP. The first ASP on the system is called the *System ASP* and includes the operating system. Customers can optionally have additional ASPs called *User ASPs*. These are typically used to hold journal receivers and savefiles, since these objects need to be separated from the objects they are protecting. The System ASP and all the user ASPs together are called *traditional ASPs*. There can be 32 of them altogether.

Starting at OS/400 V5R1, a new kind of ASP called an *independent ASP* (IASP) became available. Whereas a user ASP has its definition in the system ASP, an independent ASP holds its definition within the IASP itself. This means that an independent ASP can be varied off and switched to another system in the same cluster. Customers typically use independent ASPs to store critical applications so they can be switched to run on an alternate CPU during a planned or unplanned CPU outage. However, independent ASPs have proven to be very useful for making FlashCopy more useful on iSeries, since they allow copies of the data to be made and moved around without worrying about the load source unit (see 9.3.2, "The Load

Source Unit" on page 376). Up to 223 IASPs are supported on each system, for a total of 255 ASPs, including both traditional ASPs and IASPs.

To take advantage of IASPs, applications must first be migrated from regular ASPs to IASPs. Suitable project planning will be required. A key step is to ensure that all required object types are supported in an IASP, since IASPs only support IFS objects at OS/400 V5R1, and a few object types are still missing at OS/400 V5R2 (spoolfiles, for example). Application coding changes may or may not be required. For applications where good programming procedures have been used, changes are typically minimal.

Once IASPs are implemented, the FlashCopy task can be applied to just the IASP, not the entire system. The IASP needs to be varied off during this process. The duplicate IASP can then be varied on to a different system in the same cluster, after the cluster has been temporarily broken. This greatly simplifies the process for accessing a FlashCopy image of the system. Details are shown in 9.5.4, "FlashCopy with independent ASPs" on page 383. Note that this should only be done via the iSeries Copy Services for ESS Toolkit, which provides automation and protection against varying-on both copies of the ASP simultaneously in the same cluster, which can cause serious problems.

For more information on independent ASPs, please see the following Redbooks:

► *Clustering and IASPs for Higher Availability on the IBM @server iSeries Server,* SG24-5194

► *IBM @server iSeries Independent ASPs: A Guide to Moving Applications to IASPs,* SG24-6802

The following Web site may also be of interest:

http://www-1.ibm.com/servers/eserver/iseries/ha/

9.3.2 The Load Source Unit

Each iSeries system or LPAR partition has a special disk called the *Load Source Unit*. It holds the system's microcode and is used to IPL the system.

The iSeries lab demands that this unit reside on an integrated disk inside the iSeries system unit. This presents a challenge in an ESS Copy Services environment with traditional ASPs, since this disk is part of the single-level storage for the System ASP, and as such needs to be replicated in order to rebuild an image of the system via FlashCopy or PPRC. Therefore, a copy of this disk must also be inside the ESS.

The IBM supported method for handling this is to place the primary copy of the load source on integrated disk as required, but then to mirror it into the ESS so it can be replicated. This is done via the OS/400 remote load source mirroring function and requires a series of configuration and one-time setup tasks as follows:

► The system must have at least two buses. The load source unit must reside on one bus, and the iSeries fibre card that will contain the mirror of the load source must reside on a separate bus. The bus with the fibre card must be marked as a *remote* bus by renaming it to begin with the letter *R* using the Hardware Service Manager.

► The disk that will receive the load source mirror must be the same size as the internal load source unit and must be created as an *unprotected* LUN. This means that even though the LUN is actually protected with RAID 5 or RAID 10 inside the ESS, it reports to the iSeries as being unprotected, therefore allowing the iSeries to mirror the internal load source to it.

Note: The choice of *protected* or *unprotected* is made at the time that a LUN is defined on the ESS, and cannot be changed without recreating the LUN.

Note that the system can only be IPLd from the internal copy of the load source. Should a failure ever be encountered on this drive, the system will continue running and will automatically direct all I/Os to the mirrored copy of the drive on the external disk. However, it is important that the system *not* be powered off until the internal load source has been replaced and resynchronized.

In addition, the system can only perform a main storage dump to the internal load source drive. If the system were to crash while running off the external load source only, then it would not be able to make a main storage dump, and therefore would not be able to use that dump to reduce the recovery time or to diagnose the problem.

9.4 Positioning internal and external disk on iSeries

Because of the iSeries architecture, the positioning of integrated versus external disk on the iSeries is very different from other platforms. The following chart in Figure 9-1 helps to describe the differences in positioning:

<div style="border:1px solid">

ESS Benefits for iSeries Customers

- ESS benefits of interest to iSeries customers:
  - ► Redeploy disk among systems/platforms
  - ► Move disk among leased systems

- ESS benefits that iSeries customers already have
  - ► Data spreading across disk units
  - ► Large read/write cache

- ESS benefits of possible interest to iSeries customers:
  - ► FlashCopy (understand limitations)
  - ► iSeries Copy Services for ESS Toolkit
  - ► PPRC (though usually does not provide the availability level the customer is looking for)

- ESS benefits that apply to other platforms, but not so much on iSeries
  - ► Performance
  - ► Cost
  - ► Ease of Management

</div>

Figure 9-1   Integrated versus external disk positioning on iSeries

On other platforms, consolidated storage offers a huge benefit. Many of these platforms have typically had only limited storage management in the past, so the ESS functions provide huge operational benefits in terms of ease of management and disk allocation. In addition, these platforms have typically not had data spreading across disks or access to large read/write caches before, so they get a tremendous performance boost when they migrate to ESS. This allows them to move to much larger physical DDMs (disk drive module), therefore offering very economical disk. In addition, ESS Copy Services provide very interesting functions to these platforms that were typically not available via other techniques, which once again makes external disk very attractive.

By comparison, iSeries has had many of these functions inherent in the architecture since inception. Storage has always been system-managed, so an external disk sometimes actually increases the management workload due to the code updates that are required on SAN and external disk components. Read/write cache plus data spreading have always been available to provide impressive performance on iSeries, so moving to external disk typically does not provide any performance benefit, and therefore it is usually not possible to move to

the very large, economical disks. This affects the overall price of the solution. The save-while-active function and High Availability Business Partner software have offered attractive alternatives for high availability on iSeries, therefore making the ESS Copy Services functions less interesting. That being said, external storage does introduce two new functions on the iSeries that will be of interest in some environments. First, the ability to move storage between systems and platforms can protect your investment and be tremendously beneficial operationally in environments where disk requirements change frequently. Second, in accounts where the CPUs are leased, external storage makes it possible to move disk between CPUs while staying within the terms of the leasing agreement.

iSeries customers should understand the difference in positioning of external storage between iSeries and other platforms, so they can make informed decisions regarding the costs and benefits of external disk versus integrated disk on the iSeries platform.

9.5 Disaster recovery and high availability options on iSeries

The iSeries has a wide variety of options to help customers increase the availability of their system. Each of these options provides one or more of the following functions:

► Reduce the downtime required to run backups
► Provide a secondary system in case of system failure or site loss
► Provide a secondary system to use during planned outages

A list of the disaster recovery and high availability options available is:

► OS/400 Save-While-Active (SWA)

► OS/400 Switched Disk Clusters

► FlashCopy with Traditional ASPs

► FlashCopy with Independent ASPs

► PPRC with Traditional ASPs

► PPRC with Independent ASPs (see note)

► PPRC-XD (Extended Distance)

► Asynchronous Cascading PPRC

► Remote OS/400 Mirroring

► iSeries High Availability Business Partner (HABP) Replication Software

Note that the PPRC with independent ASPs function listed above is not currently available. However, it is discussed in order to show how it might compare to standard PPRC with traditional ASPs and the iSeries High Availability Business Partner Software.

The above options vary in both cost and function provided. This section will describe each function, and explain the advantages and disadvantages compared with the other options available. Customers can then choose the alternatives that are the best fit for their environments and their budgets. All of the options can run on an external disk if desired. For the FlashCopy and PPRC options, an external disk is mandatory. The options are not mutually exclusive: in many environments, a combination of these options will provide the best solution.

9.5.1 OS/400 save-while-active

Save-while-active (SWA) is an OS/400 function that allows backups to run while users are using the system, therefore increasing availability. It is included with the operating system at no extra charge. It can run on any type of disk, integrated or external.

Most customers use SWA by quiescing the system until they get a checkpoint, then letting the users back onto the system while the backup runs. Quiesce periods typically take 15-30 minutes. This is sometimes referred to as a *quiesced SWA* solution.

For applications that have commitment control, it is also possible to write code so the users can continue to use the system during the checkpoint period. This code will save the journal entries alongside the save, and apply those entries during the restore to ensure the restored image is at a known point. This is sometimes called an *online SWA* solution. An example of an application that allows the users to continue to work during the checkpoint period is the BRMS (Backup Recovery and Media Services)/Domino online backup function. Note that there are few other packaged applications that offer this support at the present time.

In either case, once the checkpoint has been reached, the backup starts. If users change data on the system, OS/400 makes a side-copy of the page of data prior to writing the change to the file. The save then uses the side-pages as necessary to put a time-zero copy of the data on the tape. Customers typically do not need to buy additional disks to accommodate the side pages.

Figure 9-2 explains how the save-while-active function works.
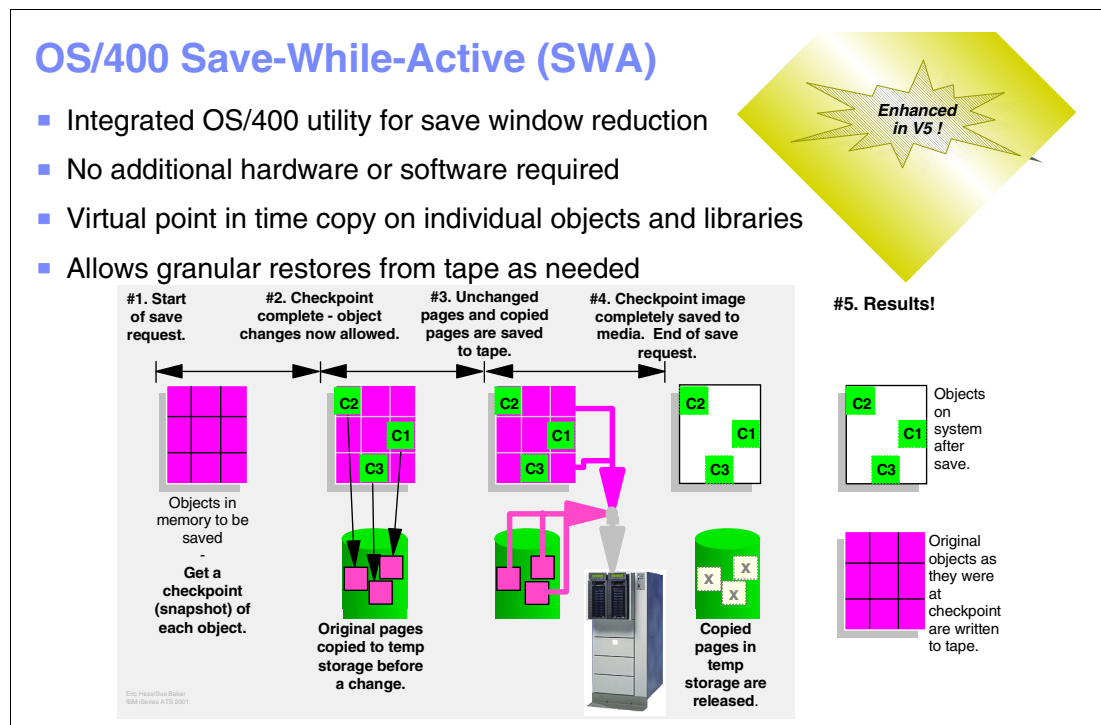
Figure 9-2   OS/400 save-while-active function

Save-while-active is an excellent choice for customers who need a near-online backup, but can afford the short quiesce period required by the typical customer. The reason it is attractive is that it offers good function and can be implemented at essentially no additional cost.

9.5.2 OS/400 switched disk clusters

OS/400 has a function called Switched Disk Clusters whereby an application or set of data can be moved into an independent ASP. That IASP can then be switched between systems in the same iSeries cluster.

Either integrated or external disk can be used for this function. Figure 9-3 shows how the Switched Disk Cluster solution works.
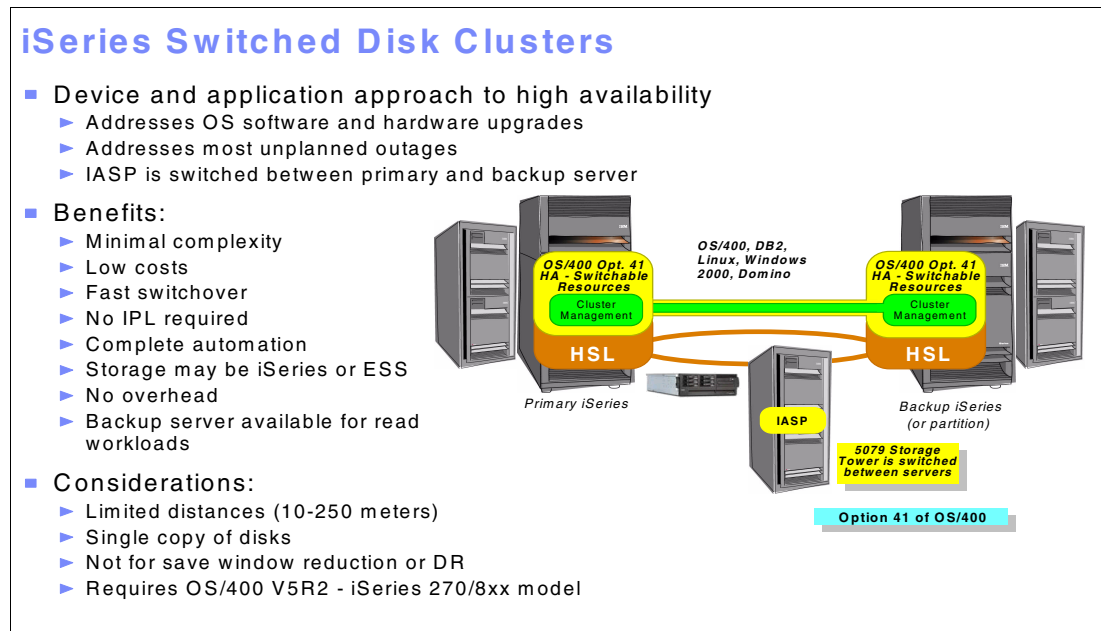
Figure 9-3   iSeries switched disk clusters

This function can be used to move a workload to a different system when hardware or software maintenance is required on the primary system. It can also be used if there is a hardware or software failure on the primary system on a component other than the IASP. The secondary system can continue running its usual workload even after it receives the IASP.

The quick switchover time provided by switched disk clusters is an advantage. Drawbacks include the limited distance (limited by the HSL loop), and the fact that there is only 1 copy of the IASP, so it does not help to run offline backups or to cover for a failure in the IASP itself.

This function would be attractive to customers who were interested in having some coverage in case of failure, without incurring too much cost or adding too much complexity to their environments.

9.5.3 FlashCopy with traditional ASPs

FlashCopy is an ESS Copy Services function that lets the user make a near-instantaneous image of the system. This image can then be used for offline backup, testing, queries, and so on. See Chapter 3, "FlashCopy" on page 33 for more details on FlashCopy.

On some platforms, individual disk units can be copied with FlashCopy one at a time. Disks can sometimes be re-attached to the same system they were flashed from, or attached to another system with minimal impact to that system. This is not true for the iSeries when using traditional ASPs. Due to the unique architecture of this platform, very specific steps must be taken in order to use FlashCopy.

Figure 9-4 shows how FlashCopy can be used on an iSeries that has traditional user ASPs only.

FlashCopy with Traditional ASPs

- Local, instant point-in-time copy of all iSeries disks
  - ► Used for offline backups, test system creation, duplicate systems for running queries, and so on
- Benefits:
  - ► Fast checkpoint.
  - ► Allows Full system backups, including *SAVSYS, with minimal downtime on primary system.
- Considerations:
  - ► Must quiesce (maybe power off) primary server.
  - ► 1/2-1 hour outage on primary server typically.
  - ► Manual processes - no CLI, LS recovery.
  - ► May impact disk performance during copy.
  - ► Secondary system is obliterated when image is attached.
  - ► Full library saves only (not incrementals).
  - ► Tape management system is not updated.
  - ► Typically does not offer any availability advantage over OS/400 Save-While-Active.

Figure 9-4 FlashCopy with traditional ASPs

In order to use FlashCopy, customers must have all their disks residing in the ESS, with the exception of the load source unit, which must remain on internal disk. The load source unit must be mirrored into the ESS using the OS/400 remote load source mirroring function.

Prior to initiating the FlashCopy, the system must be quiesced in order to move all data from main memory out onto the ESS so it will be copied. This usually requires that applications be ended and the system be powered off, although some customers may decide to end the applications and the subsystems only. In this case, the subsequent IPL of the secondary system will be abnormal and may take an extended period of time.

Once the system is quiesced, the FlashCopy command can be initiated. Note that *all* disks must be flashed at once, since the iSeries data is scatter-loaded across all the disks. Note also that the iSeries does not have a Copy Services Command Line Interface (CLI), so the FlashCopy will need to be initiated from the GUI, or from the CLI on another platform. Once the FlashCopy has been established, the primary system can be restarted and users allowed back on. All in all, the outage on the primary system is typically 1/2-1 hour, approximately the same as save-while-active.

Now the FlashCopy disks must be attached to a secondary system. Note that this procedure will obliterate the secondary system and create a clone of the primary system. To do this, a copy of the load source unit must first be pulled onto an internal disk in the secondary system. This is done using a Remote Load Source Recovery procedure. The operator will work at the console and IPL his system from his Licensed Internal Code CD. He will do a manual IPL, and select the **Recover Remote Load Source** options on the Dedicated Services Tools (DST) panel. These steps typically take about one hour.

Note that many customers choose to use a partition in a logically partitioned system for the secondary system. This way, rather than having an entire system sitting idle, awaiting the FlashCopy connection, instead the resources from this LPAR can be moved to other LPARs to do productive work, and moved back when a FlashCopy is required.

Once the duplicate set of disks have been attached to the secondary system, that system will be a clone of the original system, right down to the system name. Beware of adding this system into the network until its name has been changed.

For customers interested in using FlashCopy for offline backup, there are several considerations:

► Only full object saves are possible, not incremental saves. When OS/400 runs a save, it updates the object headers to indicate the last-saved date for each object. When an incremental save is issued, OS/400 compares the last-saved date and the last-changed date in the object header to decide whether the object should be saved. In the FlashCopy save scenario, OS/400 updates the object headers on the secondary system, not the primary system. When the save is complete, this set of disks is discarded. The next day, the original disks are flashed again, but since the object headers do not know about the previous day's save, OS/400 would not determine the proper changed objects. The result would be a cumulative save back to the date of the last save on the primary system, rather than the incremental save intended.

► Tape Management Systems such as BRMS do not get updated. When a tape management system runs a save, it keeps records in its database about all the objects that are being saved. This database is then used to create system recovery reports, media movement reports, and so on. However, in the FlashCopy save scenario, these updates are stored on the secondary disk that is later discarded. The primary system is unaware of the saves. It would be enticing to use the BRMS Networking functions to transmit the updates to the primary system, but this would be difficult operationally; since the system would need to be given a unique name before attaching to the network, and care would need to be taken since this system would be entering/leaving the network daily with long periods of inactivity. This would create a risk that the BRMS databases on peer systems could get out-of-synchronization.

► For media libraries that are aware of the system name, care would need to be taken that the FlashCopy primary and secondary systems do not both try to attach using the same system name.

Overall, the FlashCopy solution with traditional ASPs will typically be of interest to customers who want to create clone systems for application testing, since it is much faster than re-loading from tape. The offline backup function is typically less interesting, since the function it provides is very similar to that of save-while-active in terms of downtime on the primary system; yet it is much more labor-intensive to use due to the remote load source recovery, and more costly to implement since two systems, two sets of disk, and the FlashCopy software are all required. In addition, the FlashCopy offline backup solution has limitations related to incremental saves, tape management systems, and connection to some media libraries, limitations that do not exist with the save-while-active solution. The FlashCopy offline backup solution may be of interest to customers whose applications are not conducive to save-while-active due to object locking issues, and customers who need to do a *SAVSYS save while the users are on the system.

Note that the FlashCopy solution with independent ASPs bypasses many of these limitations and makes this solution far more attractive. It is described in the next section, 9.5.4, "FlashCopy with independent ASPs" on page 383.

9.5.4 FlashCopy with independent ASPs

The FlashCopy solution described in 9.5.3, "FlashCopy with traditional ASPs" on page 380 has several considerations that limit its attractiveness for a customer environment. These challenges arise because the entire disk space, including the operating system, needs to be flashed, therefore obliterating the secondary system when the copied disks are attached, creating a system with a duplicate name, and requiring a remote load source recovery. Fortunately, the independent ASP (IASP) function introduced at OS/400 V5R1 gives us a solution that bypasses these challenges and is quite attractive to an iSeries customer.

IBM has created a toolkit and services offering called *iSeries Copy Services for ESS* that automates the process for using FlashCopy with independent ASPs, and protects data integrity. It is strongly recommended that this procedure *ONLY* be done with the aid of this toolkit. Future releases of OS/400 may provide additional function so the toolkit is no longer required. The toolkit and this process are supported from OS/400 V5R2 onwards. Note that the HA Switchable Resource option of OS/400 (option #41) is also required for this solution and is chargeable.

Figure 9-5 shows the general setup for FlashCopy with independent ASPs.
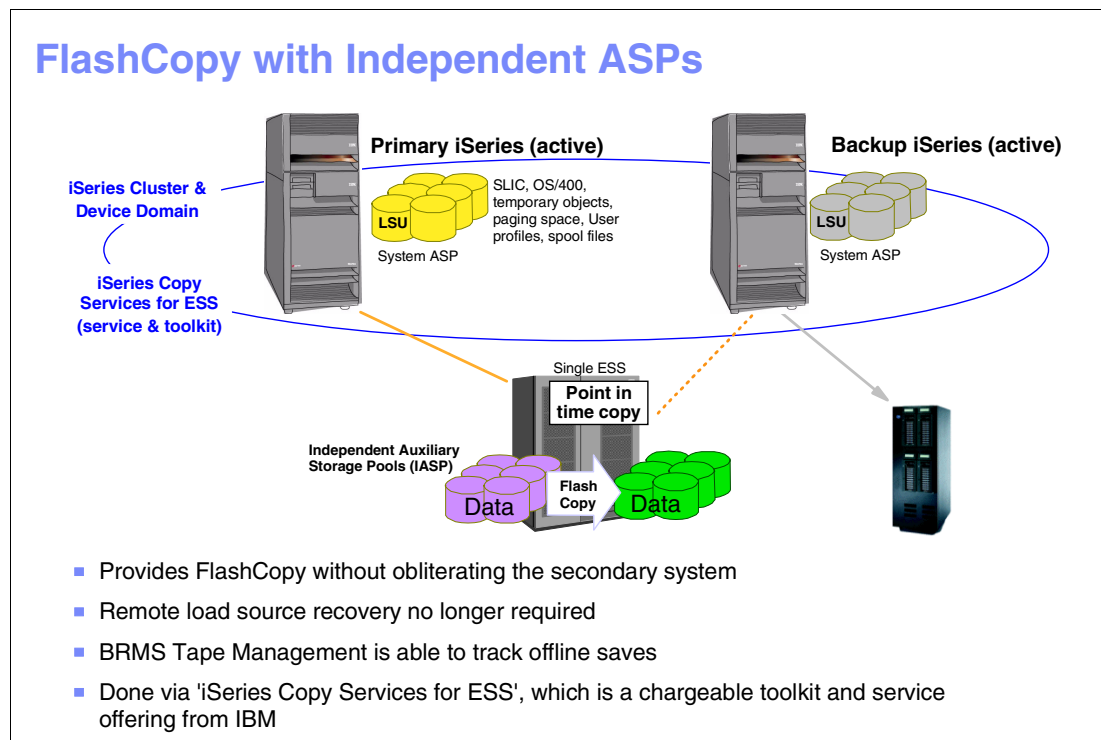
Figure 9-5   General setup for FlashCopy with independent ASPs

Both the primary and secondary systems have a system ASP that is not affected during the FlashCopy process. The primary system then has one or more IASPs that contain various applications. Rather than flashing the entire system as in the traditional ASP example, here only the independent ASP is flashed. The duplicate ASP is then attached to the secondary system and varied-on, while the system ASP on that system remains intact. The benefits of the IASP process over the traditional ASP scenario are:

► The secondary system is no longer obliterated by the process.

► Only the application running in the IASP needs to be quiesced, not the entire system.

- The remote load source recovery procedure is not required. The IASP simply needs to be varied on.

- Only the data in the IASP is copied, not the system ASP, therefore reducing the amount of disk required.

- Because the secondary system keeps its system ASP, it also keeps its name rather than becoming a clone of the primary system, so it can participate in the network.

- Because the secondary system keeps its system ASP, the BRMS Tape Management system, which runs in the system ASP, is available to run backups and to transmit information about the saves to the other systems in the BRMS network.

The FlashCopy solution with independent ASPs makes FlashCopy attractive to iSeries customers for the following reasons:

- For offline backups, it provides similar function to save-while-active due to similar downtime on the production server, and the ability to use tape management systems like BRMS. Manual load-source recovery steps are not required. It is particularly interesting for customers who are not able to use save-while-active due to object locking and application considerations. Drawbacks compared with save-while-active are that it requires extra disk and some chargeable software, and it can only do full saves, not incrementals. This should not be a big concern, given that the save is not impacting the production system.

- Customers can also make images of an IASP and vary them on another system for running queries or testing applications.

For more details on the iSeries Copy Services for ESS offering, please see the following Web site:

http://www-1.ibm.com/servers/eserver/iseries/service/itc/pdf/Copy-Services-ESS.pdf

9.5.5 PPRC with traditional ASPs

PPRC is an ESS Copy Services function that creates an ongoing hardware-level copy of the disk on a separate ESS, typically in a remote location. This function is typically used for disaster recovery at a remote site following a failure at the primary site. See Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53 for more details on PPRC.

Distances are limited to 103 km for performance reasons. For data integrity purposes, a synchronous link is used whereby the primary ESS waits for acknowledgement that the write has been successfully received on the remote ESS, before reporting the write as complete to the operating system. As a result, performance decreases as the distance increases. Performance analysis should be done as part of a PPRC install to understand the response time that can be expected at various distances and with various bandwidth networks.

PPRC is typically considered to be a Disaster Recovery (DR) versus High Availability (HA) solution since the data in memory is typically lost during a failure, and because the transition to the secondary system typically takes more than a few minutes. On iSeries, this is particularly true since a remote load source recovery procedure and abnormal IPL are required on the secondary system, therefore requiring an extended period to cut over to the secondary system. Also, this solution is only helpful for unplanned outages, not planned outages.

Figure 9-6 on page 385 shows how PPRC would work in an iSeries environment. It is very similar to the FlashCopy scenario with traditional ASPs, except that two ESSs are used, and the restart time is much longer since an abnormal IPL is required.

Peer-to-Peer Remote Copy (PPRC)

- Continuous, synchronous copy to a second ESS

- Benefits:
  - ► Simple campus/city disaster recovery (up to 103 km).
  - ► Multiple platforms supported.

- Considerations:
  - ► Does not protect transactions or data in server memory.
  - ► Unplanned failover can be many hours.
    - – IPL required, plus any DB recovery.
  - ► New complexities.
    - – Manual processes - no CLI.
  - ► May impact disk performance.
  - ► High communication costs.
  - ► Second ESS disk copy and second iSeries are unavailable for other uses.
  - ► Not viable for masking outages from software and hardware upgrades.
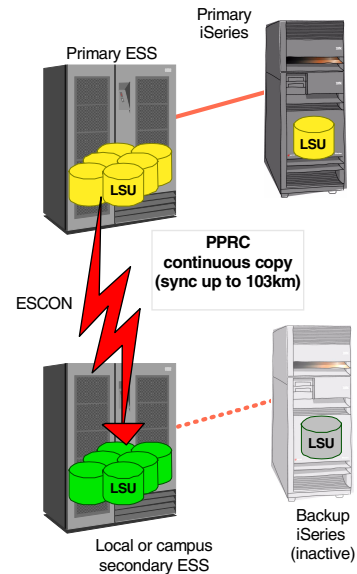
Figure 9-6   Peer-to-Peer Remote Copy (PPRC) with traditional ASPs

In order to use PPRC on iSeries, all data must reside in the ESS, with the exception of the load source unit, which must remain on internal disk. The load source unit must be mirrored into the ESS using the OS/400 remote load source mirroring function. All disks must be PPRCd in order to have a complete copy of the data to cut over to, since the iSeries data is scatter-loaded over all disks.

If a planned PPRC cutover is required, then the primary system should be quiesced to flush data from memory prior to the cutover to avoid an abnormal IPL on the secondary system. However, in DR scenarios, the failure is unexpected, so this option is not available.

When it is time to connect to the duplicate copies of the disk, a secondary system is required. Any workload that is running on this system will be obliterated when the PPRCd disks are attached, so an idle system is typically used. As in the FlashCopy scenario, customers often choose to use a partition in a logically partitioned system, so they can move the resources elsewhere except when they are needed for a PPRC cutover.

As in the FlashCopy scenario, a remote load source recovery procedure is needed to copy the PPRCd load source from the ESS down onto the internal disk in the secondary system. The operator does this by IPLing to DST using his LIC CD, then choosing the remote load source recovery procedure. However, unlike the FlashCopy scenario, this procedure will take many hours since the IPL will be abnormal. This is because the system was active when the failure occurred and the data in memory was lost. During the abnormal IPL, OS/400 will check for damaged objects, rebuild access paths, and generally try to recover from the failure. This IPL will be similar in duration to the IPL that would occur after a power failure or a hardware failure that brought the system down. It is typically in the 2-12 hour range.

Once the disks are attached to the secondary system and the system has IPLd, users can be transferred over. Once the problem on the primary system is corrected, PPRC can be reversed to migrate the data back, and the corresponding load source recovery can be done

on the primary system to prepare it for production use. On the cut-back, the source iSeries can be quiesced to avoid an abnormal IPL as part of the disk pick-up procedure.

The PPRC solution is attractive to customers who are looking to get some sort of DR capability in their environment. This solution provides a better function than a restore from tape, since the cutover will typically be quicker, and the data will be more current than the previous nights' backup. However, this solution is not as function-rich as that provided by the iSeries HA software. The differences are discussed in 9.5.10, "High Availability replication software" on page 391.

9.5.6 PPRC with independent ASPs

Theoretically, a solution similar to FlashCopy with independent ASPs could be constructed for use with PPRC and independent ASPs. It would offer the following benefits over the traditional PPRC solution:

► Cutover to the secondary system would be quicker since no load source recovery would be required, and since the vary-on of the independent ASP is typically faster than an abnormal IPL.

► A smaller communications pipe and less disk would be required, since only selected applications would need to be replicated, in other words, those in the IASP. Note that arrangements would need to be made to keep user profiles and other OS/400 objects synchronized between systems ready for a cutover.

► The secondary system would not be obliterated by the cutover.

► This solution could be used for planned outages like release upgrades and PTF applies, since the system ASP on the primary system would not be altered as the IASP was moved around.

However, as of July 2003, there is not currently a toolkit to assist with such a solution. As a result, independent ASPs should be handled the same as traditional ASPs when using PPRC, that is, by replicating the entire system and doing a remote load source recovery to cut over to the secondary system.

9.5.7 PPRC-XD

PPRC-XD is an ESS Copy Services function that provides an ongoing hardware-level copy of the disk on a separate ESS, typically in a remote location. The difference between PPRC and PPRC-XD is that PPRC is synchronous, whereas PPRC-XD is asynchronous. This means that PPRC-XD can run across dramatically larger distances without impacting performance on the primary system. However, the PPRC-XD copy of the data is *fuzzy* and is only usable after it has been made *synchronous*. This is typically done by switching into synchronous mode, then making a FlashCopy of the synchronized data, then returning to asynchronous mode. See 4.3, "PPRC extended distance (PPRC-XD)" on page 59 for further information on PPRC-XD.

On some platforms, switching into synchronous mode is fairly simple, and can be done every 15 minutes or so. Commands are used to force data out of memory onto disk, and then a brief period of lower performance follows while the ESS drops into synchronous mode. By comparison, on an iSeries, flushing memory requires that the system be powered off, and therefore it is only practical to do this periodically, typically once each evening.

As a result, PPRC-XD is really only of interest on the iSeries in the following three scenarios:

► Long Distance Backups: This is an option for customers who want to do long distance remote backups. At present, fibre tape is only supported up to 10 km on iSeries. If a

customer needed a longer distance, he could use PPRC-XD to send the data to the remote site, then use a local tape drive at that location to run a backup from the FlashCopy. Note that the considerations listed above under 9.5.3, "FlashCopy with traditional ASPs" on page 380 would apply.

► Near-line Backups: This scenario is similar to restoring the system from tape following a failure, in that the result is an image of the system the way it looked the last time it was quiesced, which is typically the previous night. The difference is that the tapes do not need to be shipped to the remote location, and the system can be brought up much more quickly via the FlashCopy procedures (for example, approximately one hour), rather than waiting for a restore from tape. Note that saves to tape should still be performed, in order to restore individual objects and to allow fallback to earlier dates than the previous night should any sort of data corruption be discovered, and so on.

► Data Center Migration: Suppose a customer needed to migrate his data center from one side of the continent to the other. The obvious solution would be to create tape backups on Friday night, fly the tapes across the country (or region) on Saturday, and reload the system at the remote site on Sunday. Clearly, this would cause extended down time on the system. An alternative strategy would be to set up a PPRC-XD connection between the two systems. Gradually over the course of several hours, the remote system would receive all the data from the primary system, all without impacting performance. Then on cutover weekend, all that would be required would be to quiesce the primary system, wait for all the changes to be transmitted, then do a remote load source recovery on the remote system, and transfer the users to the new system. The total outage would be in the neighborhood of one to two hours versus an entire weekend.

Figure 9-7 shows the PPRC-XD solution, with a backup being run from the remote FlashCopy.
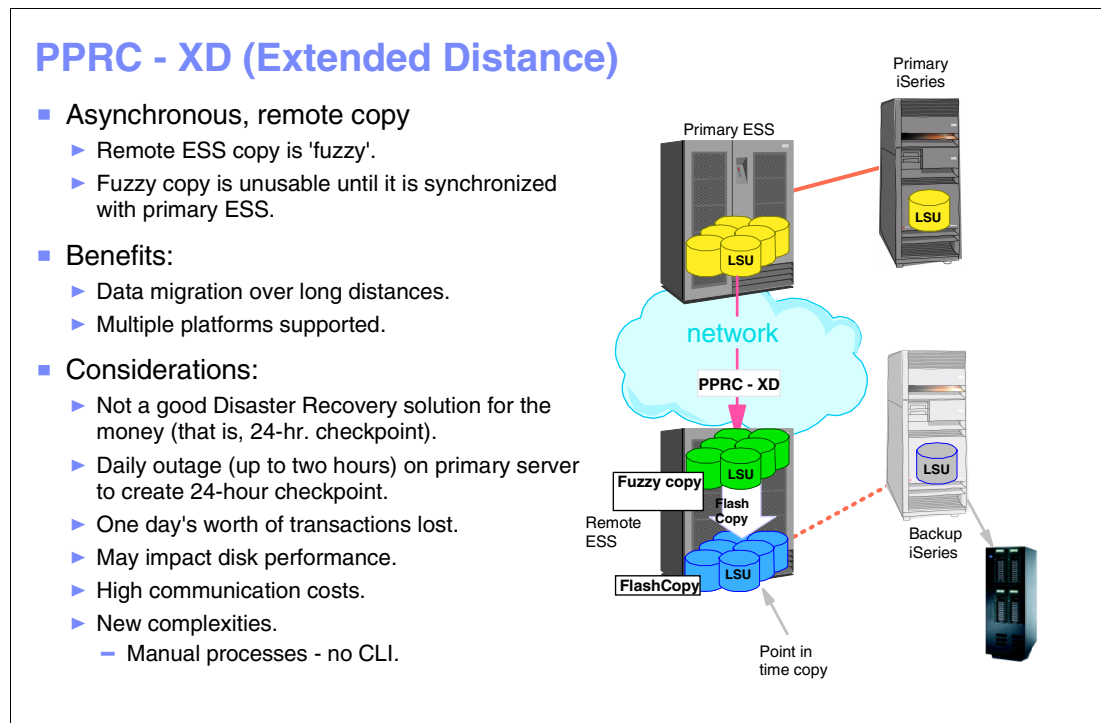
Figure 9-7   Peer-to-Peer Remote Copy Extended Distance (PPRC-XD)

9.5.8 Asynchronous Cascading PPRC

Asynchronous Cascading PPRC allows a LUN to be both a PPRC source and target. The typical use of this is to allow a multi-site PPRC environment, where the first link is synchronous, and the second link is asynchronous. Figure 9-8 shows how this works on platforms other than iSeries.

Asynchronous Cascading PPRC for non-iSeries Platforms

- A Consistent Copy may be made at a Remote Site with minimal impact to Host I/O at a Local Site.
- Existing Freeze Command used to break paths between Local and Intermediate Site and cause the Local Pairs to Suspend.
- When the Intermediate Site has sent all updates to the Remote Site, a FlashCopy is done to create the Backup Copy.
- After FlashCopy completes, the local pairs are resynchronized.
- May be performed at regular intervals.

Figure 9-8   Cascading PPRC for non-iSeries platforms

The intent of cascading PPRC is to allow a long distance remote copy solution with:

► Comparable performance to a short distance PPRC solution.

► Less impact to the host when a consistent copy is made at the remote site, therefore allowing more frequent consistent copies to allow a more up-to-date recovery point.

On most platforms, this is accomplished as follows:

► The synchronous PPRC link is across a short distance, so performance is not a problem for the host.

► When a consistency copy is required, an ESS **Freeze** command is used at the local site, rather than quiescing the applications or going to PPRC synchronous mode, as would have been done in traditional PPRC-XD. The Freeze typically takes three to five seconds, and then the users can continue using their application. This is much less intrusive to the users than the application quiesce and/or PPRC synchronize options. Meanwhile, PPRC is suspended between the local and intermediate site, and the PPRC changes queue up at the local site as the users use their application. Simultaneously, the PPRC-XD link between the intermediate and remote sites catches up, and then the FlashCopy is created at the remote site. PPRC is then resumed between the local and intermediate sites. Although the transactions that are in memory at the local site are not reflected in the consistency copy, the assumption is that database logging and recovery actions will make the consistency copy usable following a failure.

► If the primary system fails, the intermediate site is available for cutover. Alternatively, the intermediate site can go into synchronous mode with the remote site to transmit all the changes, and then the remote site can be used for cutover. If the primary and secondary

sites are both lost, then the most recent consistent copy at the remote site can be used for cutover.

On iSeries, this solution offers some additional functions over PPRC-XD, but not as much as on other platforms:

► Because the iSeries needs to be quiesced/powered off to flush the memory that is ready to make a consistent copy, rather than using the `Freeze` command, this solution does not offer any advantage in making more-frequent consistency copies.

► However, it *does* offer a long distance copy solution without affecting performance on the primary system. As above, if the primary site failed, the intermediate site could be used for a cutover, or could be switched to synchronous mode to finish updating the remote site ready for cutover. If the primary and intermediate sites failed, then recovery could be performed via the most recent consistency copy, which would typically be from the previous evening. Recall that the recoveries require a remote load source migration and an abnormal IPL, so this will take several hours, and the secondary system they attach to cannot be running any other workload.

This solution may be of interest to customers who are already using this technique for their other platforms and want to incorporate their iSeries. However, the iSeries HABP solution described later in this section offers considerably better function, since it can still cover long distances, but only requires two copies of the disk versus three, provides a much quicker switchover, and does not require the secondary system to be idle prior to cutover.

9.5.9 Remote OS/400 mirroring

OS/400 offers a disk mirroring function that is used by iSeries customers for disk protection and improved performance in a write-intensive environment. Mirroring is turned on via the Dedicated Service Tools (DST) menu. OS/400 will review the disk-related hardware available, and create a mirroring configuration that provides the maximum protection possible, that is, tower-level, bus-level, IOP-level (Input-Output Processor-level), or disk-level protection.

With ESS, OS/400 mirroring can be used to create a PPRC-like solution that offers slightly better function at a lower price than PPRC. The solution has distance limitations similar to PPRC for performance reasons. Figure 9-9 on page 390 shows how it works.

Remote OS/400 Disk Mirroring

- iSeries storage mirrored to remote ESS
- Benefits
  - ► Full redundancy - I/O path and disk protection.
  - ► Simple campus disaster recovery.
  - ► System continues running after a disk failure, unlike PPRC.
  - ► No additional software required.
- Considerations
  - ► Does not protect transactions or data.
  - ► Unplanned failover can be many hours, similar to PPRC.
    - – IPL required, plus any DB recovery
  - ► Second ESS disk copy and second iSeries are unavailable for other uses.
  - ► Not helpful for planned outages such as HW or SW upgrades.

Remote ESS - iSeries
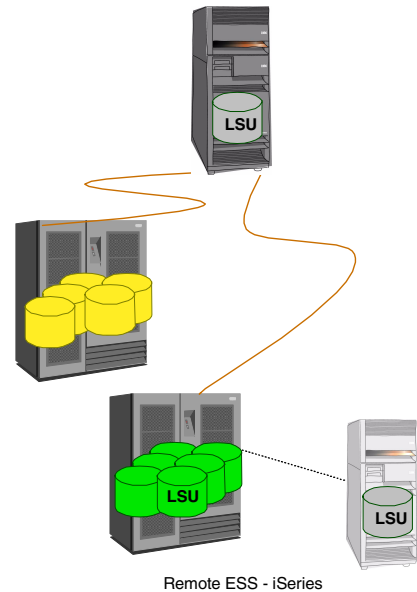
Figure 9-9   Remote OS/400 disk mirroring

In this configuration, OS/400 has half of its disk attached locally and half attached remotely. The local set of disks can be either an integrated disk or ESS. The remote disk is on ESS to allow the extended distance.

It is critical that the hardware be attached in such a way that each mirrored pair has one unit on the local disk and one unit on the remote disk. To do this, a minimum of two buses is required on the primary system. The bus that holds the fibre cards for the remote ESS is marked as a remote bus by renaming it to start with the letter "R". OS/400 will then attempt to make the appropriate pairs, assuming a suitable disk exists for it to do so. For maximum protection, it is also beneficial to arrange the hardware to achieve tower-level protection when possible, just as would be done when mirroring integrated disk.

Notice that in this configuration there are only two copies of the load source: one on integrated disk, and a mirror on the remote ESS. Contrast this with PPRC where the load source on the integrated disk is mirrored into the local Shark, then propagated via PPRC to a third copy on the remote ESS.

Once in operation, this solution is very similar to PPRC in the coverage it provides if the primary system CPU fails. A secondary system would need to be available, and it would be attached to the remote copy of the mirrored disks using a remote load source recovery procedure followed by an abnormal IPL. Any workload previously running on the secondary system would be obliterated in the process, so it is usual to have the secondary system sitting idle. Similar to the PPRC scenario, a clever solution is to use a partition in a logically partitioned system for this, since its resources can be allocated elsewhere when not needed rather than sitting idle.

The key difference between this solution and PPRC is in the scenario where a component in the disk subsystem fails. Examples would include a fibre card in the iSeries or ESS, a fibre cable, a switch, or some disk drives. In the PPRC scenario, the primary system would fail, since it was missing some of its disk, and a cutover to the secondary system would be required in order to accommodate the users. By comparison, in the remote mirroring

scenario, the system would continue running, using the data on the other unit in the mirrored pair.

In addition, this solution provides an equivalent function to *multi-path* in this environment, that is, it provides alternate paths with load balancing and failover.

The other difference between the PPRC and remote mirroring solutions is cost. In each case, a duplicate copy of the disk and a secondary CPU are required. The difference is that the mirroring software is included with OS/400 at no extra charge, whereas PPRC is a chargeable function on the ESS. Customers who already own PPRC should do their own calculations to determine which solution is most cost-effective for them, and contrast that with the function provided.

The OS/400 remote mirroring function would be attractive to a customer who was interested in PPRC, but wanted the increased resilience provided by this solution in terms of multi-path like function, and coverage for a disk/SAN failure without having to do a fail-over procedure.

9.5.10 High Availability replication software

Several IBM @server iSeries Business Partners provide High Availability software packages (HABP software) that replicate data between systems using the OS/400 journaling functions. These packages provide the premier level of function for all three of the key areas of high availability, namely, offline backups, secondary systems for hardware failures or site loss disasters, and coverage for planned outages.

Figure 9-10 shows how the High Availability Business Partner Software works.
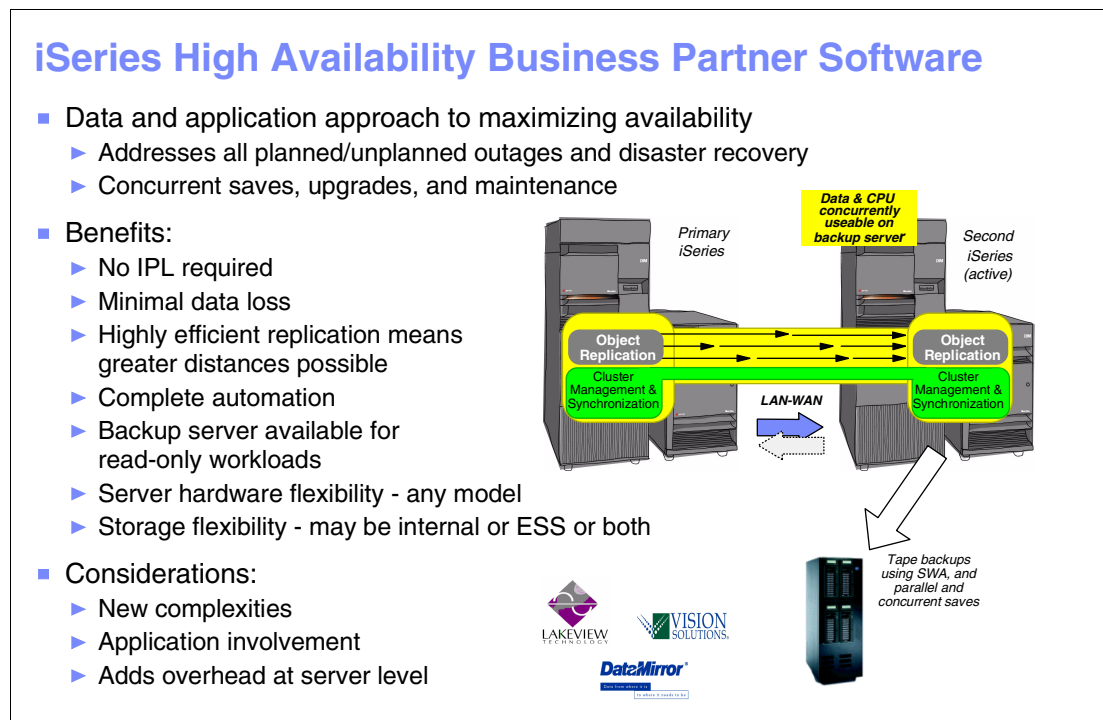
Figure 9-10   iSeries High Availability Business Partner Software

The primary system and a secondary system are connected via a LAN or WAN. Either integrated or external disk can be used. The HABP software is loaded onto the system. The customer decides which applications he wants to replicate, and selects them via the software. Remote journaling is turned on for the files associated with each application. Every time a

transaction occurs on the primary system, OS/400 remote journaling then transmits that change to the secondary system. The HABP software controls the process, and provides the checks and balances to make sure that every transaction gets across safely. It also provides functions to copy required OS/400 objects, such as user profiles, and so on, and to automate cutovers when required. On the secondary system, the HABP software receives the journaled entries and applies them to the database.

Replication is not limited to one direction. Some customers will have three systems backing one another up in a triangular configuration. Other customers will have half their applications replicating one direction, and the other half replicating the other direction. This way, if a site is lost, only half the applications need to be cut over. It is also possible to have one system replicated to multiple other systems, for example, a local replica for attended backups, and a remote replica in case of site loss. Another option is to have multiple images of an application on another system, for example, one replica for offline backups, and one replica for 7x24 read-only queries.

In the case of a planned outage, such as a CPU upgrade, release upgrade, or PTF install, the users can be switched to run on the secondary system so they are not impacted by the system work. Thereafter, the replication process can be reversed, and the users moved back to the primary system once the data is synchronized again. For applications that have been designated as *Cluster Proven*, the switchover can be automated via the HABP software and can be done in minutes. For applications that are not Cluster Proven, some manual checking is required to make sure all transactions have been propagated prior to switching the users over. Customers have been known to do the switchover in 15-30 minutes in this case.

In the case of an unplanned outage, such as a hardware failure or site loss on the primary system, the same principles are used to switch the users to the secondary system until the problem on the primary system is corrected, then switch them back.

The HABP solution also allows customers to run backups without impacting their primary system. The users continue to work without interruption, and the changes continue to be propagated to the secondary system for safekeeping. The only difference is that the apply process is suspended on the secondary system temporarily. Once the backup is complete, the "apply job" can be restarted. Should the primary system fail during the backup period, the backup can be aborted, the apply job restarted, and as soon as the apply job has caught up, the users can be cut over to the secondary system. For customers who want to minimize their cutover time in this scenario, a save-while-active save can be used, with the apply jobs being restarted once the checkpoint is reached.

This HABP solution provides the premier level of availability on the iSeries platform for the following reasons:

► For the offline backup scenario, only this solution and the online save-while-active solution (that few customers use) allow the backup to run without any downtime on the primary system. The quiesced save-while-active solution requires that applications be quiesced briefly. The FlashCopy solution with traditional ASPs requires a power-down on the primary system and the FlashCopy solution with IASPs requires a vary-off of the IASP.

► In a system or site failure situation, this solution provides the quickest cutover to the secondary system, since no IPL is required. The hardware copy solutions require an abnormal IPL.

► Planned outages, such as release upgrades and PTF applications, are covered, in addition to hardware upgrades, system failures, and site losses. PPRC does not address release upgrades and PTF applies at the present time.

► The secondary system is not obliterated on cutover, and therefore can be used for other workloads day-to-day.

► Considerably less data is replicated, since only the applications of interest are mirrored, not the operating system, temporary files, workfiles, and less-critical applications. This means that this solution is viable over longer distances and/or network costs are reduced.

► Because journaling is used, data integrity is at a transaction level; full transactions arrive at the secondary system, and the journals can be reviewed to understand which transactions made it to the secondary system prior to a failure. The hardware copy solutions do not have a concept of transactions.

► With appropriate journaling settings, transactions in memory can be protected, in addition to transactions that have already been paged out to disk. The hardware copy solutions only protect transactions that are on disk.

Depending on the customer environment, the HABP solution is typically in the same price range as the PPRC and FlashCopy solutions, given that it requires a secondary CPU, a secondary copy of disk, and some software. This pricing statement may vary somewhat in the following cases:

► Customers who only need to replicate a subset of their applications, and therefore need less extra disk in the HABP scenario than they do in the PPRC/FlashCopy scenario.

► Customers who already own a PPRC/FlashCopy license for their other platforms, and can add the iSeries without causing a significant increase in their PPRC/FlashCopy licensing fee, compared with the cost of the HABP software.

► Customers who can or cannot take advantage of the processing cycles on the secondary system to minimize the cost of idle hardware (for example, by using a logically partitioned system where the resources can be moved to/from other LPARs, or in HABP, where the secondary system can be used for another workload).

In some customer environments, the HABP software is considered more complex than the PPRC/FlashCopy solutions, so customers may need to make a bigger investment in staff skills and training.

Overall, the HABP software solution is attractive to customers who are looking for the highest levels of availability on the iSeries platform. Customers who do not need this premiere level of function may want to contrast the price and complexity of the various solutions and decide whether one of the other solutions would be a better fit in their environment.

9.6 Requirements for Copy Services on iSeries

In the following sections, we describe hardware and software requirements for AS/400 and iSeries servers using the ESS Copy Services.

9.6.1 Major characteristics

► The ESS Copy Services environment is supported on SCSI and Fibre Channel attached AS/400 and iSeries servers.

► The following Fibre Channel topologies are supported:

– iSeries OS/400 at V5R1 support:

• Point-to-Point

• Arbitrated Loop

– iSeries OS/400 at V5R2 support:

• Point-to-Point

- Arbitrated Loop
- Switched Fabric

▶ ESS Copy Services functions can be invoked and managed via the ESS Copy Services Web User Interface (WUI). The iSeries does not have a Copy Services CLI (Command Line Interface), but iSeries LUNs can be managed using the Copy Services CLI on a different platform.

9.6.2 Hardware requirements

For the latest information about AS/400 and iSeries models supported with ESS Copy Services, see the IBM Enterprise Storage Server Interoperability Matrix. The document can be found at the Web site:

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

9.6.3 Operating system requirements

▶ SCSI Attachment Support from V3R1

▶ FlashCopy and PPRC support from V4R5

▶ Remote Load Source Mirroring support from V4R5

▶ PPRC-XD support from V5R1

▶ Fibre Channel Attachment Support from V5R1

9.7 Implementing Copy Services on iSeries

For detailed information regarding the implementation steps for Copy Services on iSeries, refer to the redbook *IBM @server iSeries in a Storage Area Network: A Guide to Implementing FC Disk and Tape with iSeries*, SG24-6220, which is available at the Web site:

http://www.redbooks.ibm.com

For information regarding IASPs, clustering, and switched disk, refer to the redbook *Clustering and IASPs for Higher Availability on the IBM @server iSeries Server*, SG24-5194.

9.8 ESS Copy Services V2 functions on iSeries

Several new ESS Copy Services functions were introduced in V2 of LIC 2.2. This section will describe how these new functions might fit in an iSeries environment.

9.8.1 Improved FlashCopy establish times

The time to establish a FlashCopy has been reduced by up to 10-fold on FlashCopy V2. However, this is not of much interest to iSeries customers, since the establish time of a FlashCopy (seconds) is insignificant when compared to the time required to quiesce the system to be ready for a FlashCopy, then restart it afterwards (30-60 minutes).

9.8.2 Data Set FlashCopy

This function is only available for zSeries and S390 environments, so is not of interest to iSeries customers.

9.8.3 Multiple-Relationship FlashCopy

This function allows customers to make up to 12 FlashCopy images from a single set of source LUNs. This configuration is valid on the iSeries and may be useful to some customers. Examples of how this might be used are:

► During one quiesce cycle, a customer may want to make one FlashCopy for offline backup, and another FlashCopy for application testing

► A customer may want to keep a FlashCopy image of his system in case of system failure, since it is typically quicker to attach a FlashCopy to a secondary system than to restore from tape. Suppose the customer wanted to keep two images, one from before his batch run, and one from after his batch run, but the background copy for the first image took longer to create than the batch took to run. With V1 FlashCopy, he would have to wait for the first FlashCopy to finish its background copy, then establish the second FlashCopy before letting his users back onto the system. With Multiple Relationship FlashCopy, he can establish the second FlashCopy, let his users back onto his primary system, and let both FlashCopies finish in the background.

9.8.4 Spanning LSS boundaries for FlashCopy

Originally, the two LUNs in a FlashCopy pair had to reside on the same LSS. With FlashCopy V2, the drives can now reside on different LSSs. This provides more flexibility for configuring FlashCopy pairs on a given ESS and will be an advantage on the iSeries as well as other platforms.

9.8.5 Consistency Groups and the Freeze option

Consistency Groups and the Freeze option were created to improve the ability to get a consistent point-in-time copy across multiple volumes on some platforms. A consistency group allows a group of volumes to be associated with one another. The Freeze option causes a LUN to tell the host that its queue is full, and therefore subsequent writes will wait until the consistency group is released. Application writes that are dependent on these writes will also wait, therefore creating an application point of consistency.

Once disks are frozen, a FlashCopy or consistency copy can be made. Although transactions are still in memory and will be missing from the consistency copy, the assumption is that database logging and recovery activities can be used to make the copy usable.

This function is not of interest to iSeries customers. Much stricter rules are required in order to get a usable FlashCopy on iSeries. In a traditional ASP environment, the system must be quiesced/powered off to flush memory, and then *all* LUNs on the system must be flashed simultaneously. In an IASP environment, the application must be quiesced and the IASP varied off to flush memory, and then *all* LUNs in the IASP must be flashed simultaneously. Recall that the ESS Copy Services Toolkit should be used in the IASP environment.

9.8.6 Inband commands over PPRC links

This function is useful for customers who are doing PPRC or PPRC-XD and periodically making a FlashCopy image on the remote ESS as a consistency copy. With FlashCopy V1, the user had to issue the commands on the remote ESS. With FlashCopy V2, the FlashCopy needs to be set up initially on the remote system, but thereafter, it can be managed by issuing FlashCopy commands to the local PPRC LUN. The commands are then transmitted via the remote PPRC LUN to the remote ESS to manage the FlashCopy. This function is of interest to customers who do not have easy network access to their remote ESS, typically due to network security restrictions.

This function will be of interest to iSeries customers in this scenario.

9.8.7 Incremental FlashCopy with Reverse Restore option

Incremental FlashCopy allows a FlashCopy pair to be established, and then changes tracked so a subsequent FlashCopy can be created simply by applying the changes, rather than recopying the entire LUN. This is helpful in environments where only a subset of the data changes, therefore reducing the time to create the new image.

With Incremental FlashCopy, it is acceptable to make changes to both the source and target LUNs. When the incremental updates from a FlashCopy get applied, they overlay the changes on the secondary LUNs with the changes on the primary LUNs.

Incremental FlashCopy also offers a Reverse Restore option whereby the FlashCopy target along with any changes that have been made to it can be restored over the source. This is sometimes called a FlashBack.

Note that the original FlashCopy cannot be Reverse Restored once changes have taken place on it. Only the image with updates can be reverse-restored. If the original image might need to be restored, a Multiple Relationship FlashCopy should be used to keep both this traditional FlashCopy image and also the Incremental Image.

In an iSeries environment, this function could be useful in the following situations:

► Customers who are using FlashCopy and only have a small percentage of their data that is changing will find that an incremental FlashCopy will finish more quickly than a traditional FlashCopy.

► Customers who are doing a release upgrade or significant application changes may want to make a FlashCopy for fallback purposes in case the work goes awry. The Reverse Recovery option could be used to restore the system to its image from prior to the work, which would likely be quicker than accomplishing the same with a full FlashCopy. This would be more effective still with independent ASPs since there would be more granularity in the restore, for example, if there were multiple applications, only the failed IASP would need to be reverse-restored.

► Customers who do backups before and after their batch run could instead make two FlashCopy images prior to the batch run. One image would be kept as the before-image, and the other would be an incremental image. If the batch run failed, the before-image could be reverse-restored. If the batch run was successful, then the incremental changes could be applied to the second image and the FlashCopy pairs withdrawn before the system was restarted. This after-image could then be kept until the following evening so it could be used to recover the system if there was a problem during the day

Note that all the usual FlashCopy conventions for iSeries would need to be followed. For example, the source system would need to be quiesced to flush the memory prior to applying the incremental changes to the FlashCopy.

9.8.8 Asynchronous Cascading PPRC

The Asynchronous Cascading PPRC function was introduced at V2. It is described in detail in 9.5.8, "Asynchronous Cascading PPRC" on page 388, along with a comparison to other platforms and other iSeries DR/HA options.

9.8.9 Cache Destage performance enhancements

At the same time as the Copy Services V2 announcement, a significant change to the NVS Cache Destage algorithm was also announced. This change was written specifically for the iSeries and can provide a dramatic improvement on general disk performance in I/O intensive environments. It is included in the new ESS LIC at no additional charge.

9.9 Conclusion

The iSeries architecture was designed for ease of management, and as such, is very different from the architectures on other platforms. As a result, external storage is a very different fit on this platform compared to others.

This chapter has attempted to describe that architecture as it relates to external storage, then compare and contrast the various disaster recovery and high availability options available on both integrated and external disk. Finally, the V2 Copy Services are described as they relate to iSeries.

iSeries customers should understand the above information so they can make informed decisions regarding the best disk and availability strategies for their environments.

10

Disaster recovery

In this chapter, we discuss how to exploit ESS Copy Services for disaster recovery (DR). We analyze requirements and describe several outage scenarios.

The chapter is organized as follows:

► Disaster recovery concepts
► ESS Copy Services in DR
► Types of DR procedures
► Loss of Copy Services servers
► Planned outages
► Unplanned outages
► Guidelines for DR

This chapter helps you to apply PPRC concepts and user interfaces for disaster recovery planning.

► PPRC concepts were introduced in Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53.

► User interfaces were discussed in Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169 and Chapter 7, "ESS Command Line Interface" on page 301. The other user interface for open systems, the ESS API, is discussed in Appendix A, "ESS Application Programming Interface (API)" on page 471.

In Chapter 12, "IBM Systems Group Service Offerings" on page 449, we present an integrated solution for the Windows 2000 platform.

10.1 Disaster recovery concepts

Disaster recovery planning is part of business continuity planning. For business continuity, components like the people side of your business processes, business resumption, crisis management, and contingency plans must be prepared too. Here we introduce only some basic concepts. For a more comprehensive discussion, see *IBM TotalStorage Solutions for Disaster Recovery,* SG24-6547.

10.1.1 What is disaster recovery?

How would a shutdown of your IT system affect your business? What about a site disaster? Are your business critical processes and data protected from a site disaster? Do you put off system maintenance and upgrades to avoid system downtime? Consider Figure 10-1.
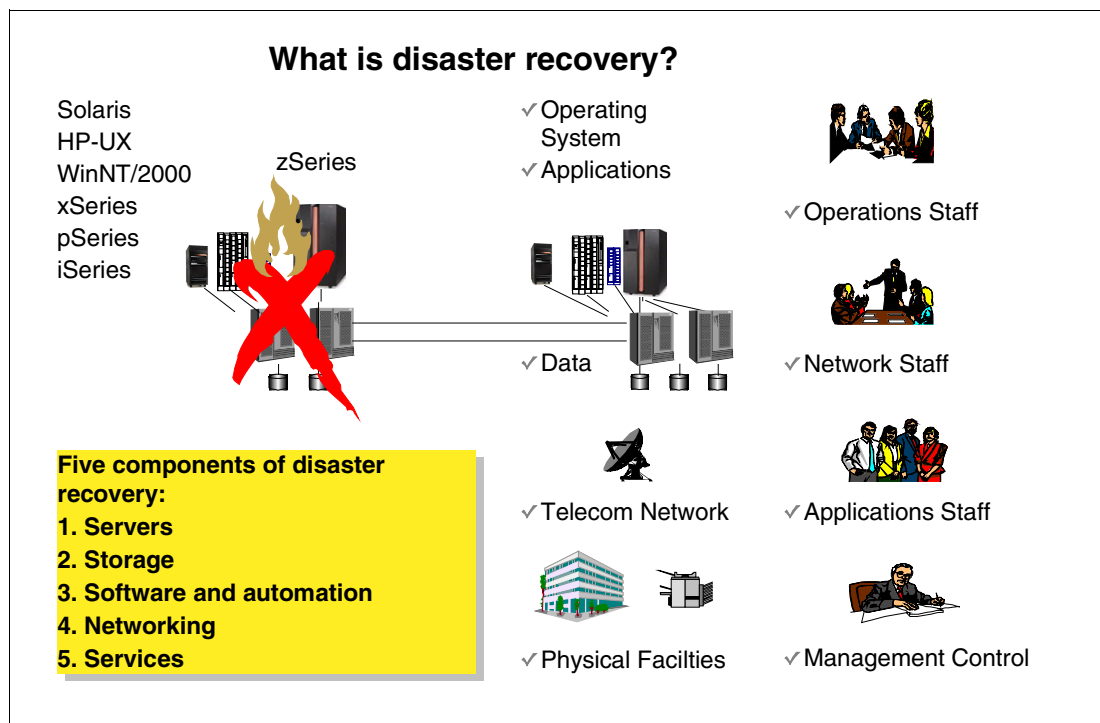
Figure 10-1   Components of disaster recovery

In today's highly competitive e-business world, outages can have a devastating impact on a business — they can even mean its demise. IBM suggests that disaster recovery is much more than just mirroring the disk data; rather, as Figure 10-1 illustrates, disaster recovery is a total business continuance solution comprising five major IT components:

► Servers
► Storage
► Software and automation
► Networking
► Services for integration

A disaster recovery implementation that only covers the storage component will leave the organization open to significant additional costs and time requirements if the other components are not covered.

10.1.2 Business objectives of disaster recovery

As shown in Figure 10-2, to design a cost-effective solution, we start determining the following objectives by application or business line:

▶ **Recovery Time Objective (RTO)**: What is the business cost-justified elapsed time to recovery?

▶ **Recovery Point Objective (RPO)**: When the Recovery Time Objective is met, what amount of data is needed to be recreated?

▶ **Network Recovery Objective (NRO)**: How long does it take to switch the entire network over to the backup data center?
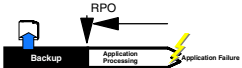
Figure 10-2   Business objectives of disaster recovery

These first two objectives (RTO and RPO) can often be balanced against each other to optimize the cost/benefit ratio. When the third objective (NRO) comes into play, networking issues will come into consideration, for example, there is no need to purchase a 30 minute RTO solution if the network provider requires two hours to switch the network.

Planned versus unplanned outages

Planned outages are just as effective at removing service from the end users as unplanned outages, and they are much more frequent. Yet, typically, disaster recovery solution cost justification is attempted based on the unplanned outage cost alone. So a realistic return of investment analysis should always consider the effects of planned as well as unplanned outages.

10.1.3 Tiers of disaster recovery

Finally, it is not only one RTO/RPO set that you are looking for. When looking for a solution, you will always want to know how much faster or slower the solution will be if you invest a little

more or less. This is the primary business issue, which leads us to the consideration of the tiers shown in Figure 10-3.
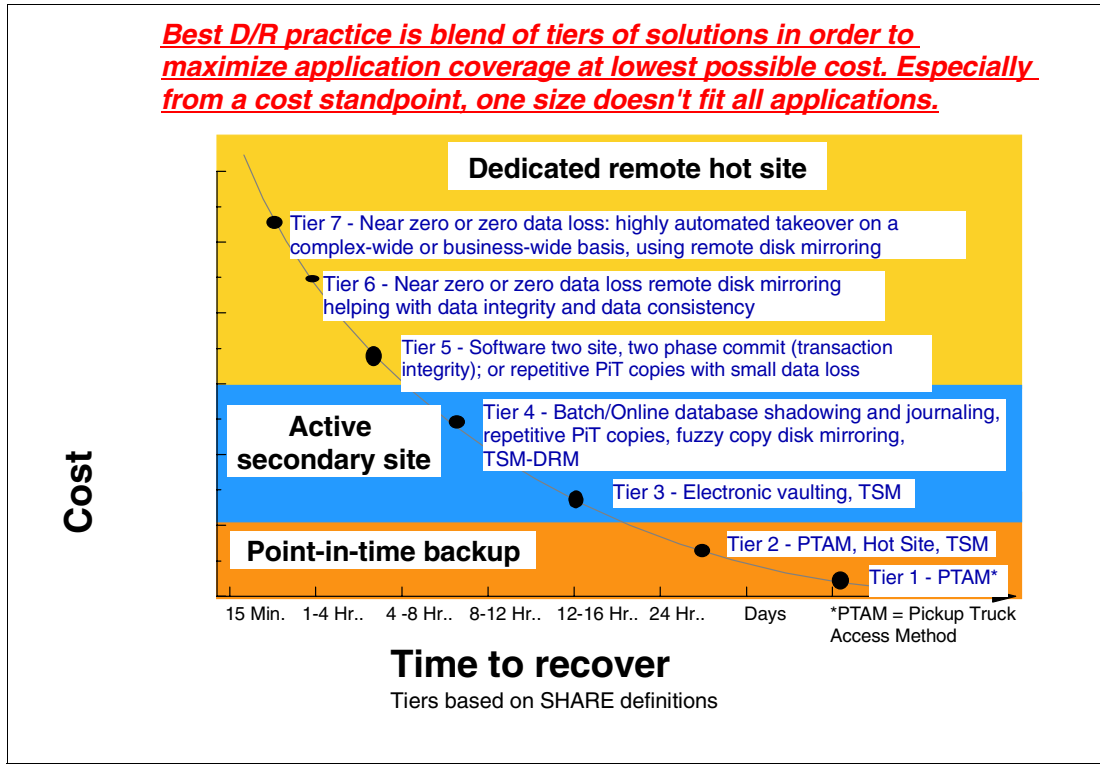
Figure 10-3   Tiers of disaster recovery

The chart in Figure 10-3 is a standard disaster recovery industry tier chart, showing that for varying types of RTO/RPO combinations, there are a variety of technologies.

In general, there are three main bands of application criticality. Within each band, there are tiers. The tiers vary as follows: tier 0 (no recovery capability), tier 1 (pickup truck access method), tier 2–3 (tape intensive recovery methods), tier 4 (disk mirroring or disk point-in-time facilities), tier 5 (software/ database specific recovery), tier 6 (near zero or zero data loss disk mirroring), and finally tier 7, which provides for complete Server/Workload/Network and Data Site Failover/Failback through end-to-end automation software.

10.2 Data consistency

Many applications, such as databases, process a repository of data that has been generated over a period of time. Many of these applications require that the repository be in a consistent state in order to begin or continue processing.

In general, consistency implies that the order of dependent writes is preserved in the data copy. For example, the following sequence might occur for a database operation involving a log volume and a data volume:

1. Write to log volume: "Data Record #2 is being updated".

2. Update Data Record #2 on data volume.

3. Write to log volume: "Data Record #2 update complete".

If the copy of the data contains any of these combinations, then the data is consistent:

► Operations 1, 2, and 3
► Operations 1 and 2
► Operation 1

If the copy of the data contains any of these combinations then the data is *inconsistent* (in other words, the order of dependent writes was *not* preserved):

► Operations 2 and 3
► Operations 1 and 3
► Operation 3

Note that "consistent" does not necessarily mean "equivalent" or up to date, only that the order of dependent writes is preserved. In addition, the order of non-dependent writes does not necessarily need to be preserved. For example, consider the following two sequences:

1. Deposit paycheck in checking account A

2. Withdraw cash from checking account A

3. Deposit paycheck in checking account B

4. Withdraw cash from checking account B

For the data to be consistent, the deposit of the paycheck must be applied *before* the withdraw of cash for each of the checking accounts. However, it does not matter whether the deposit to checking account A or checking account B occurred first, as long as the associated withdrawals are in the correct order. So for example, the data copy would be consistent if the following sequence occurred at the copy. In other words, the order of updates is not the same as it was for the source data, but the order of *dependent* writes is still preserved. For example:

1. Deposit paycheck in checking account B

2. Deposit paycheck in checking account A

3. Withdraw cash from checking account B

4. Withdraw cash from checking account A

Without the aid of host software, the Copy Services functions available in the ESS provide *power-fail* or *crash consistency*. In other words, the data is consistent (that is, the order of dependent writes is preserved) and looks to the application as if the storage and/or host lost power at some instant. Most database applications can recover from a *power-fail consistent* copy, although the recovery might take longer than a recovery from an *application consistent* copy.

Many users may choose to create an application consistent copy. However, host software must be involved in order to create this type of consistent copy. It involves quiescing application write I/Os in some way (quiesce database, put database in hot backup mode, and so on), flushing the host write cache, and then using the ESS Copy Services functions to create the copy.

10.3 Using ESS Copy Services for disaster recovery

With respect to the tiers of disaster recovery, ESS Copy Services add value mostly at tiers 3 to 6, although it can be useful also in other tiers:

Tier 0   Provides no preparation for saving information, establishing a backup hardware platform, or developing a contingency plan. From a disaster recovery perspective,

the type of data storage system that you are using is not relevant, as no system has a defined role.

Tier 1   May require full (image) backups at regular intervals to reduce recovery time. With ESS FlashCopy, the application impact of these backups can be reduced.

Tier 2   Encompasses all requirements of tier 1 and also requires a backup platform to have sufficient hardware and network to support the installation's critical processing requirements. Therefore, what we said for the tier 1 solution applies too.

Tier 3   Has the same requirements as tier 2 and, in addition, supports electronic vaulting of a subset of the information (probably the most critical). PPRC Extended Distance with FlashCopy for point-in-time consistency can be applied to this critical data.

Tier 4   Introduces the requirements of active management of the recovery data by utilizing a server host at the recovery site and bi-directional recovery. We can use PPRC-XD and FlashCopy in combination to make PIT copies.

Tier 5   Additionally maintains selected data that is always mirrored between copies (updates are committed only when applied successfully to both the local and remote copies). We configure synchronous PPRC (or Asynchronous Cascading PPRC in synchronous mode for the local pairs) to meet these requirements.

Tier 6   Encompasses a zero, or near-zero, data loss, and immediate and automatic transfer to the secondary platform. Here we need synchronous PPRC (possibly followed by PPRC-XD in cascading relationships for longest distances) for all data and cluster solutions at the host level.

Therefore the ESS brings tremendous value to your disaster recovery solution. The right balance needs to be found, depending on your business requirements. Figure 10-4 on page 405 summarizes the positioning of IBM TotalStorage disk subsystems, but please remember that the storage subsystem can guarantee only the automatic transfer of data. For recovery, the application processes must be transferred too. If the recovery time in your case must not exceed an hour, the integration of storage and server cluster technologies is necessary.

- **Enterprise solutions for Open Systems and zSeries**
  IBM TotalStorage Enterprise Storage Server
  - ► FlashCopy for point-in-time copy - *Tier 2*
  - ► PPRC-XD for remote copy and data migration - *Tier 4*
    - – Non-synchronous long or short distance
    - – IP connection capability (by channel extender vendors)
  - ► PPRC-SYNC - *Tier 6*
    - – Synchronous (supported up to 103 Km)
- **Enterprise solutions for zSeries**
  IBM TotalStorage Enterprise Storage Server
  - ► XRC - *Tier 6*
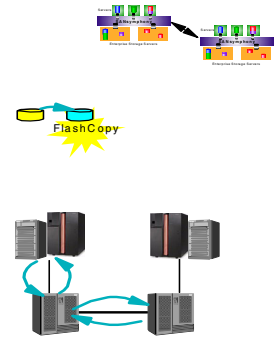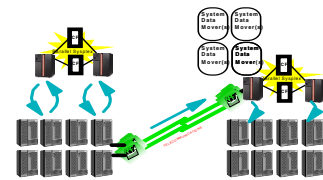
Figure 10-4   Disaster recovery storage solutions portfolio

When planning resources for your PPRC environment, you need to consider which data is critical and requires protection in the event of a disaster. Knowing this, it is possible to plan the throughput and capacity requirement of the secondary ESS.

It is important to realize that the capacity needed on the secondary ESS for disaster recovery may not have to be initially as large as the primary ESS. A disaster recovery plan (DRP) requires significant investment financially in technology, people, and processes. Every company will be different, but the I/T components of a disaster recovery plan are essentially driven by the applications and data you require for business continuity, should a disaster occur. Some applications and data will be more critical than others. An organization will typically require its core business systems to be available in a short time, whereas less critical systems quite possibly could be restored over a number of days.

Bearing in mind that the disk space you need for PPRC secondary volumes is real disk space, you need to size your secondary ESS based on your critical business requirements, possibly with some extra room for applications of intermediate importance. Create PPRC pairs for the critical data so that it is copied in real time. Then, if a disaster happens, you will have the core systems available on the secondary copies. After the initial recovery priorities have been handled, you can add more disk ranks for the applications of lower importance and restore them from tape.

10.4 Types of PPRC Failover and Failback procedures

Depending on the scenario and technical restrictions, different types of procedures are needed with respect to Copy Services server failover and volume management. Copy Services server failover depends on:

- ► The version of ESS Copy Services and server's operating mode
- ► The location of the active Copy Services server

With respect to volume management, we have to distinguish between planned and unplanned outages:

► During a *planned outage*, all storage subsystems, servers, and networks are functional. The last and current states of all components are well-defined. You can perform any operation without restriction.

Examples for planned outages are system maintenance, disaster recovery tests, and training. The objective of failover/failback procedures is to continue with consistent and current data (without any loss of transactions).

► When one or several components fail, we get an *unplanned outage*. The last state of some components may be unclear. You can operate only with components not affected by the outage (but even these operations may be restricted because of communication loss).

Examples of unplanned outages are fatal hardware failures or datacenter damage. The objective of failover/failback procedures is to continue with consistent data that is as up-to-date as possible (loss of some last transactions may be unavoidable).

So the minimum requirement for all situations is consistency at the application level. The level of currency achievable in each situation may differ.

Besides the primary difference between planned and unplanned outages, procedures differ also with respect to ESS microcode versions and PPRC operation modes:

► PPRC Extended Distance mode requires some minor changes.

► ESS LIC Version 2.1 introduced the PPRC initialization options **PPRC Failover** and **PPRC Failback**, can be used to reverse the primary and secondary roles.

– Older LIC versions do not provide these PPRC establish options.

PPRC Failover and **PPRC Failback** implement combinations of several steps in a single task. So even if your ESSs run up-to-date LIC versions, you should study the older procedures to understand the concept behind the new modes.

In the next sections, we consider these procedures in detail. 10.5, "Copy Services server failover" on page 406 discusses how to handle the loss of Copy Services servers. 10.6, "Planned outages" on page 408 and 10.7, "Unplanned outages" on page 422 outline procedures for typical scenarios with and without the usage of **PPRC Failover** and **PPRC Failback** options. At the beginning of each procedure, the situation is specified. In practice, not all situations will fit into these classes, so it is important to focus on the underlying ideas.

For simplicity, we consider only a single PPRC relationship. With Asynchronous Cascading PPRC, both relationships are handled separately, depending on the site that is affected by the outage (see 4.4, "Asynchronous Cascading PPRC" on page 65).

10.5 Copy Services server failover

If a disaster at your production site causes you to lose connections to the Copy Services server, you must switch control to the server at your recovery site. Refer back to Chapter 2, "Implementing ESS Copy Services" on page 13 to review the differences between the operation modes dual-active, mixed, and single-active. There you will also find recommendations regarding the servers locations.

Handling of the Copy Services server in a disaster scenario depends on whether your Copy Services servers are operating in dual-active, mixed, or single-active mode.

10.5.1 Dual-active configuration

If the servers operate in the dual-active configuration, then there is no fail-over or fail-back procedure necessary. You can establish a connection to the surviving Copy Services server at the recovery site and perform the PPRC path and volume procedures.

> **Attention:** When the Copy Services server at the production site is operational again, the task repository in ServerB is overwritten. Therefore you should not make any changes to the task repository if ServerA is not available. The task repository changes whenever you create, modify, or delete a task or task group.

10.5.2 Mixed configuration

If the servers operate in the mixed configuration and an outage occurs at the production site, then the procedure depends on the location of the dual-active server.

► Dual-active server at production site and single-active server at recovery site:

You have to enable the single-active backup server as the active server. Follow these steps:

a. Establish a browser connection to the ESS Launch panel on the single-active server at the recovery site.

b. Click **Tools** from the ESS Launch panel.

c. Access the **Available Actions** section of the Resetting the ESS Copy Services page, then click **Reset to Backup** to enable the backup.

You have to perform these steps for every ESS cluster at your recovery site which is running pre-LIC 2.2.

> **Attention:** When the Copy Services server at the production site is operational again, the task repository in the Backup server is overwritten. Therefore, you should not make any changes to the task repository if ServerA is not available. The task repository changes whenever you create, modify, or delete a task or task group.

► Dual-active server at recovery site and single-active server at production site:

In this case, the situation is similar to the case of two servers in the dual-active configuration. There is no failover procedure necessary. You can establish a connection to the Copy Services server at the recovery site and perform the PPRC path and volume procedures.

Because your dual-active server acts as ServerA in the mixed configuration, changes made to its task repository will be copied to the single-active server when the production site becomes operational again.

10.5.3 Single-active configuration

If you have both servers running pre-LIC 2.2 in a single-active configuration and an outage occurs at the production site, then the procedure depends upon which server you designated as the Primary Copy Services server.

► The Primary Copy Services server is at the recovery site:

There is no fail-over procedure necessary. You can establish a connection to the Copy Services server at the recovery site and perform the PPRC path and volume procedures.

Changes made to the task repository will be copied to the Backup Copy Services server when the production becomes operational again.

- ► The primary Copy Services server is at the production site: You have to enable the single-active Backup server as the active server.

  a. Establish a browser connection to the ESS Launch panel on the single-active server at the recovery site.

  b. Click **Tools** from the ESS Launch panel.

  c. Access the **Available Actions** section of the Resetting the ESS Copy Services page, then click **Reset to Backup** to enable the backup.

  You have to perform these steps for every ESS cluster at your recovery site that is running pre-LIC 2.2.

10.6 Planned outages

These procedures rely on two facts:

- ► Primary and secondary volumes are in a consistent and current state.
- ► Both ESSs are functional and reachable.

You can swap sites without any full copy operation by combining PPRC initialization modes. The exact number of steps depends on the ESS microcode level, as outlined in the following sections.

> **Note:** The procedures below are presented with creating the ESS Copy Services tasks in the moment when they are needed (on the fly). Although this works fine during a planned failover/failback operation, you should create and save these tasks in advance (to avoid mistakes and for usage also in unplanned operations).

All screenshots in the examples are made with ESS Copy Services V2 (LIC Version 2.2). The user interface for V1 differs with respect to icons and some options, but the steps to perform the procedures are the same.

10.6.1 Planned failover without PPRC Failover/Failback options

During normal operation, data is being replicated from the production site to the recovery site. We assume the following:

- ► Both ESSs (production and recovery) are functional and reachable.
- ► Applications are updating the primary volumes located in the production site.
- ► Paths are established from the production to recovery site.
- ► All PPRC volume pairs are in the duplex state (with PPRC-XD: duplex pending XD).

A planned failover involves the following steps (for PPRC, see Figure 10-5 on page 409; for PPRC-XD, see Figure 10-6 on page 410):

1. When the planned outage window is reached, the application must be quiesced to cease all write I/O updating the primary volumes.

   Depending on the host operating system, it may be necessary to dismount the primary volumes. (The secondary volumes will appear with the same volume identification as the primaries but as different devices. This may cause errors with some systems; see Chapter 8, "Open systems specifics" on page 327).

2. The PPRC copy pairs must be allowed to go to the duplex state.

   If synchronous PPRC is being used, then this will only apply when the relationships have just been established and the pairs are in the duplex pending state.

   If PPRC-XD is being used, then the volumes must go into the duplex state by issuing a `go-to-sync` (see 4.3, "PPRC extended distance (PPRC-XD)" on page 59). Alternatively, you can check to see if all data has been copied by using the volume information panels or the ESS Copy Services CLI.

3. Terminate the PPRC pair relationships once the copy pairs are in the duplex state.

4. Depending on your path design, terminate the paths between the primary and secondary LSSs.

   > **Attention:** Be careful when removing paths because the paths between primary and secondary LSSs may be used by PPRC relationships for other host systems that are not affected by the outage.

5. Depending on your path design, establish PPRC paths in the opposite direction from the recovery site to the production site.

6. Establish PPRC relationships with volumes in the recovery site as primaries. Ensure that you select the initialization mode **Do not copy volume**.
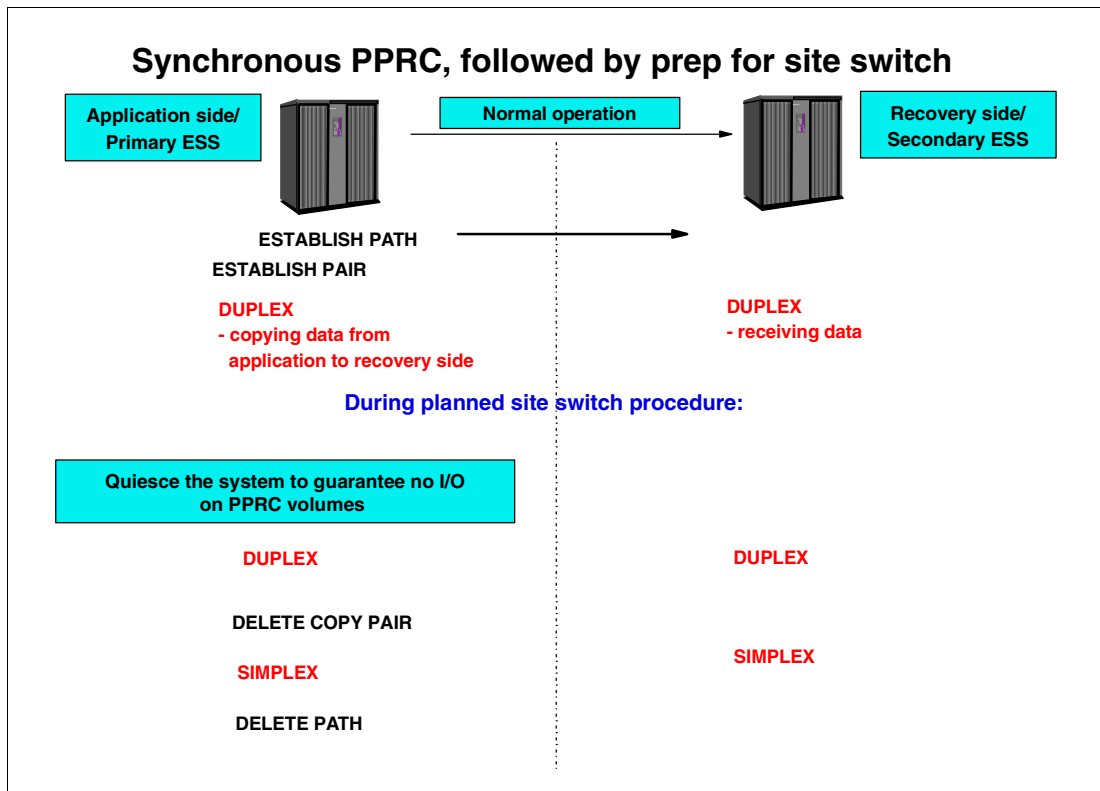
Figure 10-5   PPRC site switch preparation
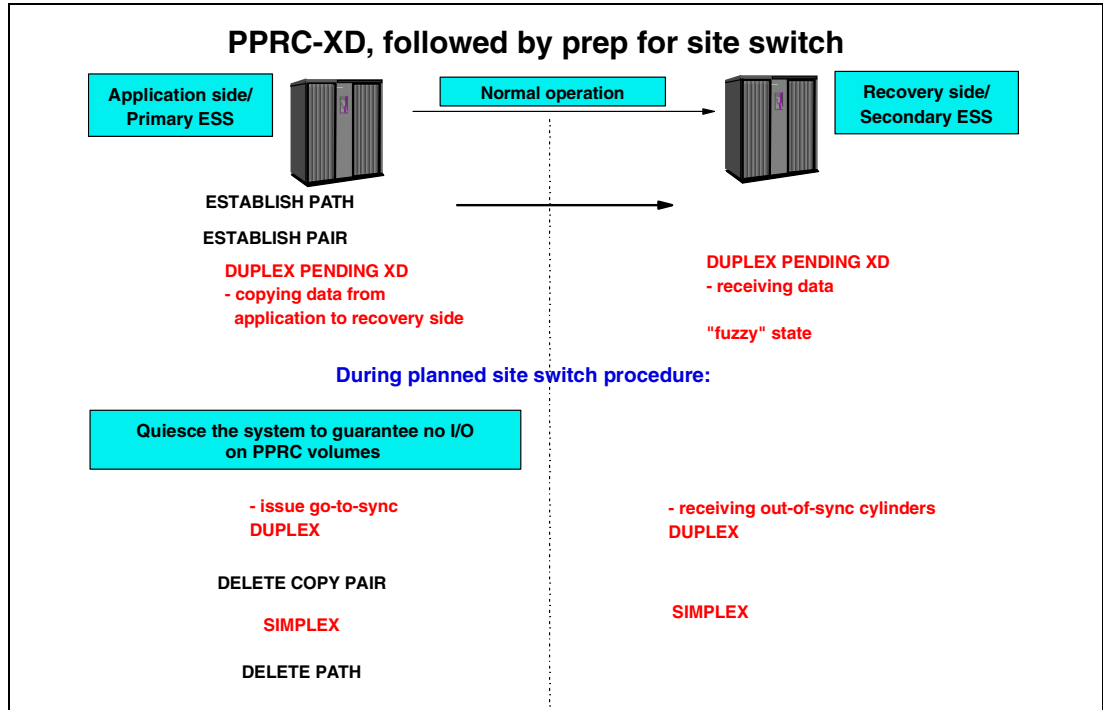
Figure 10-6   PPRC-XD site switch preparation

7. Immediately suspend the primary volume to start recording changes to the primary volume. (The option **Suspend PPRC after establish complete** is not available with the **Do not copy volume** initialization mode, so this step requires a new task.)

8. Depending on your operating system, it may be necessary to rescan Fibre Channel devices (to remove device objects for the production site volumes and recognize the new primaries) and mount the new primary volumes. Start all applications (see Figure 10-7 on page 411).

Reversing direction and tracking incremental changes

Application side/ Primary ESS

Switched operation

Recovery side/ Secondary ESS

ESTABLISH PATH

ESTABLISH PAIR *MODE(NOCOPY)*

DUPLEX

DUPLEX
- *no* data copy

SUSPEND PRIMARY

SUSPENDED

SUSPENDED
- PPRC is tracking incremental changes

After above steps complete, run production at secondary site
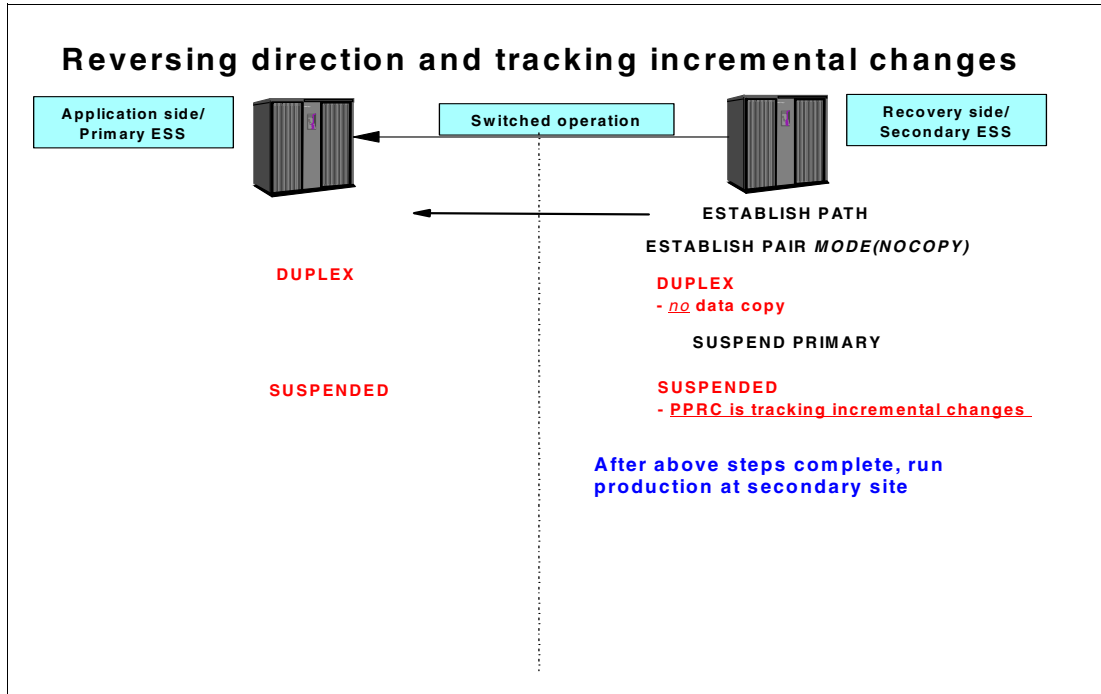
Figure 10-7   Initiating reverse synchronization

9. Now that the application has started, all the write I/O to the new primary volumes is being tracked by the bitmap created in step 7.

   Depending on your plans regarding the production site, the PPRC pairs can remain suspended (if you want to do offline maintenance). Alternatively, the updated tracks can now be propagated by re-establishing the synchronous PPRC with initialization mode **Copy out-of-sync cylinders only**, or go into the duplex pending XD mode.

There are two ideas to recognize in this procedure:

► In Step 6, it is assumed that the once primary and secondary volumes have not been updated while in the simplex state and are identical in contents (because we stopped all updates in Step 1).

► In Step 7, we prepare for failback with minimum data transfer. Additionally, we do not update the production site volumes until we are sure that the applications found are consistent with the current data at the recovery site.

10.6.2 Failback after planned failover without PPRC Failover/Failback options

Once you are ready to swap back to the production site, we have the following situation:

► Both ESSs (production and recovery) are functional and reachable.
► Applications are updating the primary volumes located in the recovery site.
► Paths are established from recovery to production site.
► The primary and secondary volumes are in the same state, either duplex or suspended.

The failback actions are a mirror of the fail-over actions. You perform the same steps as described in the section above, but with the production and recovery site exchanged. Step 7 can be omitted (but consider the tip below). Establishing the copy pair with the option **Do not copy volume** places the synchronous PPRC pair immediately into a duplex state. Therefore, on the start of the application, the write I/O will occur in a synchronous manner without

needing to go through the duplex pending. A PPRC-XD pair will be placed in the state duplex pending XD.

These fail-back steps can be seen in Figure 10-8 for PPRC and Figure 10-9 on page 413 for PPRC-XD.

Tip: Step 7 during the fail-back procedure is to suspend the PPRC pair, which has been newly established, after the fail-back procedure is essentially done (and before we run the applications using the volumes at the production site). This is a little trick that comes from good practices in a datacenter: you know that the volumes at the recovery site are consistent and current (because you have used them before switching sites back).

Now in step 8, you run the applications for the first time (after the failback) with the volumes at the production site. You check the applications. When you now see that something went wrong during the fail-back process and you followed the fail-back process without step 7, then you have overwritten (via PPRC) your last known good copy.

By including step 7, you have some more tricks in your bag. If all looks well, then in step 9 you transfer the new updates (made during the application check) and keep both the primary and secondary up. So you should consider including the suspend operation (step 7) in every case to not modify volumes known as consistent and current until you are sure that the volumes to be used after the swap are consistent and current too.
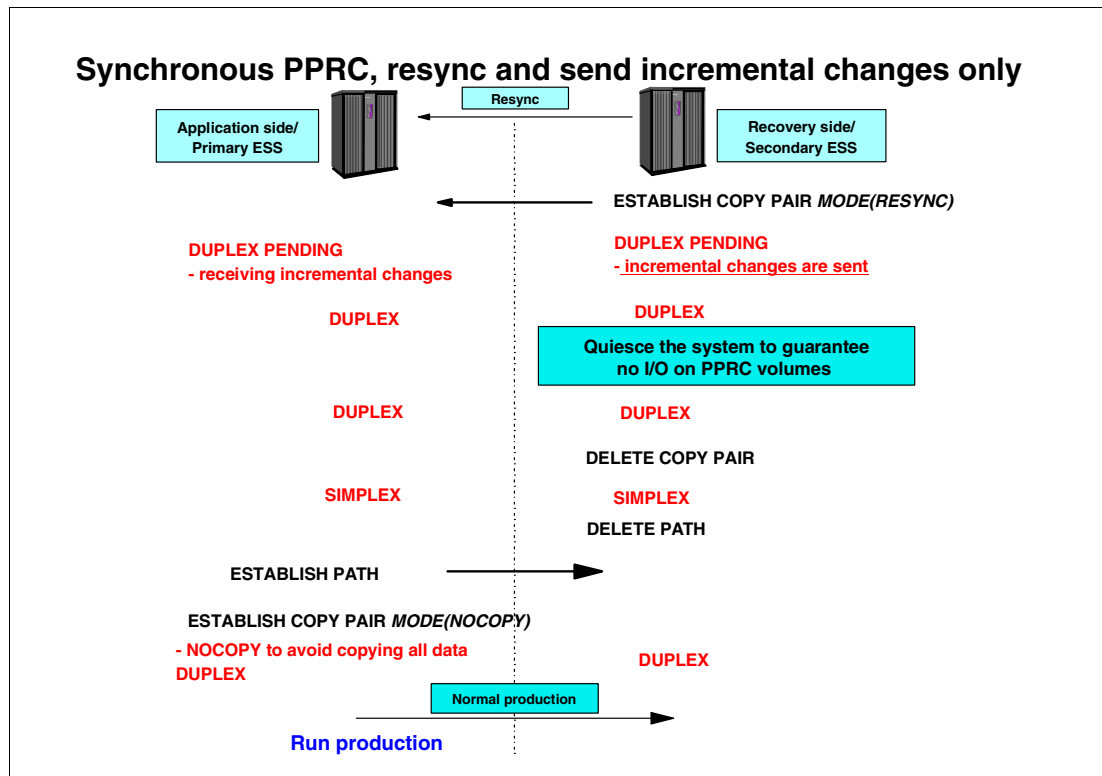
Figure 10-8   PPRC resumes at application site
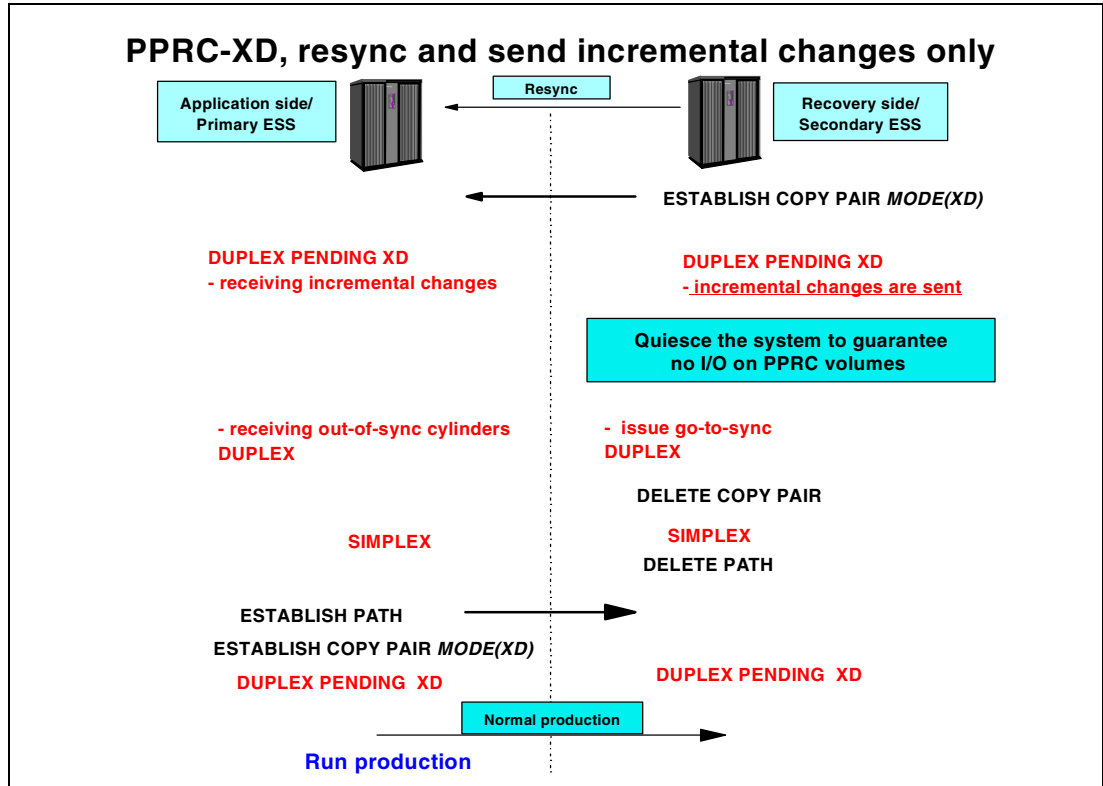
Figure 10-9   PPRC-XD resumes at application site

10.6.3 Planned failover with PPRC Failover/Failback options

With ESS LIC Version 2.1.0, the PPRC initialization modes **PPRC Failover** and **PPRC Failback** have been introduced to reduce the number of steps when switching sites.

In our example, we have used the ESS with serial number 18767 as the production site (with volume A) and the ESS with serial number 22331 (with volume B). Volume A is our primary volume with volume serial number 000-60018767 (Volume 000 on LSS 16). Volume B is the secondary volume with volume serial number 009-30922331 (Volume 009 on LSS 13). During normal operation, data is being replicated from volume A to volume B using PPRC. We assume the following:

► Both ESS (production and recovery) are functional and reachable.
► Applications are updating the primary volumes located in the production site.
► Paths are established from the production to the recovery site.
► All PPRC volume pairs are in the full duplex state.

A planned failover using the fail-over task involves the following steps:

1. When the planned outage window is reached, the application must be quiesced to cease all write I/O from updating the primary volumes.

   Depending on the host operating system, it may be necessary to dismount the primary volumes. (The secondary volumes will appear with the same volume identification as the primaries but as a different devices. This may cause errors with some systems; see Chapter 8, "Open systems specifics" on page 327).

2. The PPRC copy pairs must be allowed to go to the duplex state. For PPRC-XD, this means issuing a `go-to-sync` operation. Alternatively, you can check to see if all data has been copied by using the volume information panels or the ESS Copy Services CLI.

   If synchronous PPRC is being used, then this will only apply when the relationships have just been established and the pairs are in duplex pending state.

3. Create the fail-over task. You need to perform the following steps at the Volumes panel (Figure 10-10):

   a. Bring up the volumes of LSS 13 (containing recovery volume B) in the source column.

   b. Bring up the volumes of LSS 16 (containing production volume A) in the target area.

   > **Important:** For creating the fail-over task, you have to select the primary and secondary volume according to their roles after the failover has been completed.

   c. Click on volume B; the volume ID turns blue.

   d. Right-click on volume A, and the volume ID turns red.

   e. Right-click on volume A again to open the Task Wizard.
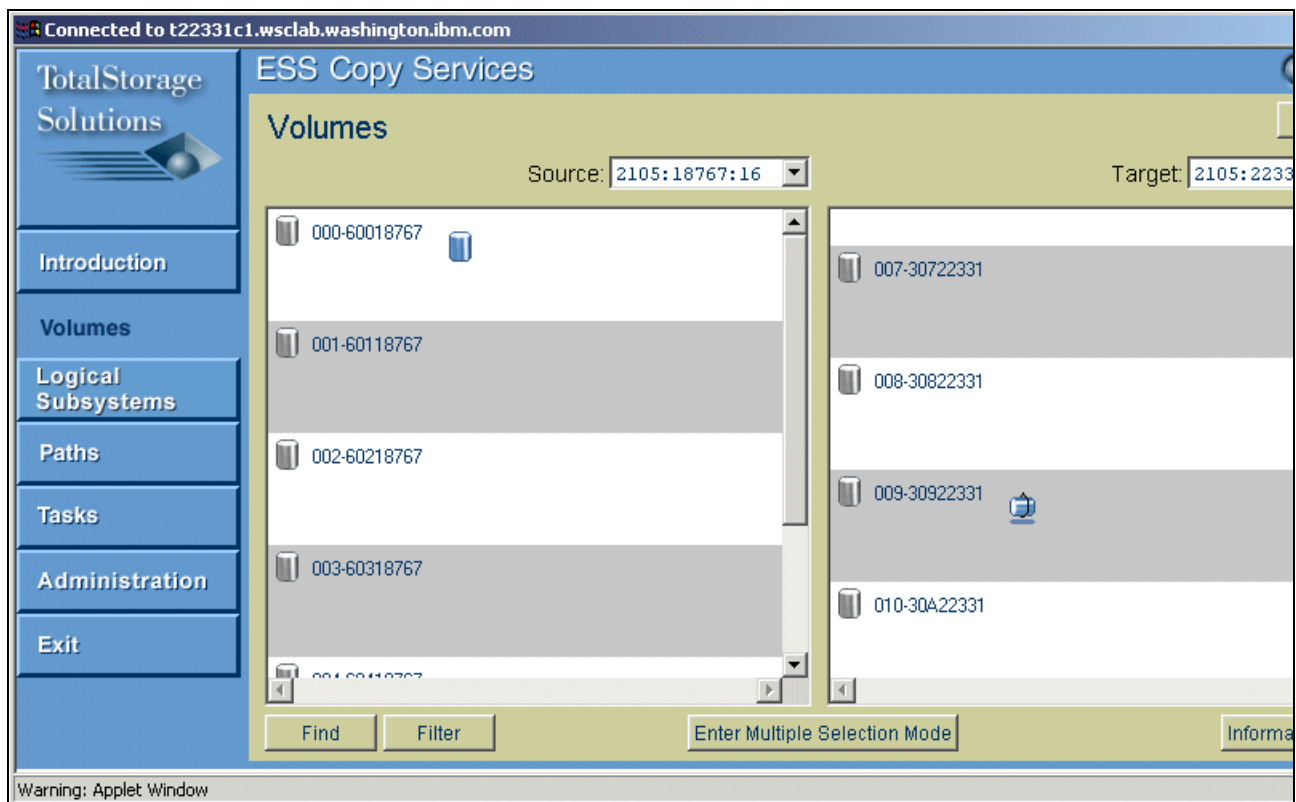
Figure 10-10   Create fail-over task

   a. Select the task type **Establish synchronous PPRC copy pair**, and click **Next**.

   a. On the next window, select the **PPRC Failover** option (Figure 10-11 on page 415).

   > **Restriction:** For PPRC-XD, the **PPRC Failover** option is not available, so you have to follow the procedure from 10.6.1, "Planned failover without PPRC Failover/Failback options" on page 408.
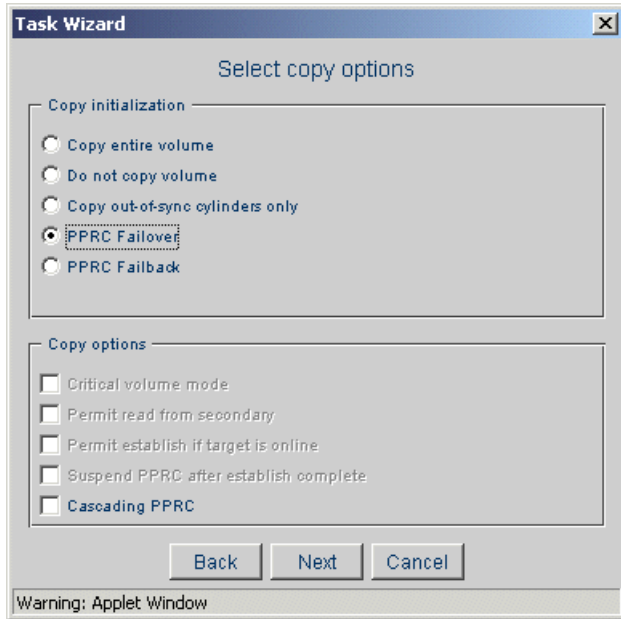
Figure 10-11   Task Wizard fail-over option

   b. Click **Next**. In the next window, you can save the task in the task repository.

4. Run the fail-over task.

   After you have successfully executed the fail-over task, volume B is the new primary volume in suspended state (Figure 10-12 on page 416). Volume A remains in the primary full-duplex state.

Figure 10-12   Volume state after failover

5. Depending on your path design, terminate the paths between the primary and secondary LSSs. You have to choose the option **Force removal of PPRC path even if pairs exist**.

> **Attention:** The paths between primary and secondary LSS may be used by PPRC relationships for other host systems that are not affected by the outage.

6. Depending on your path design, establish PPRC paths in the opposite direction from recovery site to production site.

   Although it is not strictly necessary to reverse the paths, we recommend you do so to have a well-defined situation at the end of the procedure. Additionally, you will need the paths in the opposite direction if you want to transfer the updates back to the production site.

7. Depending on your operating system, it may be necessary to rescan Fibre Channel devices (to remove device objects for the production site volumes and recognize the new primaries) and mount the new primary volumes. Start all applications.

8. Now that the application has started, all the write I/O to the new primary volumes is being tracked by the fail-over task.

   Depending on your plans regarding the production site, the PPRC pairs can remain suspended (if you want to do offline maintenance). Then the fail-back task will initiate the data transfer back to the production site.

   Compared with the procedure for the older LIC (10.6.1, "Planned failover without PPRC Failover/Failback options" on page 408), we realize the following differences:

   – The **PPRC Failover** option combines the three steps to terminate the old PPRC relationship and establish and suspend the new relationship into a single task.

– The production site volume remains in the primary full-duplex state instead of changing to secondary suspended.

– There is no PPRC termination at any time. Therefore, when switching path directions, the option **Force removal of PPRC path even if pairs exist** is needed.

10.6.4 Failback after planned failover with PPRC Failover/Failback options

Once the production site has been restored, you must move your application back. We use the same volumes as in 10.6.3, "Planned failover with PPRC Failover/Failback options" on page 413 as an example on how to create the fail-back task.

In our example, we have used the ESS with serial number 18767 as the production site (with volume A, serial number 000-60018767) and the ESS with serial number 22331 (with volume B, serial number 009-30922331). During normal operation, data would be replicated from volume A to volume B using PPRC, but now volume B is our primary volume after the failover.

We assume that, during an operation at the recovery site, data has not been replicated from volume B to volume A:

► Both ESSs (production and recovery) are functional and reachable.

► Applications are updating the primary volumes located at the recovery site.

► Paths are established from the recovery to the production site.

► Volumes at the recovery site are in the suspended state (primary), volumes at the production site are in the duplex state (primary).

Remember that the outcome of 10.6.3, "Planned failover with PPRC Failover/Failback options" on page 413 was that volume B was in the primary suspended state and volume A was still in the primary duplex state. A failback using the **PPRC Failover** and **PPRC Failback** options involves the following steps:

1. Create the fail-back task from the recovery site to the production site. You need to perform the following steps on the Volumes panel:

   a. Bring up LSS 13 (with the recovery site volume B) in the source column.

   b. Bring up LSS 16 (with the production site volume A) in the target column.

   c. Click on volume B; the volume ID turns blue.

   d. Right-click on volume A; the volume ID turns red.

   e. Right-click on volume A again to open the Task Wizard (Figure 10-13 on page 418).
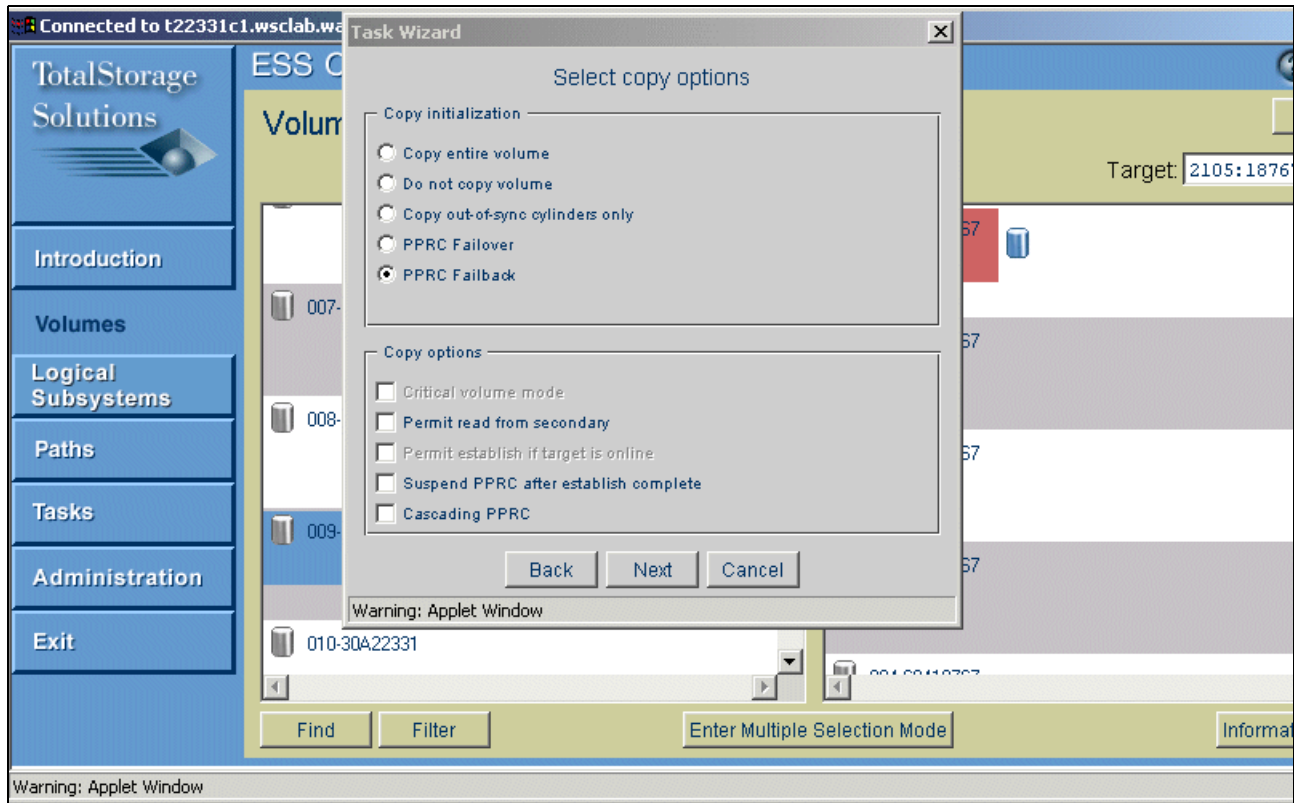
Figure 10-13   Create tailback task

f.   Select the **Establish synchronous PPRC copy pair** task type and click **Next**.

g.   On the next window, select the **PPRC Failback** option (Figure 10-11 on page 415). In contrast to the **PPRC Failover** option, this is available also with PPRC-XD.

h.   Click **Next**. In the next window, you can save the task in the task repository.

2.   Run the fail-back task from the recovery site to the production site. The fail-back task will copy all the tracks modified on volume B to volume A.

If you did not reverse the PPRC paths during failover, you have to establish paths from the recovery to the production site before running the fail-back task.

> **Note:** The PPRC Failback initialization mode resynchronizes the volumes.
>
> ► If a volume at the production site is in simplex state, all of the data for that volume is sent from the recovery site to the production site.
>
> ► If a volume at the production site is in the full-duplex or suspended state and without changed tracks, only the modified data on the volume at the recovery site is sent to the volume at the production site.
>
> ► If a volume at the production site is in a suspended state and has tracks that have been written to, the volume at the recovery site will discover which tracks were modified on any site and send both the tracks changed on the production site and the tracks marked at the recovery site.
>
> The volume at the production site becomes a write-inhibited secondary volume. This action is performed on an individual volume basis.

Figure 10-14   Volume state after the failback

3. Wait until data transfer has been completed (Figure 10-14). After the volumes are in full duplex, volume B (serial number 009-30922331) is a primary volume, and volume A (serial number 000-60018767) is a secondary volume.

4. Before returning to normal operation, the application (still updating volumes in the recovery site) must be quiesced to cease all write I/O from updating the primary volumes.

   Depending on the host operating system, it may be necessary to dismount the primary volumes. (The secondary volumes will appear with same volume identification as the primaries but as a different devices. This may cause errors with some systems; see Chapter 8, "Open systems specifics" on page 327).

5. To accomplish the return to normal operations, you should now execute one more failover. This converts the full-duplex secondary volumes at the production site to suspended primary volumes, and the volumes will also remember their correct partners' address.

   To create this additional fail-over task, you need to perform the following steps at the Volumes panel:

   a. Bring up LSS 16 (with the production site volume A) in the source column.

   b. Bring up LSS 13 (with the recovery site volume B) in the target column.

   > **Important:** To create the fail-over task, you have to select the primary and secondary volume according to their roles after the failover has been completed.

   c. Click on volume A; the volume ID turns blue.

   d. Right-click on volume B; the volume ID turns red.

   e. Right-click on volume B again to open the Task Wizard.

f.   Select the task type **Establish synchronous PPRC copy pair**, and click **Next**.

g.   Within the next window, select the **PPRC Failover** option and click **Next**.

> **Restriction:** For PPRC-XD, the **PPRC Failover** option is not available, so you have to use multiple tasks to terminate the A–B PPRC relationship and establish and suspend the B–A relationship.

h.   Within the next window, you can save the task.

6. Run the additional fail-over task. Volume A will become the primary volume in the suspend state (Figure 10-15).
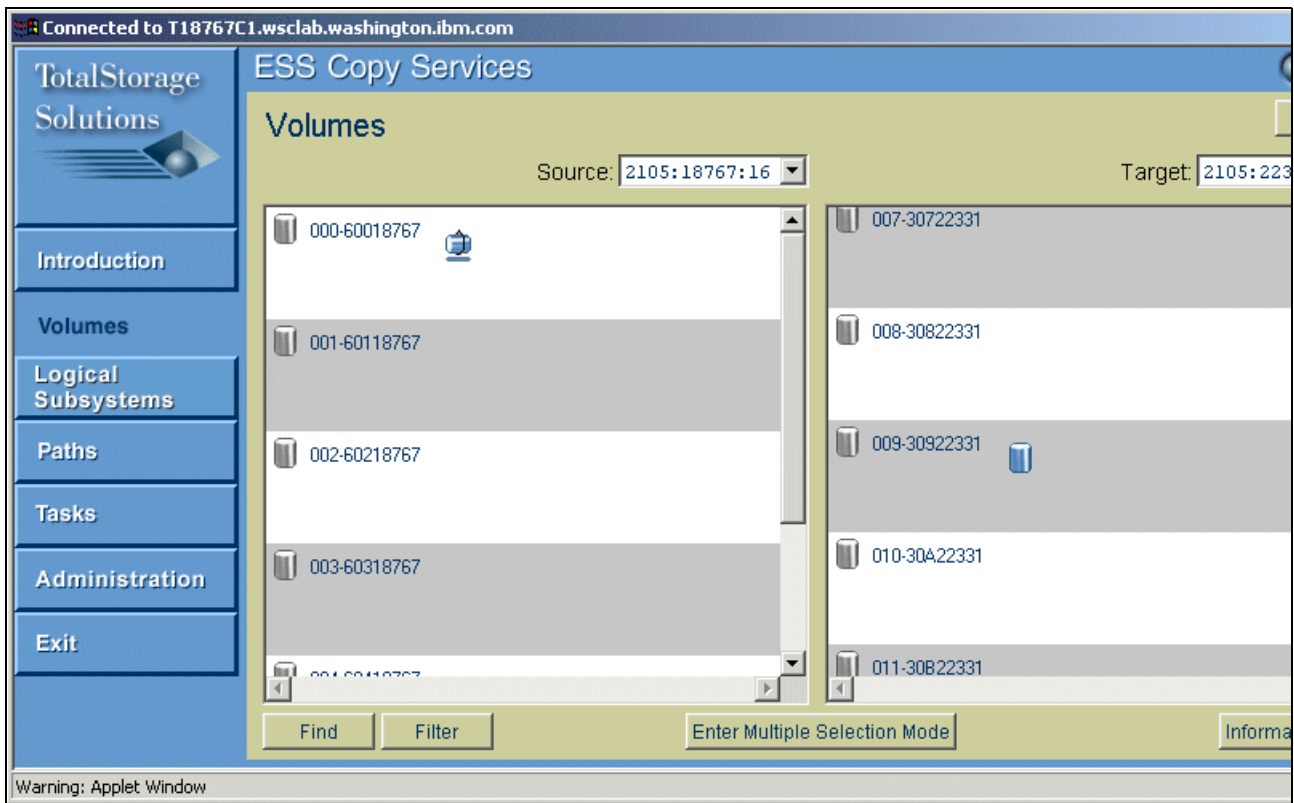
Figure 10-15   Failover after recovery

7. Depending on your operating system, it may be necessary to rescan Fibre Channel devices and mount the new primary volumes at the production site. Start all applications and check for consistency.

   Now that the application has started, all the write I/O to the new primary volumes is being tracked by the bitmap created. You should verify the application's integrity.

8. Depending on your path design, terminate the paths from the recovery to the production LSSs. Because the PPRC relationship has not been terminated, you have to choose the option **Force removal of PPRC path even if pairs exist**.

> **Attention:** Be careful when removing paths because the paths between the primary and secondary LSS may be used by PPRC relationships for other host systems that are not affected by the outage.

9. Depending on your path design, establish PPRC paths in the opposite direction from production site to recovery site.

10. Then, issue another fail-back task to return to normal operations. Perform the following steps to create the fail-back task at the Volumes panel:

    a. Bring up LSS 16 (with the production site volume A) in the source column.

    b. Bring up LSS 13 (with the recovery site volume B) in the target column.

    c. Click on volume A; the volume ID turns blue.

    d. Right-click on volume B; the volume ID turns red.

    e. Right-click on volume B again to open the Task Wizard.

    f. Within the Task Wizard select **Establish synchronous PPRC copy pair** and click **Next**.

    g. Select the **PPRC Failback** option in the next window and click **Next**.

    h. Within the next window, you can save the task.

11. If you run the task, volume A (which was a suspended primary) will become a full duplex primary, and volume B will become a full duplex secondary. You are returned to normal operation.

This difficult-looking procedure can be summarized as follows:

► The **PPRC Failover** option combines the three steps to terminate the old PPRC relationship, and establish and suspend the new relationship into a single task. The state of the old primary volume is preserved. This design takes into account that the old primary LSS may be no longer reachable.

   Considering the whole failover/failback process, this option has been used twice:

   – In the fail-over procedure's step 3 on page 414, by selecting the recovery site volume as primary and the production site volume as secondary (performing the fail-over role switch).

   – In the fail-back procedure's step 5 on page 419, by selecting the production site volume as primary and the recovery site volume as secondary (performing the fail-back role switch).

► The **PPRC Failback** option checks the preserved state of the old primary volume to determine how much data to copy back. Then either all tracks or only out-of-sync tracks are copied, and the old primary volume becomes a secondary full-duplex (the outcome of step 3 on page 419).

   Consider the whole fail-over/fail-back process, which has been used twice too:

   – In the failback procedure's step 1 on page 417, by selecting the recovery site volume as primary and the production site volume as secondary (beginning the fail-back process).

   – In the fail-back procedure's step 10 on page 421, by selecting the production site volume as primary and the recovery site volume as secondary (completing the fail-back process).

Figure 10-16 on page 422 summarizes the whole process of fail-over and fail-back with these four tasks.

▶ Therefore, when establishing a PPRC volume pair with mode **PPRC Failover**, you have to select primary and secondary volumes according to their desired roles after the task has been completed. The task enforces the new primary's role.

Conversely, when establishing a PPRC volume pair with mode **PPRC Failback**, you have to select primary and secondary volumes according to their actual roles before the task will run. The task enforces the new secondary's role.
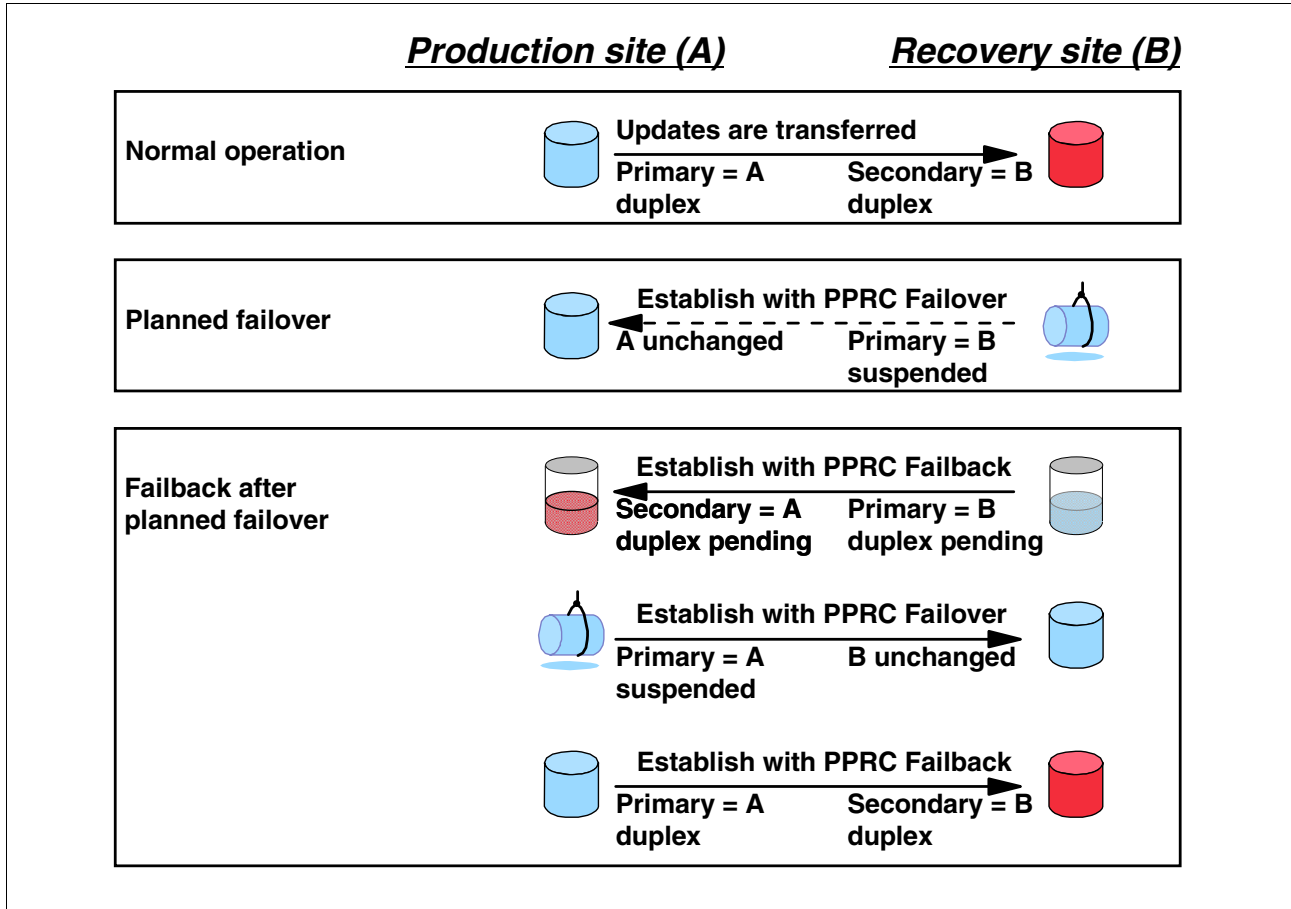
Figure 10-16   Combining the options PPRC Failover and PPRC Failback

In this manner, **PPRC Failover** and **PPRC Failback** are dual operations. It is possible to implement all site switch operations with two pairs of PPRC Failover and PPRC Failback tasks (one pair for each direction).

10.7 Unplanned outages

In contrast to the assumptions for planned outages (11.6, "Planned outages" on page 376), the situation in a disaster is more difficult:

▶ At the beginning of a planned outage, both ESSs are functional and reachable.

In a disaster situation, only one ESS (at the recovery site) is functional and reachable. The production ESS may be lost or unreachable.

▶ The procedures for planned failover ensure that the volumes at the production and recovery site are in the same state, and that is sufficient to transfer back only the changed tracks during the failback.

After an unplanned failover, volumes at the production and recovery site may be in different states.

► When planning an outage, you stop all I/O at the production site and check that all changes are transferred to the recovery site. So the secondary volumes are in a consistent and current state before beginning the failover procedure.

In a disaster situation, secondary PPRC-XD volumes are typically not consistent and not in a current state. Synchronous PPRC volumes are in a current state. But without using consistency groups, you can assume power-failure consistency only at the level of a single volume, not at the application level.

The problem that only the recovery site ESS is functional and reachable imposes some operational restrictions. For example, all tasks must be scheduled to the secondary ESS. You cannot establish and suspend PPRC in the opposite direction prior to ESS LIC Version 2.1.0.

The same applies to the issue of volume states: it requires some additional care at the beginning of the fail-back procedure. There are situations requiring a full copy from the recovery to the production site. (The task options **PPRC Failover** and **PPRC Failback** handle that automatically.)

But the problem of consistency and currency is a critical one: we cannot recover from a secondary volume that is not consistent. Even if the single volume was kept consistently up-to-date by synchronous PPRC, the recovery may fail at the application level. Consider the following scenario:

1. A fire at the production site first destroys some (not all) fiber links to the recovery site, but the ESS and servers are still working. Without consistency grouping, this causes some primary PPRC volumes to change their state to suspended, creating change bitmaps in the production ESS.

2. When the fire spreads out, more volumes are affected, and finally the ESS or servers are stopped.

3. You start the site failover, resulting in suspended primary volumes in the recovery ESS (with change bitmaps too).

Even with synchronous PPRC, you cannot guarantee the consistency of the secondary volumes. The key to avoid such situations is the usage of consistency groups (see "PPRC data consistency" on page 71). If the volumes belong to consistency groups, all primary volumes will enter the Queue Full state at the first occurrence of an update failure of a secondary volume. Transactions are blocked immediately, and the sites can be switched with application consistent volumes. (Some customers set the long busy timer to be days long so it is guaranteed that there would be no updates to the primary volumes.)

> **Important:** Usage of consistency groups is essential to recovery from an unplanned outage with synchronous PPRC (refer back to 4.5.2, "PPRC consistency groups" on page 74). With PPRC-XD, you have to maintain consistent point-in-time copies, as described in 4.3.3, "Creating a consistent point-in-time copy" on page 61 and 4.4.4, "Creating a consistent copy with PPRC-XD as a remote pair" on page 69.

Similar to 10.6, "Planned outages" on page 408, the procedures in the sections below describe only the volume-level operations. If the outage also affects the active Copy Services server, then you must follow the guidelines for Copy Services server availability from Chapter 2, "Implementing ESS Copy Services" on page 13 before taking volume-level actions.

Important: The procedures below are presented with creating the ESS Copy Services tasks in the moment when they are needed (on the fly). In practice, this will not work for an unplanned failover! When one ESS is not functional or reachable, you cannot create these tasks (although they will run with only one ESS available). Therefore you must create the tasks in advance before the disaster occurs.

All screenshots in the examples are made with ESS Copy Services V2 (LIC Version 2.2). The user interface for V1 differs with respect to icons and some options, but the steps to perform the procedures are the same.

10.7.1 Unplanned failover without PPRC Failover/Failback options

During normal operation, data is being replicated from the production site to the recovery site. We assume the following:

► Only the recovery ESS is functional and reachable.

► Applications were updating the primary volumes located in the production site, but these volumes may not be available.

► Paths have been established from the production to recovery site, but may be no longer functional.

► Volumes at the recovery site are in the duplex state (secondary) or suspended. Volumes at the production site were in the duplex state (primary), but they may be now suspended. The exact combination of volume states depends on the nature of the outage.

An unplanned failover involves the following steps:

1. Terminate the PPRC pair relationships from the production to the recovery site. Because the production ESS is no longer reachable, you have to select the option **Schedule task with target logical subsystem** for this task.

   The recovery site volumes will go into the simplex state. When the production ESS becomes available, their volumes may appear still as duplex (source), suspended, or simplex, depending on the nature of the outage.

2. Depending on your operating system, it may be necessary to rescan Fibre Channel devices (to remove device objects for the production site volumes and recognize the new primaries) and mount the recovery site volumes.

3. Start all applications.

Compared with 10.6.1, "Planned failover without PPRC Failover/Failback options" on page 408, you miss all the steps to prepare an easy failback. Depending on the exact situation, you should try to include these steps. But operations like establishing new paths and PPRC relationships in the opposite direction fail if the other ESS cannot be reached, and you cannot create change bitmaps in the recovery site when the volumes are not PPRC primaries.

These problems are solved with the **PPRC Failover** and **PPRC Failback** options, introduced with ESS LIC Version 2.1.0.

10.7.2 Failback after unplanned failover without PPRC Failover/Failback options

Once you are ready to swap back to the production site, we have the following situation:

► Both ESSs (production and recovery) are functional and reachable.

► Applications are updating the volumes located in the recovery site.

► The path situation may be unclear.

► Volumes at the recovery site are in a simplex state; the volumes at the production site are in the duplex state (source), suspended, or simplex.

The procedure to return to normal operation depends on the exact nature of the disaster:

1. Repeat the termination of PPRC relationships from production to recovery site, but now with the **Schedule task with source logical subsystem** option.

   This is necessary when the production volumes are in any other state than simplex, blocking all further PPRC operations.

2. Depending on your path design and the exact nature of the disaster, paths still existing from the production to the recovery site must be removed.

3. Depending on your path design, establish PPRC paths from the recovery site to the production site.

4. Establish PPRC relationships with volumes in the recovery site as primaries. Ensure that you select initialization mode **Copy entire volume** to copy all tracks from the recovery to the production volume, overwriting the whole secondary volume.

   The volume pair changes from state simplex to duplex pending and finally to duplex after the copy is completed.

   > **Important:** Use the initialization mode **Copy out-of-sync cylinders only** if you are absolutely sure that there are no changes made to the production site volumes.

5. Now the application must be quiesced to cease all write I/O updating the recovery site volumes.

6. Depending on the host operating system, it may be necessary to dismount the primary volumes. (The secondary volumes will appear with same volume identification as the primaries but as a different devices. This may cause errors with some systems; see Chapter 8, "Open systems specifics" on page 327).

7. Depending on your path design, terminate the paths from the recovery to the production LSSs.

8. Depending on your path design, establish PPRC paths in the opposite direction from the production site to recovery site.

9. Establish PPRC relationships with volumes in the production site as primaries. Ensure that you select initialization mode **Do not copy volume**.

10. Immediately suspend the primary volume to create a cylinder update tracking bitmap. (The option **Suspend PPRC after establish complete** is not available for **Do not copy volume** initialization mode, so this step requires a new task.) See the Tip in 10.6.2, "Failback after planned failover without PPRC Failover/Failback options" on page 411 for more discussion about the need to suspend.

11. Depending on your operating system, it may be necessary to rescan Fibre Channel devices (to remove device objects for the production site volumes and recognize the new primaries) and mount the new primary volumes. Start all applications at the production site.

12. Now that the application has started, all the write I/O to the new primary volumes is being tracked by the bitmap created in step 10. When you have checked the application consistency, re-establish the original PPRC operation with initialization mode **Copy out-of-sync cylinders only**, or go into the duplex pending XD mode.

There are two details to notice about this procedure:

► In step 9, it is assumed that the once primary and secondary volumes have not been updated while in the simplex state and are identical in contents (because we stopped all updates in step 5).

► Step 10 is not strictly necessary, but you should include the suspend operation for this reason:

– You do not modify recovery site volumes known as consistent and current (because they were used before the swap) until you are sure about the volumes in the production site after the swap.

10.7.3 Unplanned failover with PPRC Failover/Failback options

During normal operation, data is being replicated from the production site to the recovery site. We assume the following:

► Only the recovery ESS is functional and reachable.

► Applications were updating the primary volumes located in the production site. But these volumes may not be available.

► Paths have been established from production to recovery site, but may be no longer functional.

► Volumes at the recovery site are in state duplex (secondary) or suspended. Volumes at the production site were in state duplex (primary), but they may be now suspended. The exact combination of volume states depends on the nature of the outage.

With ESS LIC Version 2.1.0, the PPRC establishment modes **PPRC Failover** and **PPRC Failback** have been introduced, which allow you to prepare for easy failback even if only one ESS is available. The fail-over procedure is basically the same as for planned failover, as explained in 10.6.3, "Planned failover with PPRC Failover/Failback options" on page 413, so we provide Figure 10-17 on page 427 as a summary.
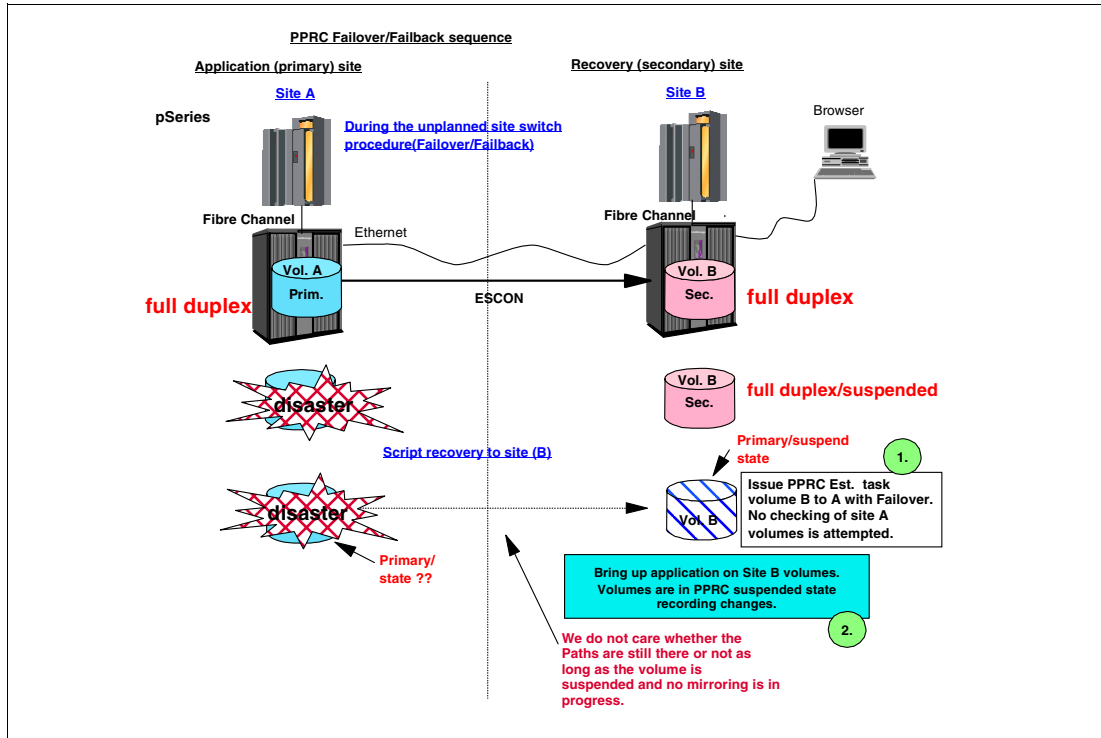
Figure 10-17   PPRC Failover option

Compared with 10.7.1, "Unplanned failover without PPRC Failover/Failback options" on page 424, the PPRC Failover option turns out to be useful for two reasons:

► The **PPRC Failover** option allows the new primary volume to change state and establish a relationship even without having sight of the new secondary volume, by using relationship information maintained in the ESS database.

  The establish operation will still succeed even if the new secondary is in a bad state or non-existent or even if the paths are down; however, the state of the old primary volume is preserved.

► Because a PPRC relationship from the recovery to the production site is established, the recovery site volume can become a suspended primary. This gives a chance for failback without full copy.

10.7.4 Failback after unplanned failover with PPRC Failover/Failback options

We assume that, during operation at the recovery site, data has not been replicated from recovery site volume B to production site volume A:

► Both ESSs (production and recovery) are functional and reachable.

► Applications are updating the primary volumes located at the recovery site.

► The path situation may be unclear.

► Volumes in the recovery site are in the suspended state (primary), and volumes in the production site are in the duplex state (source), suspended, or simplex.

The fail-back procedure is essentially the same as after a planned failover, as explained in 10.6.4, "Failback after planned failover with PPRC Failover/Failback options" on page 417. Depending on the exact nature of the disaster, some details vary:

► Before running the fail-back task, you may have to remove old paths still existing from the production to the recovery site with the **Force removal of existing path** option selected, which will force the paths to be removed even if there are PPRC copy pairs using those paths. (Refer to Figure 10-18.)
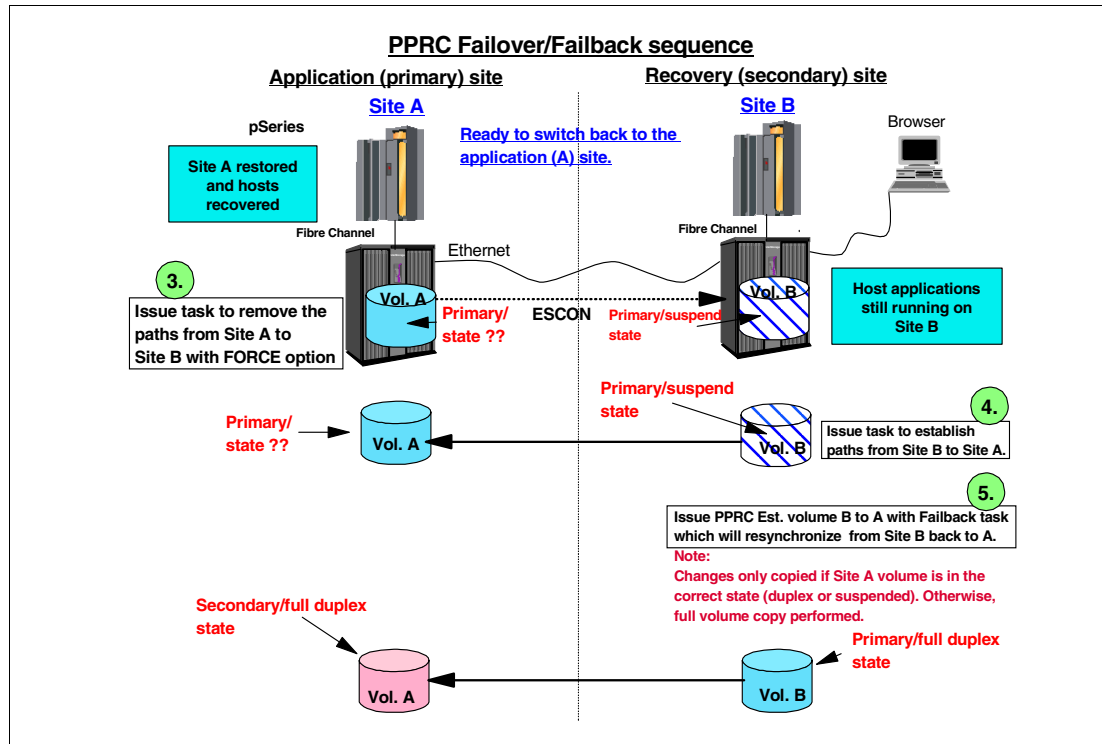
Figure 10-18   PPRC Failback option

► For data to be transferred back from the recovery site, it is necessary to establish PPRC paths from the recovery to the production site.

► When establishing the copy pair with the **PPRC Failback** option, the ESS at the recovery site will query the preserved state of the original primary volume (volume A) to determine how much data to copy back, as shown in Figure 10-18.

> **Note:** The PPRC Failback initialization mode resynchronizes the volumes.
>
> ► If a volume at the production site is in the simplex state, all of the data for that volume is sent from the recovery site to the production site.
>
> ► If a volume at the production site is in the full-duplex or suspended state and without changed tracks, only the modified data on the volume at the recovery site is sent to the volume at the production site.
>
> ► If a volume at the production site is in a suspended state and has tracks that have been written to, the volume at the recovery site will discover which tracks were modified on any site and send both the tracks changed on the production site and the tracks marked at the recovery site.
>
> The volume at the production site becomes a write-inhibited secondary volume. This action is performed on an individual volume basis.

▶ Once volume B and volume A are in the duplex state and you are ready to switch back to your production site, you can continue according to the fail-back method described in 10.6.4, "Failback after planned failover with PPRC Failover/Failback options" on page 417. This can be seen in Figure 10-19.
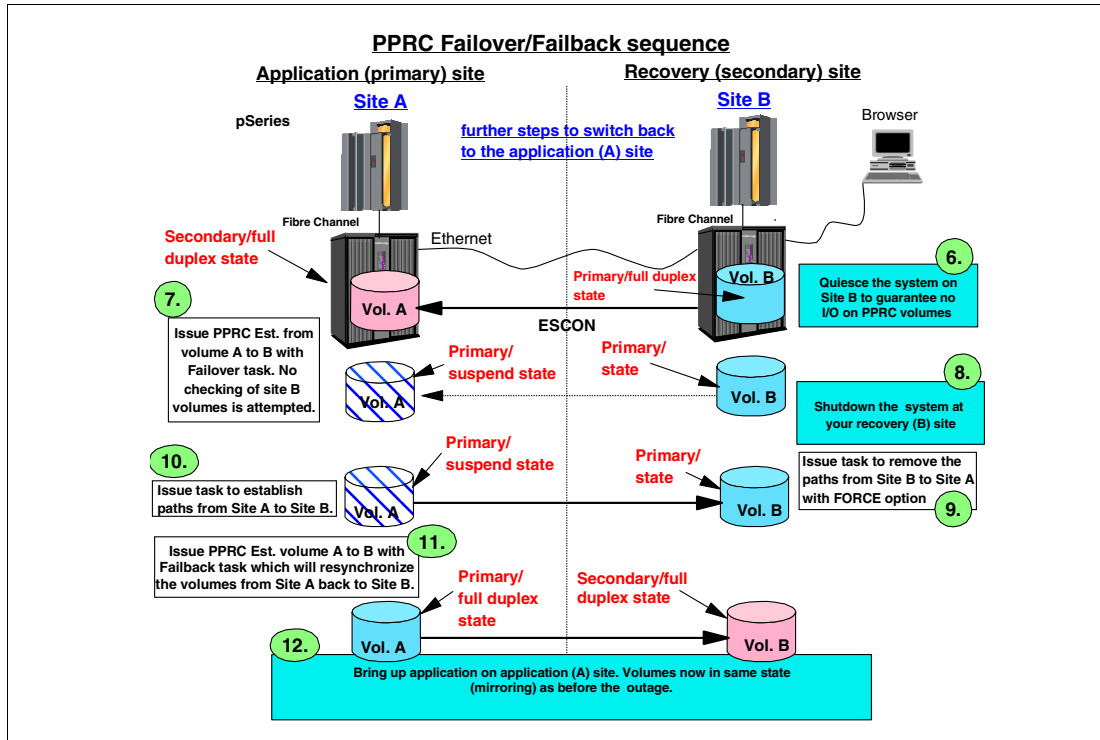
Figure 10-19   Resume application site

10.8 Failover/failback with Asynchronous Cascading PPRC

The exploitation of PPRC Failover and Failback modes helps reduce time required to synchronize PPRC volumes after switching between sites during a planned or unplanned outage.

In this section, we describe how the PPRC Failover and Failback modes can be used in an Asynchronous Cascading PPRC environment, like the one shown in Figure 4-8 on page 66. We show how PPRC Failover and Failback modes can be used to swap application I/O to site B after I/O to site A has ceased, for example, due to a planned outage for site A maintenance. The discussion in this section can be complemented with the example presented in 6.12.16, "PPRC Failover/Failback with Asynchronous Cascading PPRC" on page 283, which shows the actual Web user interface panels and options used to execute the failover and the failback.

Starting with LIC Version 2.3.0, PPRC Failover and Failback mode for Asynchronous Cascading PPRC is supported by the ESS, and can be managed using the ESS Copy Services Web user interface (WUI) as well as the ESS Copy Services Command Line Interface (CLI) via a saved task.

In the following scenario, we assume that all I/O to site A volumes has ceased due to a planned outage. The following steps describe the process of switching application I/O from site A to site B and then, when site A is available again, switching application I/O back from site B to site A while maintaining the PPRC-XD sessions between site B and site C.

10.8.1 Switching to site B

We assume that I/O to site A has ceased due to a planned outage, and in this section, we show how PPRC Failover can be used to switch to site B.

PPRC Failover to site B volumes

The first step in switching application I/O from site A to site B is to issue a PPRC Failover to site B. This can be done via the ESS Copy Services WUI. Examples of this procedure can be found in 6.12.13, "Performing a PPRC Failover" on page 266 and in 6.12.16, "PPRC Failover/Failback with Asynchronous Cascading PPRC" on page 283.

The ESS will perform the following when establishing PPRC from B to A with the **PPRC Failover** option specified:

1. Detects that the site B volume is a cascading volume.

2. Suspends the A → B session but does not reverse the direction.

3. Activates Change Recording (CR) for the site B volumes.

4. A new cascading volume flag, "Enable Secondary Host writes", is enabled.

This is different from PPRC Failover processing in a non-cascading environment, which would make B a primary in a B → A PPRC session. This could not be done for a cascading PPRC volume, as the B volume cannot be a primary for B → A and B → C at the same time. The B volume's Out-Of-Synch (OOS) information is required to maintain the B → C PPRC-XD session, and so Change Recording (CR) information is required to track updates to the B volumes for the imminent B → A PPRC session.

The state of the Asynchronous Cascading environment after the PPRC Failover processing has completed is shown in Figure 10-20.
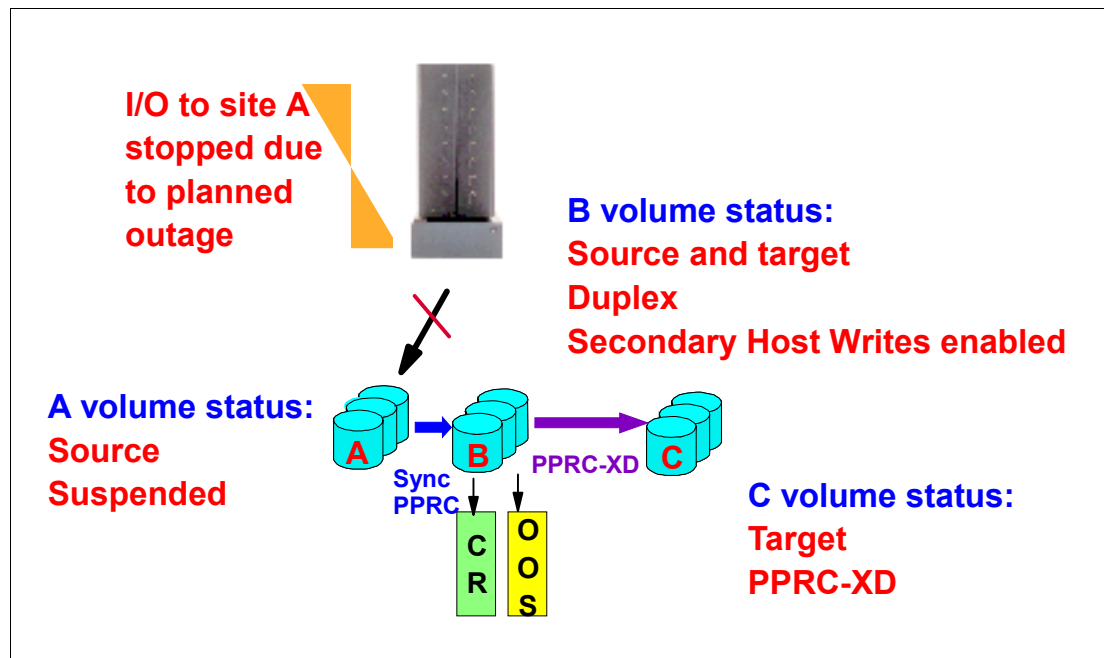
Figure 10-20   Volumes states after Failover command to site B volumes

Swap I/O to site B volumes

Host I/O can now be enabled to the site B volumes. The CR and OOS information is tracked for updates to site B volumes, as the PPRC-XD session B → C is still active. The go-to-SYNC operation can still be done to synchronize the site C volumes and FlashCopy can be used to create consistent copies of the site C volumes during this phase. This can continue until site A is ready to become the primary site again for application I/O. Figure 10-21 shows the status after application I/O has been switched to site B.
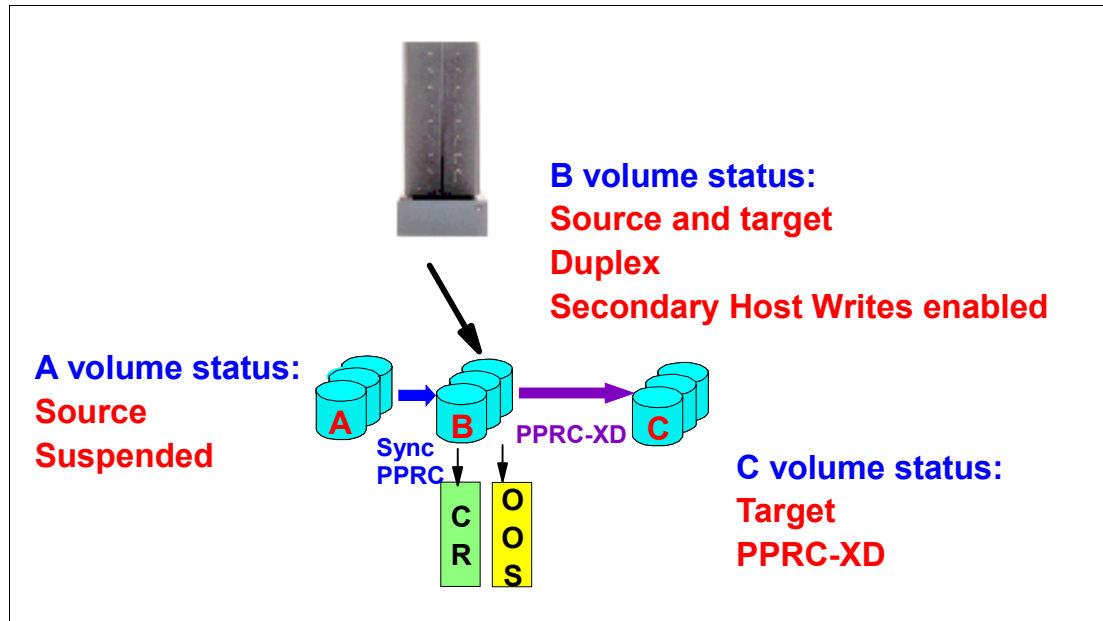
Figure 10-21   Swap I/O to site B volumes

10.8.2 Switching back to site A

We assume now that the planned outage for site A has completed and we are ready to switch application I/O back to site A. This section describes the series of step to achieve this.

Suspend the B -> C PPRC-XD session

After site A is repaired or ready again, the site A volumes need to be synchronized with the site B volumes before application I/O can be moved back to site A. In order to achieve this, a B → A PPRC session must be started. As the site B volumes cannot be a primary for both B → A and B → C at the same time, the B → C PPRC-XD session must be suspended.

Reverse paths B -> A and PPRC Failback to site B volumes

With the B → C sessions already suspended, the PPRC Failback from B to A can now be done:

1. Establish B → A PPRC paths.

2. Issue a PPRC Failback command to the site B volumes. PPRC Failback processing does the following:

   a. As the ESS is aware of the Asynchronous Cascading PPRC status, the OOS information will be used for the B → A PPRC session and the CR information will be used for the imminent B → C PPRC-XD session.

   b. Site A volumes are synchronized with site B.

c. Although B is no longer primary for C, the "trusted resync" flag is set to allow a later resynchronization.

Quiesce application I/O to site B

The site A volumes are now in full duplex with the site B volumes, so the site B volume's OOS information is no longer required. CR is still maintained to track updates to site B volumes, as the updates will have to be copied to the site C volumes when the B → C PPRC-XD session is re-established.

Application I/O to the site B volumes for all appropriate open system hosts can now be quiesced in preparation to switching back to site A. Figure 10-22 shows the status after the B → A PPRC session has been established and application I/O has been quiesced.
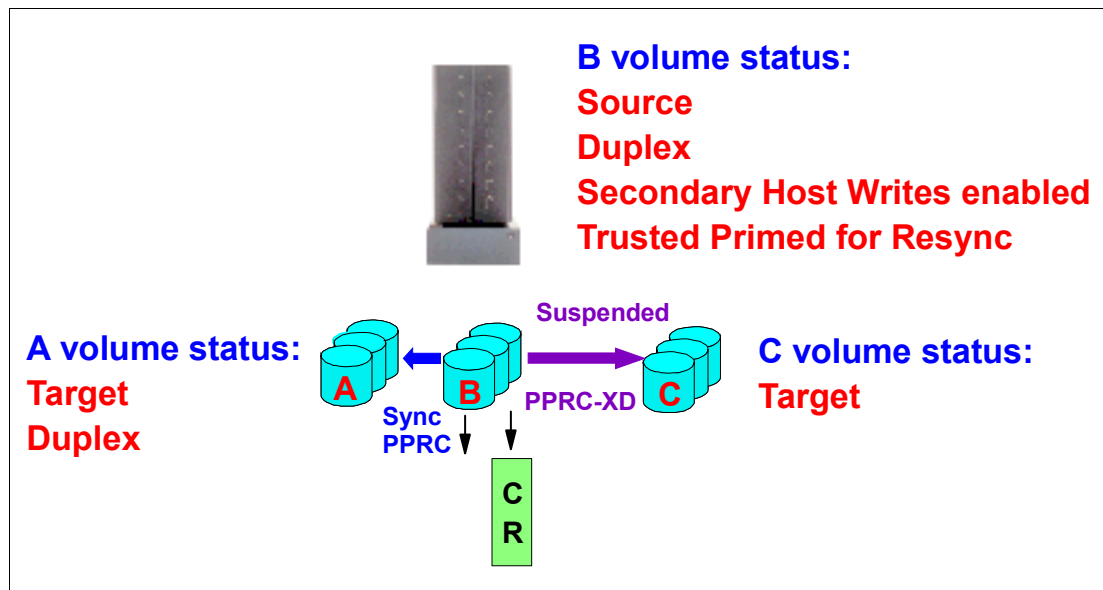
Figure 10-22   Volumes states after failback to site B volumes and quiesce application I/O

PPRC Failover to site A volumes

Now site A volumes need to become primary volumes again before resuming application I/O updates:

1. Establish the A → B PPRC paths.

2. Issue the PPRC Failover command to the site A volumes. This will:

   a. Terminate the B → A PPRC relationship.

   b. Establish the A → B PPRC relationship.

3. Issue a PPRC Failback command to the site A volumes. This will:

   a. Complete the restore of the A → B PPRC relationship. No resynchronization is necessary, as the volumes were in full duplex state.

   b. The cascading volume flag, "Enable Secondary Host writes", is reset, as host I/O will be directed to the site A volumes.

Figure 10-23 on page 433 shows the volumes states after the PPRC Failover to site A.
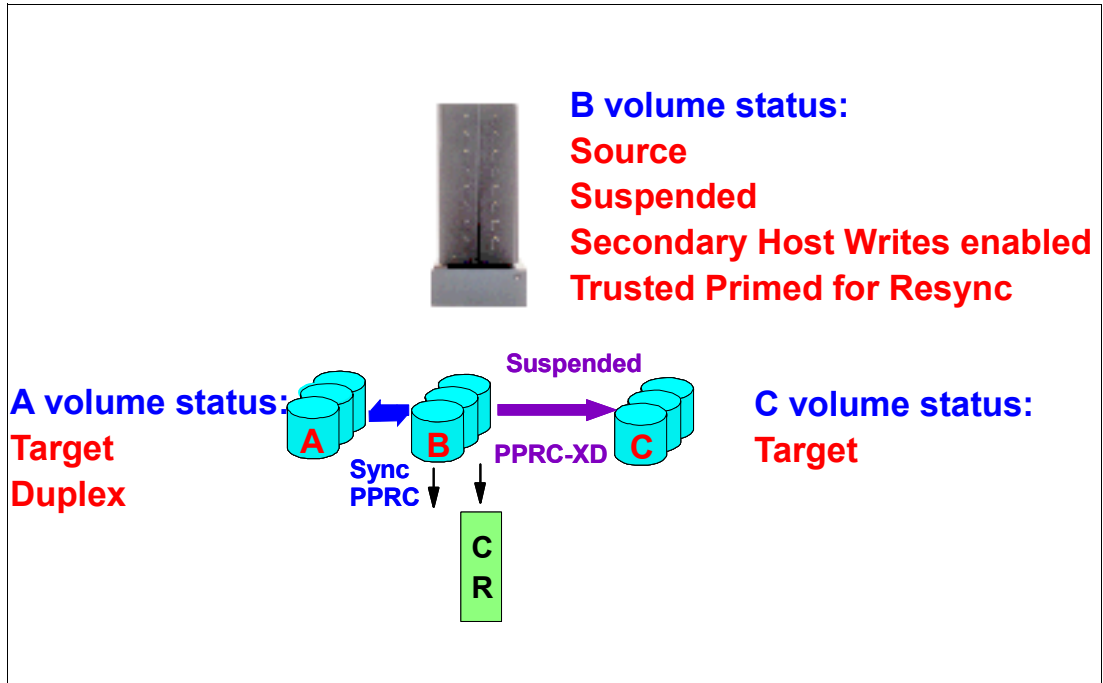
Figure 10-23   Volumes states after PPRC Failover command to site A

Figure 10-24 shows the volume states after the PPRC Failback to site A.
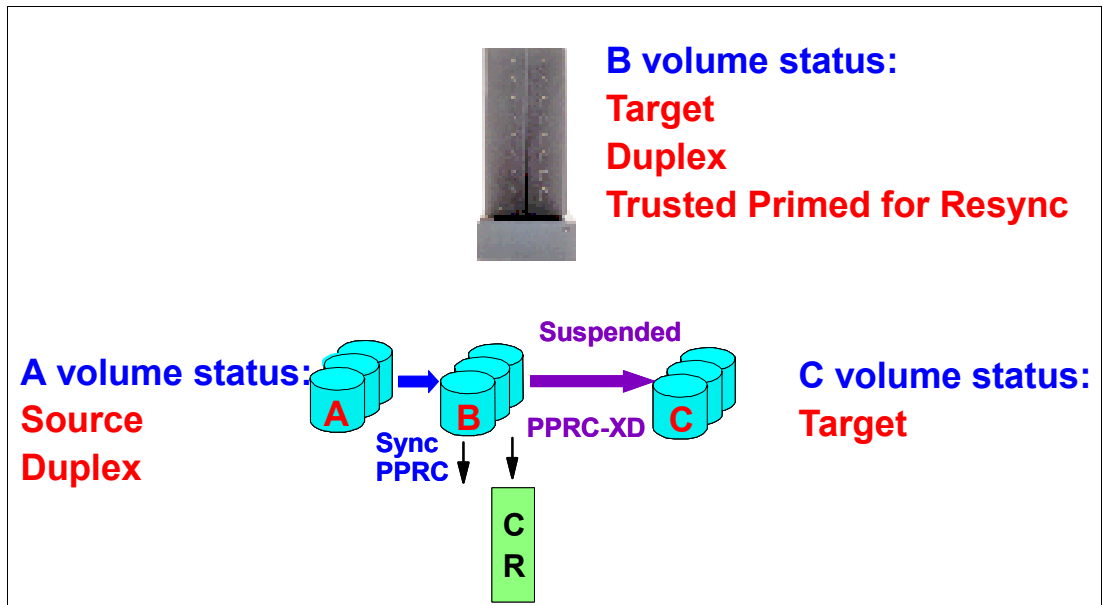
Figure 10-24   Volumes states after PPRC Failback command to site A

Return to full Asynchronous Cascading PPRC relationship

Now we can re-establish the normal asynchronous cascading PPRC environment, with site A, B, and C fully functional.

1. Restart application host I/O to the site A volumes. If changes have been made at site B, then the following list describes host considerations and checks you should make before restarting the applications at site A:

   – Check for additional volumes added at site B, to any applicable hosts, and establish the PPRC relationship between any new LUN additions and make the appropriate LUN to host assignments at site A.

   – Configure the host O/S to recognize the new LUNs (on AIX, rerun `cfgmgr`; on HP, rescan for the disks, and execute the host specific commands for each host and host type).

   – Follow the host specific procedure to bring volume, and disk groups back on line.

   – Before mounting the file systems, execute the `fsck` or I/O commands specific to the hosts.

2. Application writes are synchronously PPRC mirrored to site B. CR information has been maintained for the site B volumes.

3. Re-establish the B $\rightarrow$ C PPRC-XD pairs with "resync". The "resync" option is allowed even though the site B volumes are SIMPLEX to the site C volumes. This is because the "trusted resync" flag was set during the Failback processing. The flag is reset during this processing. The CR information is transformed to OOS information, which is used to re-establish the B $\rightarrow$ C PPRC-XD relationship.

Figure 10-25 shows the status after returning to full Asynchronous Cascading PPRC.
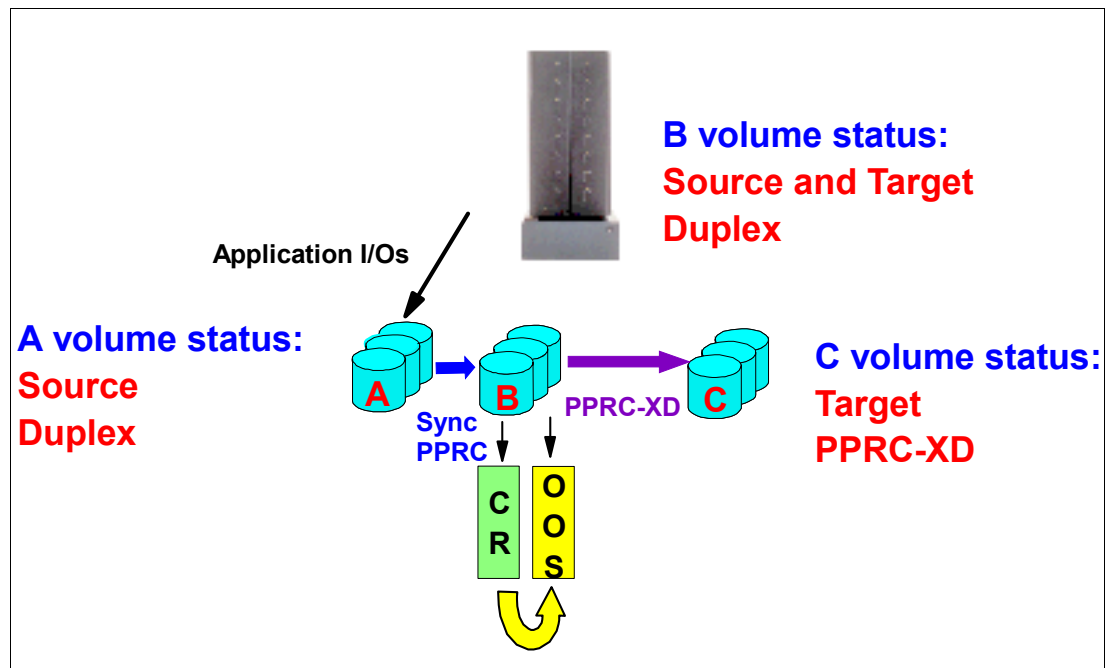
Figure 10-25   Return to full Asynchronous Cascading PPRC relationship

Normal operation for Asynchronous Cascading PPRC has now been re-established.

► Application I/O is directed to the site A volumes.

► Synchronous PPRC A $\rightarrow$ B has been re-established.

Non-synchronous PPRC-XD relationships B $\rightarrow$ C have been re-established. OOS information is maintained, which can be used to synchronize the B $\rightarrow$ C sessions.

10.9 Good practice guidelines for disaster recovery planning

In every case, disaster recovery planning is complex; nothing can understate the importance of rehearsals and testing your environment. You only get one shot at getting it right when a real disaster hits.

► Carefully set up your PPRC tasks. Create in advance all necessary tasks that will establish source-target relationships for paths, volumes, and LSSs (see Chapter 6, "ESS Copy Services Web User Interface: LIC 2.2.0. and 2.3.0" on page 169). For each task that establishes a source-target relationship, you should create a corresponding task that terminates the relationship.

► Each task type must be available for both directions: from production to recovery site and vice versa. This shows the importance of planning issues like naming conventions (Appendix E, "Sample naming convention" on page 499).

► Ensure that the tasks are well tested and documented.

► Periodically swap sites and run production operations at the recovery site to test the whole disaster recovery plan.

► Prepare your documentation as if it were intended for someone else; you may not be around when a disaster strikes.

► Be sure you understand any operating system specific issues related to bringing your PPRC secondaries online. This may include operations like removing old configuration data from the system, rescanning Fibre Channel devices, importing volume groups, and mounting them. These steps differ depending on your operating system and volume manager. Have them well documented in your recovery operations control book.

10.9.1 CLI scripts

For several reasons, it is highly recommended that you prepare all Copy Services tasks in advance and then run them by CLI scripts:

► Scripts can be reviewed and optimized in advance to avoid human errors.
► Scripts help to handle complexity and time pressure in a disaster situation.
► Script output and messages can be saved in log files, which are useful for later analysis.

Coding guidelines

There is a wide range of script languages available. Independent from the operating system and script language you choose, we recommend you apply some general guidelines:

► The management station to run the scripts must not be affected by the outage.

This means that you need at least two machines with all CLI software and scripts installed. Additionally, these machines should not rely on ESS volumes.

► You may not want to automate everything.

Disaster situations vary with respect to components affected and their state. You must not assume that situations evolve in a way that you expected. For example, a script may not be able to distinguish safely between a site failure and a total loss of intersite communication. Some critical decisions may need to be made by a human administrator.

► Requiring personal authentication with unique user name/password for each script reduces the risk of starting critical operations inadvertently, and it provides some minimum level of recording.

► Design your scripts robustly with structured error handling for problems that may arise during the fail-over process (for example, copy services server not reachable or some volume pair not in the required state for the next step). But if the error situation is not

absolutely clear, then let the script pause until an operator has diagnosed the problem. Wrong error recovery actions may put data consistency at risk.

► Grouping of tasks (in the Tasks panel of the Web User Interface) may not be appropriate for outage scenarios. If a single task fails, it is necessary to ungroup the overall group task in order to gain access and to re-execute the individual task.

Grouping of tasks does improve the performance, so you must weigh both sides of the argument for or against the grouping of tasks.

► In our discussions, we always assumed that production volumes as well as recovery volumes are assigned permanently to their hosts. You should consider assigning and un-assigning volumes strictly on a need-to-access basis (using the ESS Storage Management CLI in addition to the ESS Copy Services CLI).

If volumes are assigned permanently, then only the PPRC secondary volumes are protected against write operations, but the **PPRC Failover** option results in two primary volumes, both writable. A failed task may create simplex volumes. So in difficult situations, problems with operating system device names or unintended writes to the wrong volume might be fatal, slow down the recovery process, or even lead to loss of transactions.

Additionally, there are good coding practices for scripts published in general computer literature. Here we show some pieces of code for the Windows 2000 command line interface. But they can be adopted easily for other platforms.

Synchronous and asynchronous tasks

From a programmer's point of view, running an ESS Copy Services task via the ESS Copy Services CLI can be considered as a remote procedure call. The remote procedure is executed either synchronously or asynchronously:

► **Synchronous calls** give the flow control back to the calling program when the procedure's execution is completed. The return status code indicates if the whole operation's result was successful.

► **Asynchronous calls** give the flow control back when the procedure's execution has been initiated. The return status code indicates if the initiation was successful. It does not give any information about the operation's final result.

Activations of tasks for path management, terminating or suspending PPRC relationships, or establishing PPRC relationships with the option **Do not copy volume** are synchronous calls. Activations of tasks that possibly involve data transfer (establishing PPRC relationships with the options **Copy entire volume**, **Copy out-of-sync cylinders only**, **PPRC Failover**, or **PPRC Failback**) are asynchronous calls.

Therefore, the CLI command `rsExecuteTask` ("rsExecuteTask (rsExTask) command" on page 304) gives different results, depending on the task provided as the command parameter. In case of synchronous calls, the next script line is executed only when the task has completed its execution, and the status variable reflects the final result. But with asynchronous calls, you must not assume that the data transfer has been completed when the next line is executed.

The asynchronous nature of long-running tasks is useful to initiate multiple tasks in parallel, but it is dangerous to begin the next phase of a failover/failback procedure without checking for outstanding transfers. The `rsQueryComplete` command accepts task names and determines whether all volumes defined in the given tasks have completed the PPRC initial synchronization or FlashCopy background copy. Example 10-1 on page 437 shows a piece of code on how to use `rsExecuteTask` and `rsQueryComplete` in a Windows 2000 script.

Example 10-1   Using rsExecuteTask and rsQueryComplete

```
...
call :SubroutineExecuteTask "Task_1"
call :SubroutineExecuteTask "Task_2"
call :SubroutineExecuteTask "Task_3"
call :SubroutineQueryComplete "Task_1,Task_2,Task_3"
...

:SubroutineExecuteTask
if %DEBUG_TRACE% echo --- %~f0 [%0 %*]
setlocal enableextensions & pushd %ESS_CLIHOME%
  %DEBUG_SIMULATE% rsExecuteTask.exe /s %ESS_CSS% /u %ESS_USERNAME% /p %ESS_PASSWORD% %1
  if errorlevel 1 (echo *** Error %errorlevel% when starting task %1 & pause)
popd & endlocal
goto :EOF

:SubroutineQueryComplete
if %DEBUG_TRACE% echo --- %~f0 [%0 %*]
setlocal enableextensions & pushd %ESS_CLIHOME%
  %DEBUG_SIMULATE% rsQueryComplete.exe /s %ESS_CSS% /u %ESS_USERNAME% /p %ESS_PASSWORD% %1
  if errorlevel 1 (echo *** Error %errorlevel% when waiting for task %1 & pause)
popd & endlocal
goto :EOF
```

Note 1

Note 2

Note 3

Note 4

Note 5

Please note with respect to coding practice:

1. Task names are case-sensitive, so use quotes. In this examples, multiple tasks are started to run in parallel, and then we wait for the completion of all tasks.

2. The purpose of the global variable DEBUG_TRACE is to provide automatic trace functionality. If it has the value 1 before the subroutine is executed, then the script's filename (expression ~%f0), the subroutine name (expression %0) and all subroutine parameters (expression %*) are printed.

3. Both subroutines are parts of a larger package and may be called in different contexts, so no assumptions must be made about the current directory. The global variable ESS_CLIHOME contains the path name to the CLI program files. The line creates a local scope for variables as well as the current drive and directory (via the **setlocal** and **pushd** commands). After the subroutine's body, a corresponding set of commands, **endlocal** and **popd**, close the local scope. This makes the subroutine "well behaved" and preserves the current drive and directory, as well as all environment variables. It also means that code within the subroutine is free to alter any variable, as any changes made are automatically restored when the procedure completes.

4. The global variables ESS_CSS, ESS_USERNAME, and ESS_PASSWORD provide the address, user name, and password for the active ESS Copy Services server. Connectivity to that server should be tested with **rsTestConnection** in advance. The parameter %1 is the name of the task (or list of task names) to be executed in SubroutineExecuteTask and to be monitored in SubroutineQueryComplete.

   The DEBUG_SIMULATE variable can be defined as **echo** when you want to test the script's logic without doing something at the ESS.

5. The subroutines are paused in case of errors to allow further investigation and correction (like restarting the task manually) before the script continues.

In this example, SubroutineExecuteTask returns immediately, but SubroutineQueryComplete returns only when all tasks are completed.

10.9.2 Script modules

Reviewing the scenarios, you recognize that the number of scripts necessary to handle at least typical situations makes maintenance difficult. Adding new volumes or host systems requires modifications in several scripts, which will lead to errors. There are two techniques you should consider for your scripts:

► Separation of script logic and configuration information
► Script libraries

Separation of script logic and storage configuration

The programming technique to resolve issues like adding new volumes is the separation of the program's logic from the objects to which this logic is applied. You should use a configuration file describing all volumes and their PPRC relationships (see Example 10-2). Then all of your scripts parse this configuration file and apply some operation to all volumes

Example 10-2   Configuration file

```
...
#Keyword    Address
CSS_Primary=10.160.74.185
CSS_Backup=10.160.74.85

# Keyword Datacenter   Datacenter  ESS       ESS
#         longname     sortname    name      serial
PPRC_Node=PRODUCTION   P           CO        23052
PPRC_Node=RECOVERY     R           CO        23165

# Keyword Production   Recovery    Host      Drive
PPRC_Pair=10023052     10023165    FP3_DB    E    # DB exe
PPRC_Pair=20023052     20023165    FP3_DB    F    # DB log 1
PPRC_Pair=30023052     30023165    FP3_DB    D    # Pagefile
PPRC_Pair=40223052     40223165    FP3_DB    C    # Boot
PPRC_Pair=50023052     50023165    FP3_DB    G    # DB log 2
PPRC_Pair=60023052     60023165    FP3_DB    H    # DB tables 1
PPRC_Pair=40023052     40023165    FP3_DB    L    # DB tables 2
PPRC_Pair=30123052     30123165    FP3_DB    N    # DB archive
...
```

The file describes several elements of the configurations with the same syntax:

► Comments (from # to end of line) and spaces can be added anywhere.

► Each line begins with a keyword (terminated by =), followed by parameters (separated by blanks) in the order defined for that keyword.

In our example, first the Copy Services server addresses are given. Then the lines with keyword PPRC_Node provide for each ESS:

► A datacenter longname (used in display text)
► A datacenter shortname (used in task names)
► A mnemonic name for the ESS pair to which this ESS belongs
► The ESS serial number

Similarly, the keyword PPRC_Pair indicates that the line describes a PPRC volume pair. Each PPRC pair is defined by:

► The ESS volume serial number in the production site
► The peer's volume serial number in the recovery site
► The host or cluster name abbreviation

► The drive name used by the host (in our case, this is the Windows drive letter)

The configuration file, in conjunction with script arguments, allows you to control script actions in a very flexible manner. Example 10-3 on page 439 is a piece of code from a Windows 2000 script ESS_Show.cmd, which shows the state of volumes (using the **rsQuery** command) for a subset of PPRC volumes. For simplicity, we have omitted code for routine actions like providing help text, checking input parameters, setting default values of optional parameters, or testing global variables.

Example 10-3   Using the rsQuery command with configuration file

```
if %DEBUG_TRACE% echo --- %~f0 %*
setlocal enableextensions
set DATACENTER=%1
set ESSNAME=%2                                                          Note 1
set HOST=%3
set DRIVE=%4
...
for /f "eol=# tokens=1-5 delims== " %%I in (%ESS_CONFIG%) do (          Note 2
  if /i %%I==CSS_Primary set css1=%%J
  if /i %%I==CSS_Backup  set css2=%%J
  if /i %%I==PPRC_Node (if /i %%J%%L==%DATACENTER%%ESSNAME% (
                            set ESSSERIAL=%%M
                            goto :CSS_Found)
                       )
)
:CSS_Found                                                             Note 3
...
for /f "eol=# tokens=1-5 delims== " %%I in (%ESS_CONFIG%) do (          Note 4
  if /i %%I==PPRC_Pair (call :SubroutineQuery %%J %%L %%M
                        call :SubroutineQuery %%K %%L %%M)
)
...
endlocal
goto :EOF


:SubroutineQuery
if %DEBUG_TRACE% echo --- %~f0 [%0 %*]
setlocal enableextensions
set VOLUME=%1
if not %VOLUME:~-5%==%ESSSERIAL% goto :EOF
if not %HOST%==* (if /i not %HOST%==%2  goto :EOF)                      Note 5
if not %DRIVE%==* (if /i not %DRIVE%==%3 goto :EOF)
pushd %ESS_CLIHOME%
  %DEBUG_SIMULATE% rsQuery.exe /s %ESS_CSS% /u %ESS_USERNAME% /p %ESS_PASSWORD% /q %VOLUME%
  if errorlevel 1 (echo *** Error %errorlevel% when querying volume %VOLUME% & pause)
endlocal & popd
goto :EOF
```

Notes regarding Example 10-3 on page 439:

1. The two mandatory parameters of the script specify a datacenter longname and ESS mnemonic name from a PPRC_Node definition in the configuration file. Optional parameters are a host name and drive letter according to a PPRC_Pair definition. The script displays the state of all volumes matching this node/pair specification.

   For example, with our configuration file from Example 10-2 on page 438, running the script with the parameters PRODUCTION CO * * shows all volumes in the production ESS of the ESS pair CO. The parameters PRODUCTION CO FP3_DB * restrict the output to volumes

belonging to host FP3_DB. Otherwise, PRODUCTION CO * C takes volumes with drive letter C belonging to any host defined for this ESS.

2. The first `for /f` loop scans the configuration file (path in variable ESS_CONFIG) to find the matching PPRC_Node definition. This parse statement recognizes comments (from # to end of line) and puts up to five tokens (separated by = or space) into the additional variables I to M. Then the first token I is checked as a keyword:

   – If the keyword is CSS_Primary or CSS_Backup, then this Copy Services server definition is saved as valid for all PPRC nodes below until the next server definition is found.

   – If the keyword is PPRC_Node, then this definition is checked for a match with datacenter and ESS name. If we have found the right ESS, then the ESS serial number is saved.

3. Therefore, the code part :CSS_Found is reached either when we have found the required ESS or when reaching the end of the configuration file. Here we have omitted the code for handling the end-of-file case and testing the connections to the Copy Services servers.

4. The second `for /f` loop scans the configuration file for all PPRC_Pair definitions. For every pair found, :SubroutineQuery is called for the first as well as for the second volume. The subroutine will check if this volume matches the script's input parameter (because string manipulation with `for /f` loop variables is easier at the lower call level).

5. The subroutine checks if the combination of volume serial number, host name, and drive letter match the input parameters (note the wildcard handling). If yes, then the **rsQuery** command is issued.

In large environments, the configuration file could include several ESS Copy Services domains and contain additional sections about path design. Depending on ESS roles and naming conventions in your environment, other styles of wildcard support in scripts may be useful. However, the idea is that any change of volume numbers requires only modifying one single file describing all volumes, and changes in the operation's logic affect only one script file.

Script libraries

Example 10-3 on page 439 about querying volume states by a generic ESS_Show.cmd script leads to the approach that every copy operation can be mapped to a generic ESS_Copy.cmd script too. Using a uniform task naming structure as outlined in Appendix E, "Sample naming convention" on page 499; the task names can be generated by this script from the specification of volume, operation type, and operation direction. Example 10-4 shows the interface for such a script.

Example 10-4   Generic copy script ESS_Copy.cmd

```
ESS_Copy.cmd <op> <dc1> <ess1> <dc2> <ess2> <host> <drive>

Parameters:  <op>     Operation code according to task naming rules.
                      A preceding "W" indicates that the procedure waits
                      for completion of this task instead of starting the task.
             <dc1>    Primary datacenter
             <ess1>   Primary ESS
             <dc2>    Secondary datacenter
             <ess2>   Secondary ESS
             <host>   Host or cluster name
             <drive>  (optional) drive letter, default = *
```

This script uses the coding techniques shown in Example 10-1 on page 437 and Example 10-3 on page 439 to parse the configuration file, build the task name according to

the naming rules, and then run this task (or wait for its completion). Looking back to the fail-over/fail-back procedures discussed in 10.6, "Planned outages" on page 408 and 10.7, "Unplanned outages" on page 422, you see that every step of such procedure can be implemented by a single call of ESS_Copy.cmd, processing all volumes belonging to that host.

This leads to a two-level modular script library to handle fail-over/fail-back procedures:

- ► The lower level of the library consists of single-operation scripts like ESS_Show.cmd, ESS_Copy.cmd, and similar scripts for path handling. They apply a simple operation to all volumes matching a wildcard pattern. Usually, this pattern specifies all volumes belonging to a host/cluster (if the unit of failover is the host) or to a whole ESS (if the unit of failover is the ESS).

- ► The upper level of the library consists of scripts implementing the logic of a scenario like planned failover, unplanned failover, and failback after planned or unplanned failover. These scenario scripts call the lower-level single-operation scripts to apply every step of the procedure to all affected volumes.

Example 10-5 is a piece of code from a Windows 2000 script for planned failover with ESS LIC Version 2.1.x or earlier. It implements the procedure from 10.6.1, "Planned failover without PPRC Failover/Failback options" on page 408, using an ESS_Copy.cmd script like Example 10-4 on page 440, with a configuration file like Example 10-2 on page 438. The operation codes are explained in Appendix E, "Sample naming convention" on page 499.

Example 10-5   Scenario script for planned failover

```
...
echo --- Planned failover (1): Terminate old PPRC ...
call ESS_Copy.cmd PS %DC1% %ESS1% %DC2% %ESS2% %HOST%

echo --- Planned failover (2): Switch paths ...
call ESS_Paths.cmd %DC1% %ESS1% %OLD_PATHS% %NEW_PATHS% %DC2% %ESS2%

echo --- Planned failover (3): Prepare PPRC in opposite direction ...
call ESS_Copy.cmd PN %DC2% %ESS2% %DC1% %ESS1% %HOST%
call ESS_Copy.cmd WPN %DC2% %ESS2% %DC1% %ESS1% %HOST%

echo --- Planned failover (4): Suspend new PPRC ...
call ESS_Copy.cmd PH %DC2% %ESS2% %DC1% %ESS1% %HOST%
...
```

For a detailed discussion, see 10.6.1, "Planned failover without PPRC Failover/Failback options" on page 408. Please note that the script does not wait for host operations like dismounting volumes (it assumes that these steps are already completed). Such delay could be inserted at the beginning by waiting for some operator confirmation. Path handling in step (2) depends on your path design. In step (3), establishing the new PPRC pairs with the **Do not copy volume** option might be done without synchronization. We recommend waiting for all PPRC-establishing operations, regardless of the expected amount of data transfer.

In this way, new scenarios can be implemented easily. Changes of procedures are restricted mostly to one scenario script (adding or modifying calls to single-operation scripts). The single-operation scripts are useful by themselves as troubleshooting tools. All scripts are independent from volume changes.

HACMP Remote Copy and ESS PPRC

High Availability Cluster Multi-Processing (HACMP) Remote Copy is a component of IBM HACMP/XD (extended distance). HACMP Remote Copy increases data availability for IBM TotalStorage Enterprise Storage Server (ESS) volumes that use Peer-to-Peer Remote Copy (PPRC) to copy data to a remote site for disaster recovery purposes. HACMP Remote Copy takes advantage of the PPRC Failover/Failback functions and HACMP cluster management to reduce downtime and recovery time during disaster recovery.

The material in this chapter comes from the *HACMP Remote Copy: ESS PPRC Guide*, SC23-4863.

This chapter provides an overview of HACMP Remote Copy. The chapter contains the following sections:

► HACMP Remote Copy and PPRC

► HACMP Remote Copy in an HACMP Cluster

► Sample Configuration

11.1 HACMP Remote Copy and PPRC

HACMP Remote Copy lets you extend HACMP cluster management for highly available applications and servers to support the disaster recovery mechanism supplied by PPRC. The HACMP Remote Copy integration with PPRC increases data availability by providing continuing service during hardware and/or software failures for ESS systems that are in different locations.

PPRC allows mirroring to be suspended and restarted without affecting data integrity.

11.1.1 PPRC mirroring

When PPRC is activated, it establishes a synchronous mirror between a specified volume on one ESS, and a specified volume on another ESS. The PPRC mirrored volumes are referred to as a PPRC pair. In a PPRC pair, the volume that data is being written to is the source volume, and the volume that contains a mirrored copy of the data is the target volume. The paths over ESCON links that connect these pairs are called PPRC paths.

PPRC mirrors at the disk subsystem level, making it transparent to hosts. Applications access mirrored volumes as they do any volume.

11.1.2 HACMP Remote Copy integration with PPRC

HACMP Remote Copy takes advantage of the fail-over and fail-back functions of PPRC to provide automatic fallover and recovery of a pair of ESS volumes. The PPRC Failover and fallback features help to reduce the downtime and the recovery time during disaster recovery.

> **Attention:** The ESS Copy Services Web Interface refers to a fallover from a source to a target volume as a *failover*, and a fallback (a resynchronization) as a *failback*.

HACMP Remote Copy provides high availability and disaster recovery through:

► Automatic fallover of PPRC-protected volume pairs between nodes within a site
► Automatic fallover of PPRC-protected volume pairs between sites
► Automatic recovery/reintegration of PPRC-protected volume pairs between sites
► Support for cascading and rotating resource group management policies
► Support for cascading, ignore, and rotating intersite management policies
► Support for VERBOSE_LOGGING and EVENT_EMULATION
► Support for the Subsystem Device Driver (SDD)
► Support for cluster verification and synchronization
► Support for C-SPOC

HACMP Remote Copy relies on PPRC tasks for managing PPRC volume pairs and the associated PPRC paths. Tasks are a set of actions to be executed on an ESS. They let you automate a series of steps that otherwise would be performed in the ESS Copy Services Web Interface. To access these tasks, nodes in an HACMP cluster that supports HACMP Remote Copy require access to ESSNet to execute commands.

11.2 HACMP Remote Copy in an HACMP Cluster

In a single-site HACMP environment, all cluster nodes sharing volume groups have physical connections to the same set of disks. In an HACMP Remote Copy environment, the cluster nodes access the same shared volume groups, but the nodes at each site access them from different physical volumes, which are the two volumes in a single PPRC pair.

HACMP Remote Copy lets you include PPRC-mirrored volumes in an HACMP cluster. This requires two HACMP sites. Each site contains one ESS and the nodes attached to it. PPRC pairs are defined so that one volume in the pair resides on the ESS at one site, and the other volume in the pair resides on the ESS at the other site.

A PPRC pair associated with an HACMP site definition and defined to HACMP is referred to as a PPRC replicated resource. The definition for a PPRC replicated resource contains the volume identifier and the name of the ESS. HACMP recognizes which volumes mirror each other for each PPRC replicated resource.

> **Note:** PPRC copies the volume information, including the PVID, from one volume in a PPRC pair to the other. The volumes at both sites contain the same logical volumes, and must therefore be imported with the same volume group name. This also allows single-name entries in a resource group definition.

11.2.1 Resource groups that include PPRC replicated resources

With HACMP Remote Copy, a resource group:

- ► Contains a shared volume group and the PPRC replicated resource(s) associated with the individual volumes in the volume group
- ► Contains nodes that all have access to the ESSNet for each ESS system
- ► Supports cascading or rotating resource group management policies
- ► Requires an intersite management policy to handle a resource group during site recovery.

The policies supported are cascading, rotating, or ignore. (HACMP Remote Copy does not support concurrent intersite management.)

The policies are conceptually the same as the HACMP resource group management policies, but they relate to the movement of the resource group between sites.

11.2.2 HACMP sites for HACMP Remote Copy

HACMP Remote Copy requires two HACMP sites for use within a resource group to control which volume in a PPRC pair a node can access. Although nodes at both sites can access a volume group, PPRC allows access to only one volume (the source volume) in a PPRC pair at a time. This prevents nodes at different sites from accessing the same volume group at the same time. Typically, a number of volumes are mirrored via PPRC from the one site to the other.

HACMP Remote Copy handles the automation of fallover from one site to the other should a failure occur. This minimizes the recovery time after an outage or a disaster. When a site fails, the resource group configuration determines access to disk resources.

Within a resource group, the nodes at one site may handle the PPRC replicated resource differently than the nodes at the other site, especially in cases where the states (suspended or full-duplex) of the volumes are different at the two sites.

11.2.3 Primary and secondary volumes

The terms primary and secondary are specific to each pair of PPRC-mirrored volumes as determined by the pair definition. In configurations that designate one site as a production site and the other as a backup site, the production site holds the primary volumes for the PPRC pairs, and the backup site holds the secondary volumes. In a mutual recovery configuration, in which nodes from both sites are active, each site contains primary volumes for some PPRC pairs and secondary volumes for others. Each PPRC pair is included in a resource group, whose home node is located at the site that is primary for the PPRC pair.

11.2.4 Failover and fallback

HACMP Remote Copy and PPRC manage how a volume falls over from one volume to another. Both volumes are in a full copy state, with the source volume varied on.

If nodes at one site fail (for example, if they go offline), the nodes at the surviving site send commands to ESS Copy Services to execute a PPRC Failover task. This lets the nodes that are taking over to access the target volume, and puts the PPRC mirrored pair into a suspended state. When the nodes at the failed site recover, the nodes send commands to ESS Copy Services to execute a PPRC Failback task, which resynchronizes the volumes.

The following process summarizes how ESS volumes fall over should a volume become unavailable, then fall back when the failed volume is available again. In this example, ESS_A is the ESS at Site A that initially holds the source volume. ESS_B is the ESS at Site B which initially holds the target volume, and to which ESS_A falls over. If ESS_A goes offline, the following occurs:

1. ESS_A is suspended.

2. The nodes at Site B execute a PPRC task (Establish Pair - PPRC Failover).

3. Applications then write to ESS_B, with PPRC keeping track of the new writes. No mirroring takes place.

4. ESS_A goes back online:

   – ESS_A releases any stale reservations it still holds.

   – The nodes at Site B initiate a PPRC task (Establish Pair - PPRC Failback) to generate a data synchronization between the volumes by PPRC.

5. When the task in Step 3 completes, the source volumes are on ESS_B, and the target volumes are on ESS_A.

> **Attention:** Fall-back behavior is different for resource groups and for PPRC-protected volumes. Fallback for a resource group indicates that ownership of the group falls back to another node. Fallback for PPRC pairs indicates that suspended pairs are resynchronized with the original target volume becoming the new source volume.

11.3 Sample configuration

HACMP Remote Copy Version 5.1 supports a PPRC configuration that has two ESS systems with mirrored volumes connected by ESCON connections.

Figure 11-1 on page 447 shows a sample two-site configuration using PPRC between the two ESS systems and HACMP with HACMP Remote Copy for cluster management. This example shows a single production site and a single recovery site and assumes that the all primary volumes are on the production site.

In Figure 11-1:

► The HACMP Production Site includes:

   – Server A and Server B

   – The ESS labeled Primary ESS

► The HACMP Recovery Site includes:

   – Server C and Server D

   – The ESS labeled Secondary ESS

► The HACMP resource group contains:

   – The four server nodes connected over serial and IP networks

   – One or more shared volume groups

   – The PPRC replicated resource associated with the volumes in the volume group

Note: All nodes also need to be able to connect to the ESSNet for each ESS.
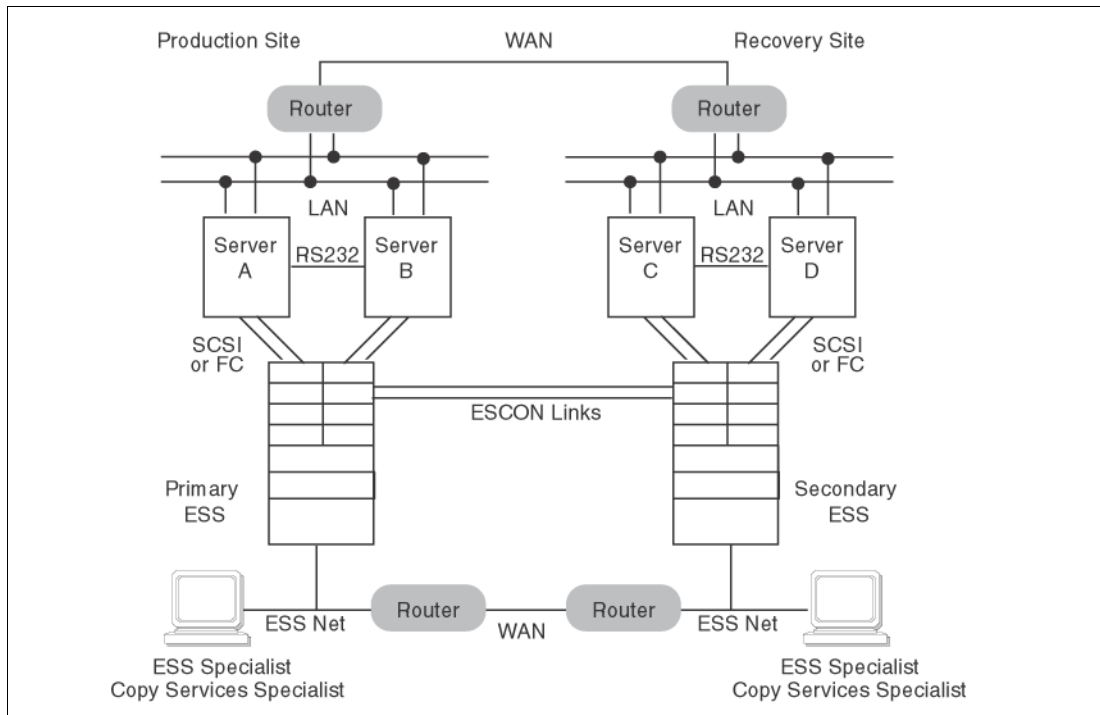
Figure 11-1   Sample PPRC configuration with HACMP Remote Copy

12

IBM Systems Group Service Offerings

This chapter provides additional information on using the ESS Copy Services functions in the IBM @server xSeries® environment. We will discuss a new IBM Systems Group Service offering, IBM TotalStorage Support for Geographically Dispersed Sites for Microsoft Cluster Service (GDS for MSCS) that uses the ESS's PPRC capability to allow Windows clustering across metropolitan distances, as well as other storage solutions.

In this chapter, the following topics are discussed:

► IBM TotalStorage Support for Geographically Dispersed Sites for Microsoft Cluster Services (GDS for MSCS)

► IBM TotalStorage Data Manager for SQL Server 2000 (TDM for SQL 2000)

► Other Windows offerings, including IBM TotalStorage ESS - Exchange Integration Resource Pack, IBM TotalStorage ESSSync, and IBM TotalStorage Volume Set Manager.

12.1 GDS for MSCS

IBM TotalStorage Support for Geographically Dispersed Sites for Microsoft Cluster Services integrates Microsoft Cluster Service (GDS for MSCS) and the Peer-to-Peer Remote Copy (PPRC) advanced function of the IBM TotalStorage Enterprise Storage Server (ESS). GDS is implemented using two sites (one primary, one backup), with each site having one cluster node. Should the primary site suffer a failure, the data and applications automatically fail over to the backup site, typically in minutes. Once the primary site is repaired, the resources can be failed back to the original site.

GDS offers high availability for servers and applications plus disaster recovery for files and data. It improves data availability with continued service during hardware and software failures. It provides disaster-tolerant capabilities by allowing the cluster servers and associated mirrored storage to be geographically separated by distances up to 103 km. It can also be used to reduce or eliminate planned downtime, such as during a software or hardware upgrade. We can see the topology in Figure 12-1.

The ESSs are connected to each other using PPRC links, and the host servers are fibre attached to the ESS at each site. The cluster shared disks are mirrored between two ESSs. The solution uses PPRC to coordinate the ownership and movement of disk resources between sites.
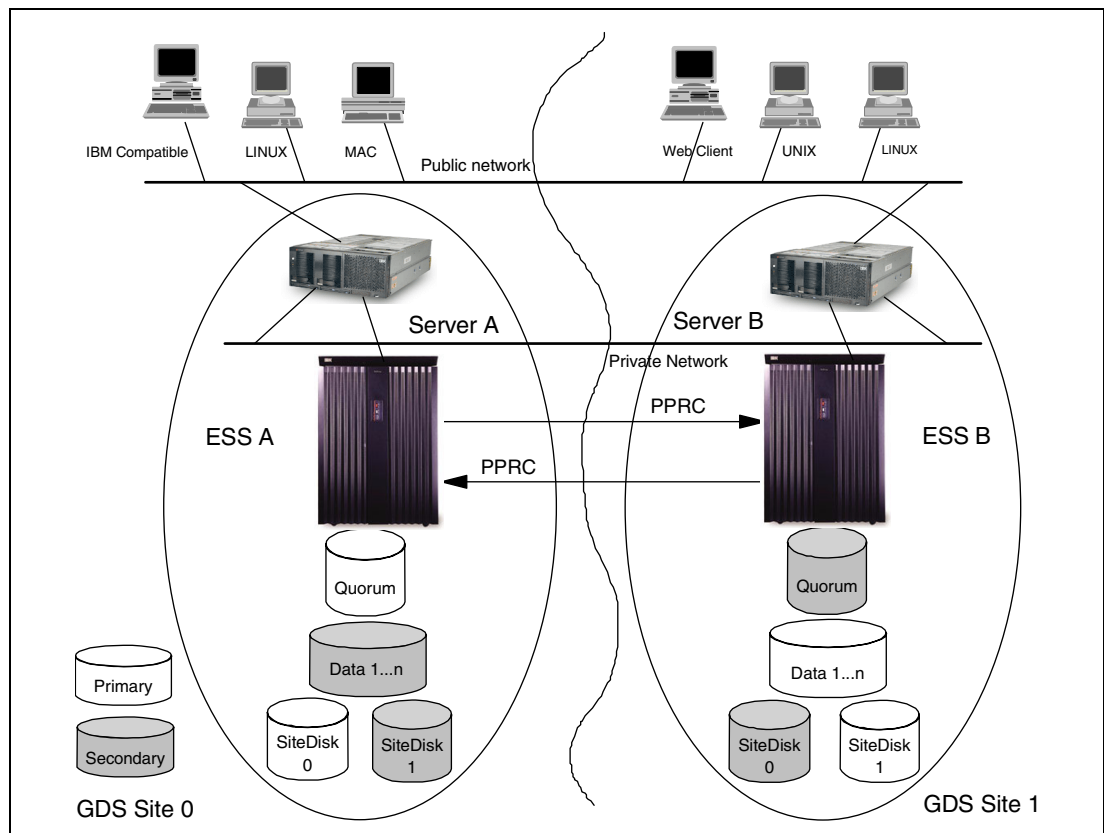
Figure 12-1   GDS topology

12.1.1 Software modules

There are four software modules making up the GDS software package. All four must be installed on each node in the cluster. They are:

▶ IBM PPRC Cluster Resource DLL and extension

GDS add a new resource type into Microsoft Cluster, which is called an IBM_PPRC cluster resource. Each non-quorum physical disk resource in the MSCS depends on an IBM_PPRC resource. This software module manages the interactions between MSCS and the PPRC resource and monitors the status of these resources. If one node is offline, the PPRC cluster resource fails over the physical disk resources from the offline node site to the online node site by using fail-over scripts that reverse the PPRC direction.

▶ IBM Subsystem Device Driver (SDD) Server Service

The SDD Server Service is used to communicate with the ESS subsystem. It is used to retrieve the state of the PPRC volumes. This is installed at the same time as the SDD bus driver code.

▶ IBM GDS Service

The quorum disk is handled by the IBM GDS Service, since it cannot depend on any resource in a fail-over group. This service's core functionality is:

– Monitors the state of the cluster (quorum, group, and membership), using MSCS cluster notification events

– Monitors the health of the PPRC link

– Controls the PPRC heart beat process

– Manages the quorum failover

On each cluster node, the fail-over service will monitor the status of the networks, the peer nodes, and the quorum resource. If this service determines that all the nodes at the other site are unavailable, then it makes the quorum disk available. MSCS event monitoring will identify moves involving the quorum disk.

ESS PPRC handles issues of data consistency. See Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53.

▶ GDS Configuration Utility

– Used to update GDS settings after installation

– Manages the settings for the quorum resource

12.1.2 GDS Demo

In this section, we have an example of how GDS works. Two servers, RootBeer and MountainDew, form a GDS. They are both connected to the public network and linked directly to each other via a private network. The two servers are connected to two different ESSs by host bus adapters (HBA) and the ESSs use PPRC for remote mirroring of the data over PPRC links between them. The two sites can be up to 103 km apart.

Disk E: is the quorum resource and there are two data disks (F: & G:) which are mirrored on the remote ESS using PPRC. The topology is displayed in Figure 12-2 on page 452.
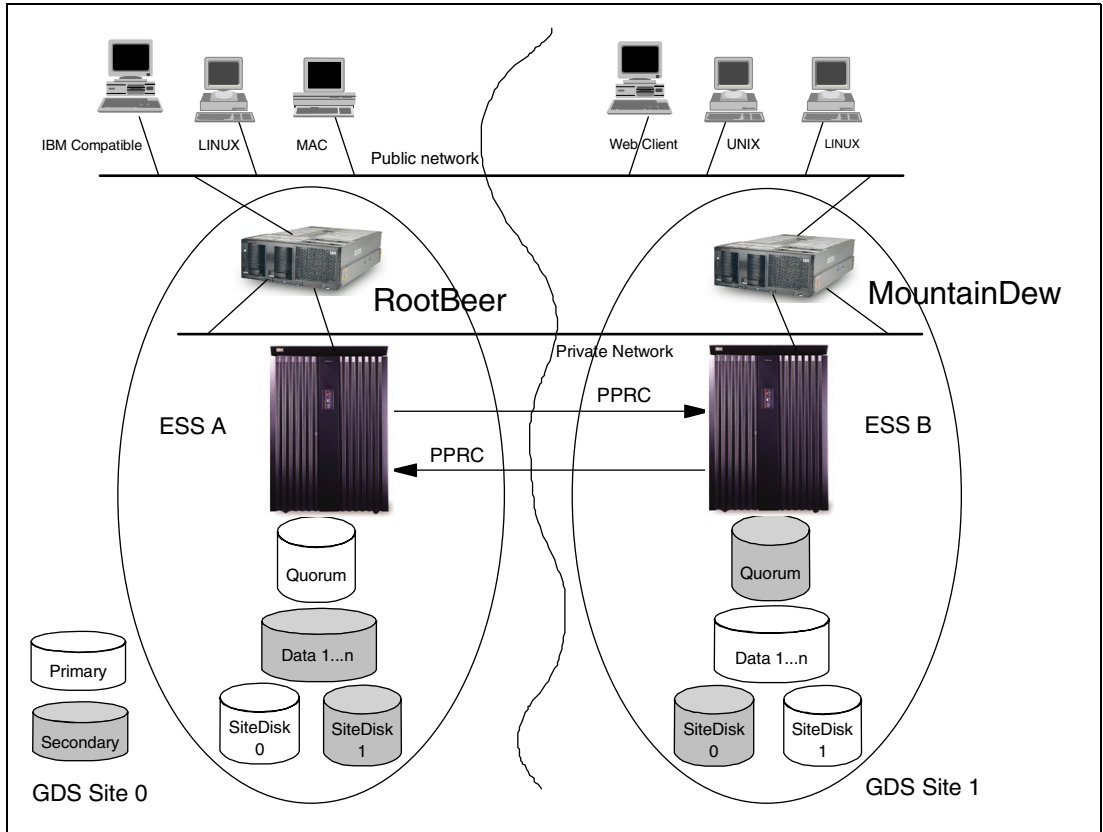
Figure 12-2   GDS example

1. RootBeer is the node that owns all the cluster resources as demonstrated by Figure 12-3, Figure 12-4 on page 453, and Figure 12-5 on page 453.
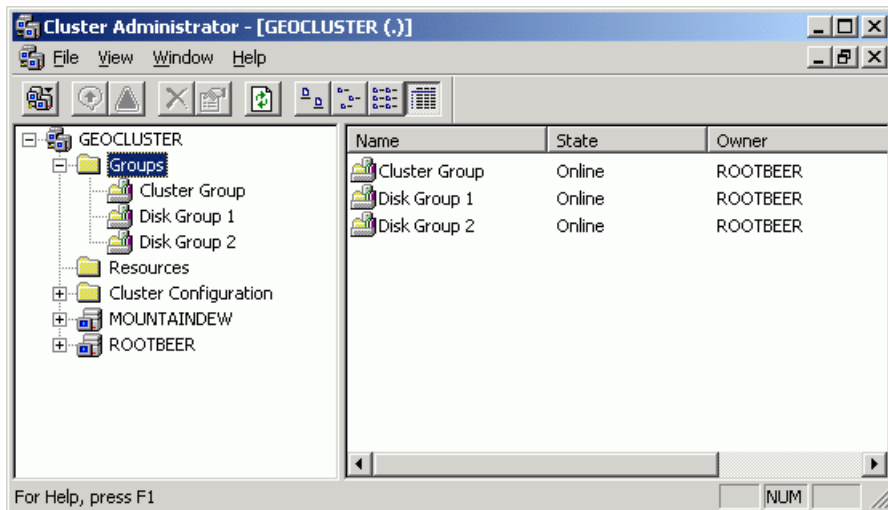
Figure 12-3   All three groups are owned by RootBeer

We can see that the cluster quorum resource is now owned by the RootBeer node (Figure 12-4 on page 453).
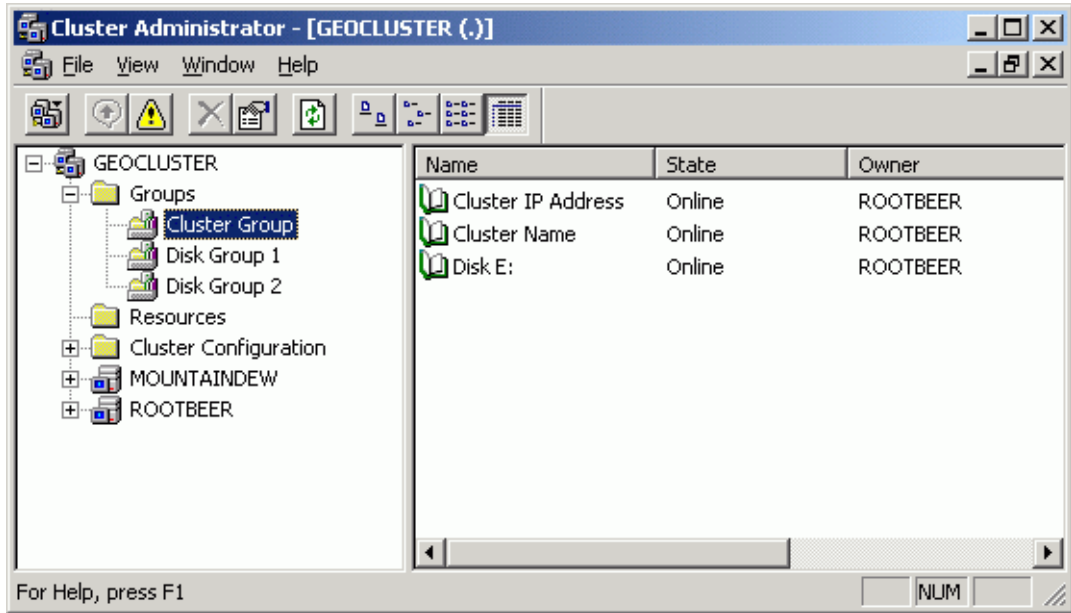
Figure 12-4   Quorum resource (Disk E:) is on RootBeer

All of the disks of the cluster are visible to RootBeer. MountainDew cannot see any shared disks. This is the same behavior as a typical shared-disk MSCS cluster.
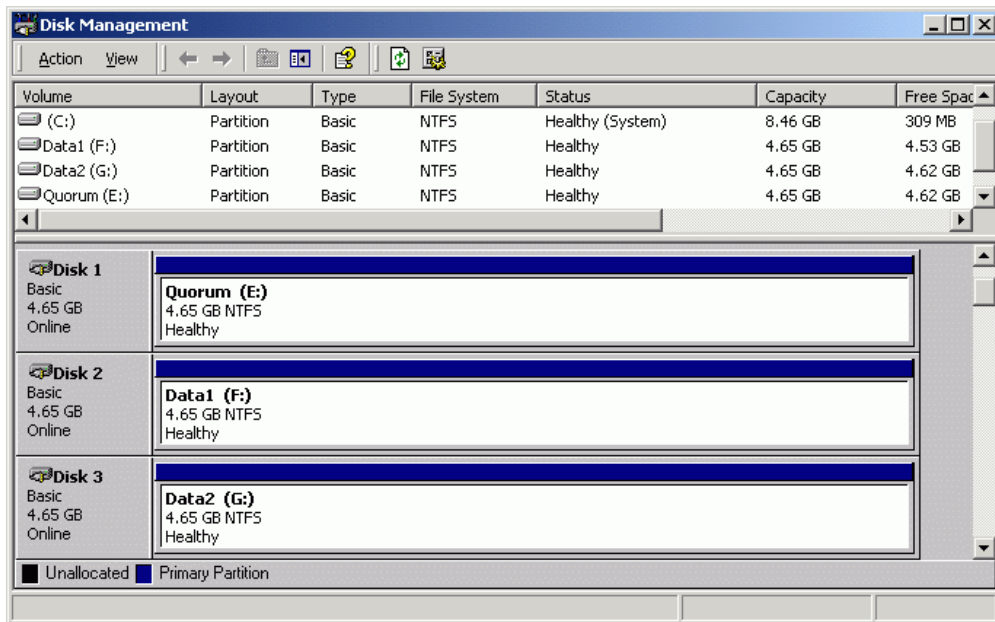
Figure 12-5   The Disk manager of RootBeer sees every disk

The ESS Specialist Web interface displays the PPRC states of the cluster volumes. RootBeer's volumes are PPRC Primary volumes (blue icon), while MountainDew's volumes are PPRC Secondary (red icon) (see Figure 12-6 on page 454).
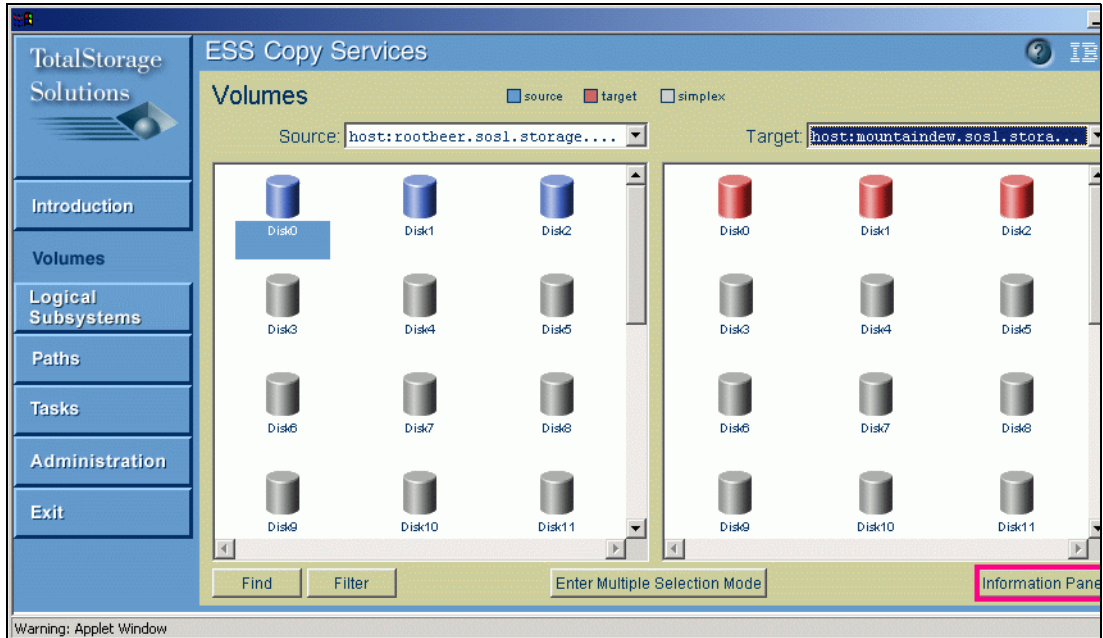
Figure 12-6   PPRC primary and secondary role

2. We now move all of the resource groups to MountainDew. The result is seen in Figure 12-7.
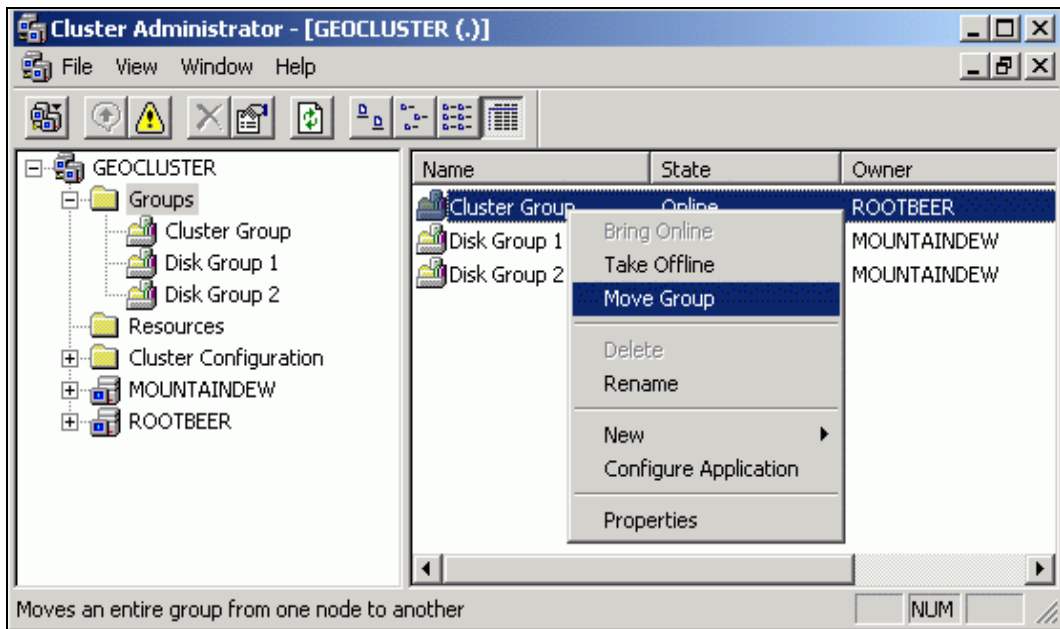
Figure 12-7   Initiate move group on RootBeer for all three groups

The move typically takes only a few seconds, but may be longer depending on the hardware, the work load, and the complexity of the applications running on MSCS.
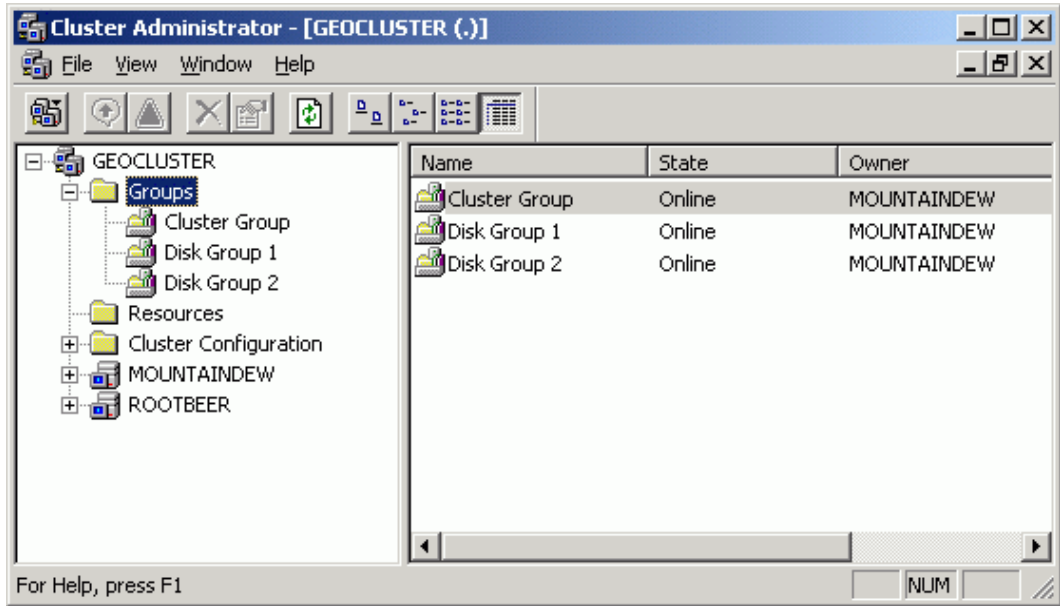
Figure 12-8   All groups are now owned by MountainDew

MountainDew's volumes are now PPRC primary volumes, while RootBeer's volumes are PPRC secondary volumes (The direction of the PPRC relationship has been reversed). See Figure 12-9.
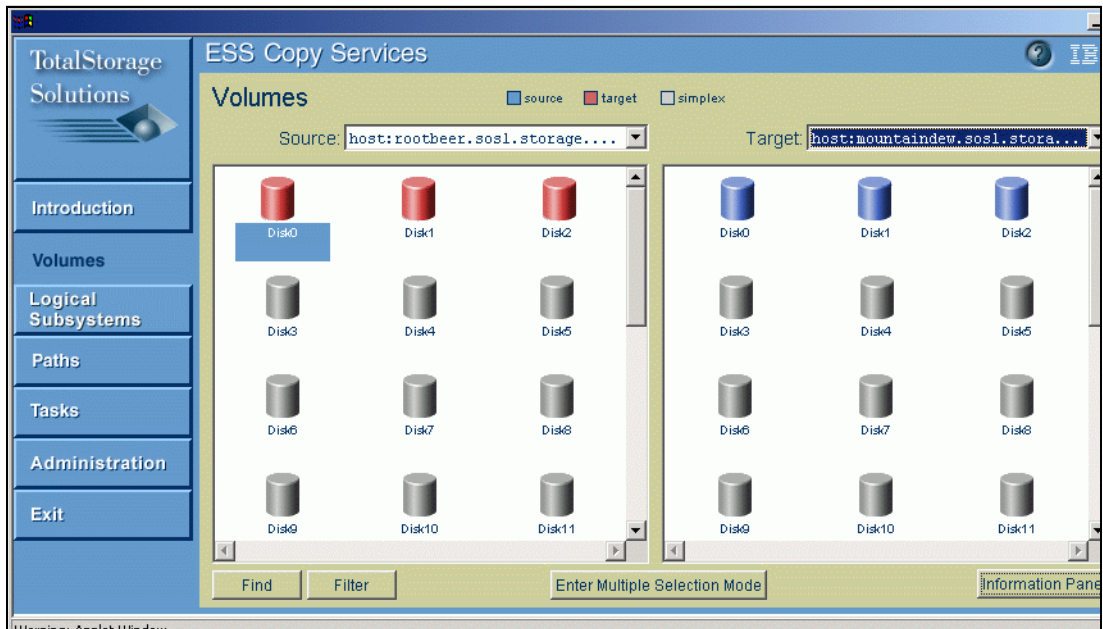
Figure 12-9   PPRC primary and secondary role exchanges

3. Now power off MountainDew to force a cluster failover. The groups and resources will fail over to RootBeer.
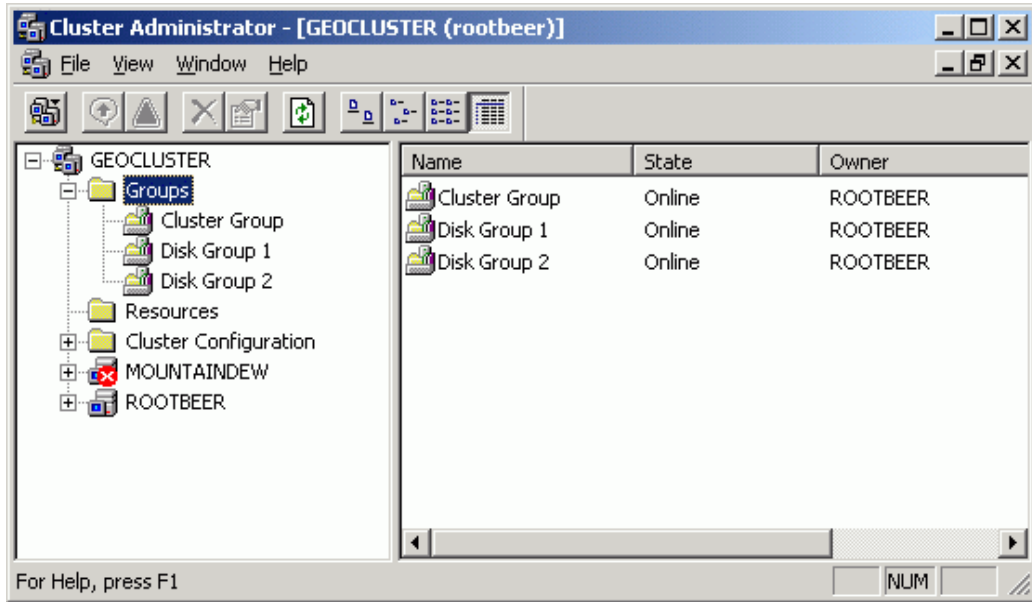
Figure 12-10   The groups are now again owned by RootBeer

RootBeer's volumes are now PPRC primary volumes, while MountainDew's volumes are PPRC secondary volumes. (The direction of the PPRC relationship has been reversed again.)
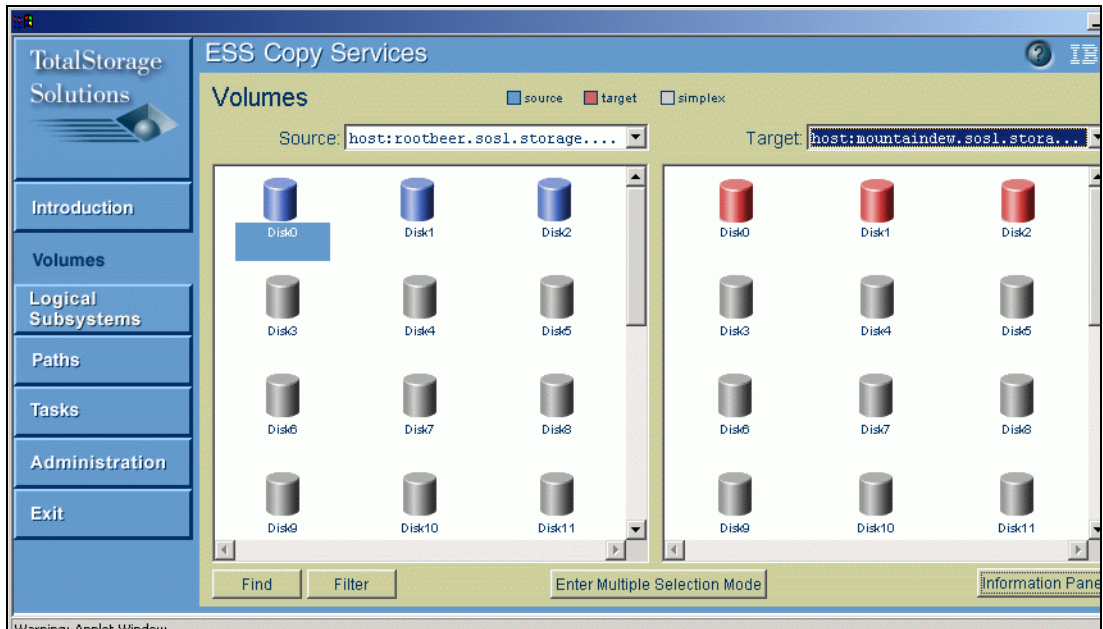
Figure 12-11   PPRC role exchange again

12.1.3 Scenarios

We have seen an overview of GDS, so let us discuss how it works in possible scenarios. In all of the situations described in the following scenarios, node "Server A" is the original owner of the quorum disk, which is a PPRC primary volume on "ESS A".

Scenario 1: Total network communication fault

Both the public and private network links are broken, and all communication is lost between Server A and Server B (see Figure 12-12).
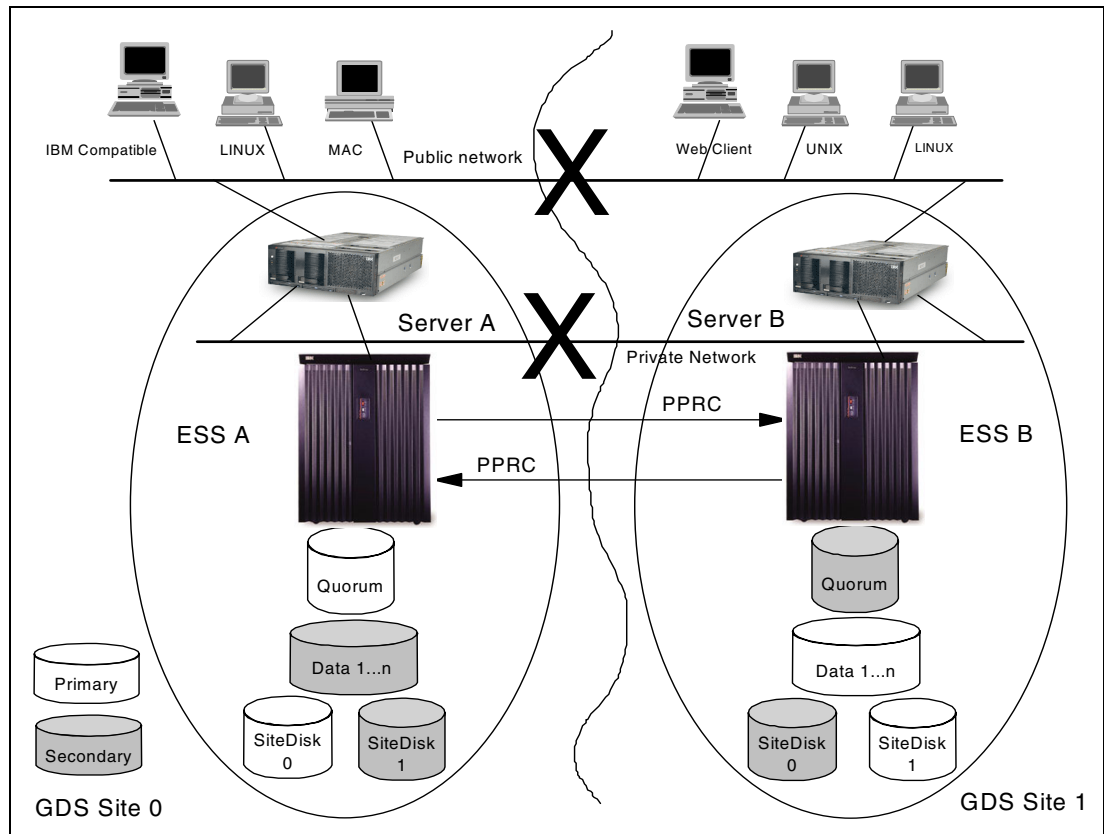
Figure 12-12   Total network communication fault

Normal MSCS operation

When both the private and the public networks have failed, then the cluster nodes will start the quorum resource arbitration process, in which all the cluster nodes try to reserve the quorum resource. The cluster node that 'wins' the quorum resource arbitration will then bring all resources online. The MSCS is stopped, and then Windows tries to restart the cluster service on a predetermined schedule.

IBM GDS operation

The fail-over service determines if the node at the other site is up or not through the PPRC link, even without networking available. No failover is necessary if the quorum owning node is still up. This adds another layer of protection compared to normal MSCS operation.

After the network connection is back to normal, no additional data replication occurs, as the PPRC link was always up and the data was continuously being replicated.

Scenario 2: Storage communication fault

In the event of an ESS communication fault, see Figure 12-13 on page 458.

A PPRC communication fault indicates that no site-to-site mirroring is occurring. If the PPRC communication is broken, but the cluster nodes are still able to communicate, then based on a user selected setting, **PPRC Fault Mode**, GDS takes different actions.

► If the PPRC Fault Mode is set to **Failstop**, then all PPRC resources are taken offline, and MSCS on both nodes is shut down. This option is used in the environment where data integrity and data consistency is the most important.

► If the PPRC Fault Mode is set to **Failover**, the quorum owning node stays online, the peer node shuts down, and all the resources fail over to the quorum owning node. This option is used in customer environments where data availability is the most important.

► If the PPRC Fault Mode is set to **No_New_Onlines**, the quorum owning node stays online and the resources stay online. The cluster service at the peer node shuts down and all of its resources go offline; these resources will not fail over to the quorum owning site. When the PPRC connection re-establishes, only the changes are replicated. This is the default setting.

► If the PPRC Fault Mode is set to **Online_No_Moves**, then both nodes stay online, and no resources are taken offline. No moves are allowed until the PPRC link is repaired.
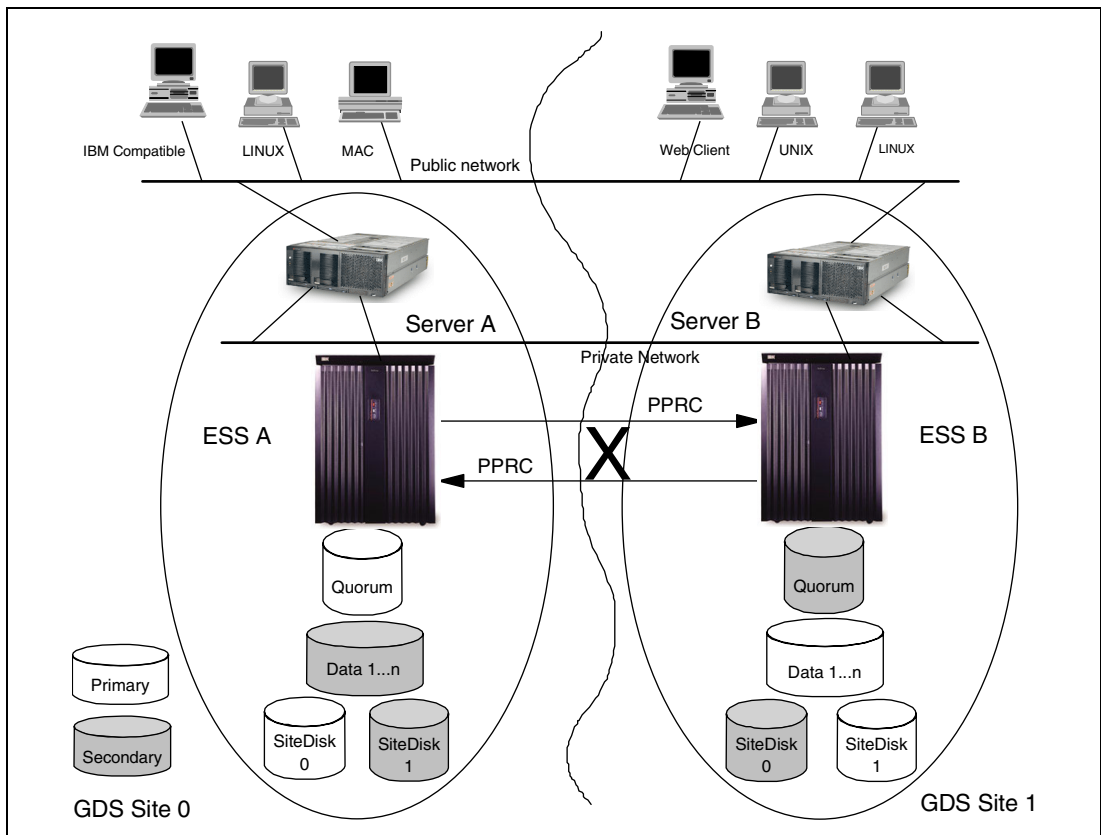
Figure 12-13   Storage communication fault

Scenario 3: Server fault

When either of the cluster nodes fail or the node cannot bring up the corresponding resource, failover happens. The topology is displayed in Figure 12-14 on page 459.

In the case of failure:

1. The remaining cluster node arbitrates for the quorum resource, and then brings the resources from the failed cluster node online by reversing the PPRC direction.

2. The PPRC Cluster resource DLL fails over the other disks' PPRC directions.

3. When the server is repaired, the resources may fail back to their preferred owner (based on the user-defined fail-over group's fail-back policy).

After the fault server is fixed, no PPRC replication happens, since data is always synchronized.
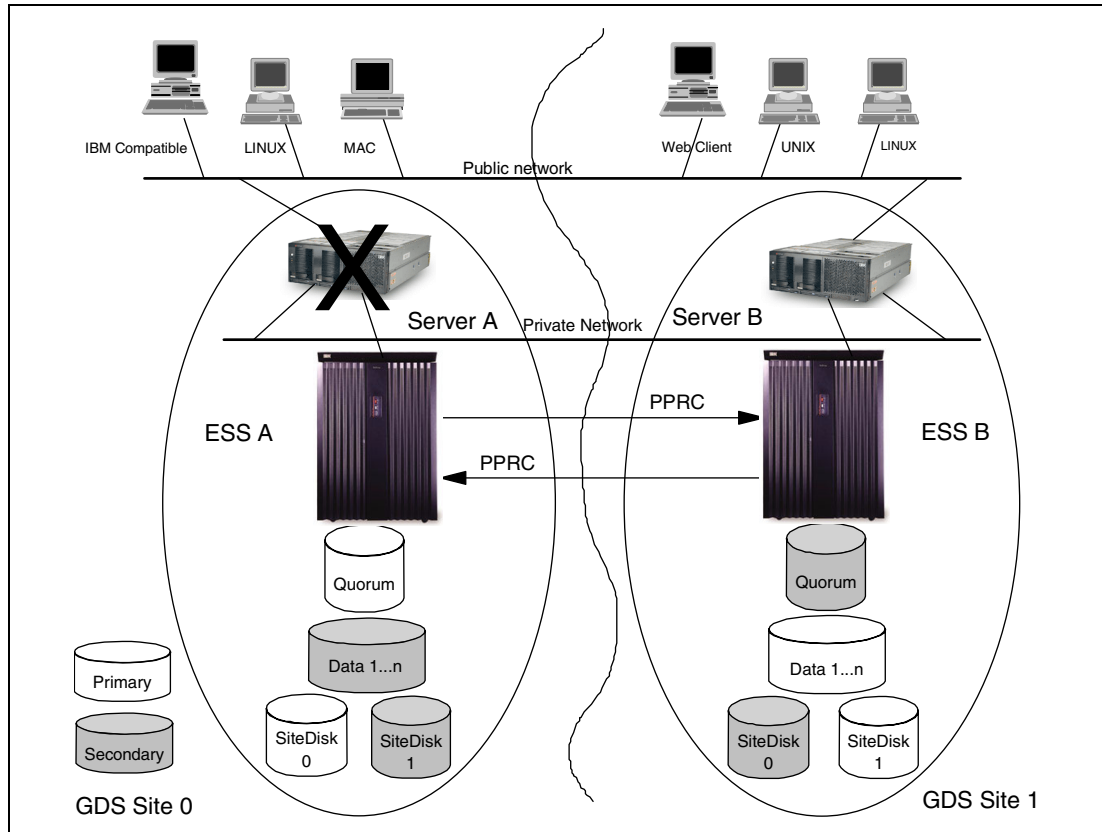
Figure 12-14   Server fault

Scenario 4: ESS fault

When one of the ESSs fails (this is an extremely rare case; ESS is highly, highly reliable), a failover happens. The topology is displayed in Figure 12-14. As mentioned above, the ESS has an internal data protection mechanism, which will protect data integrity. In the case where an entire ESS fails, an IsAlive query will fail on all cluster nodes at that site and the PPRC link status will change.

In the case of failure:

1. The cluster node at the site with the failed ESS will take its resources offline.

2. The cluster node at the site of the functioning ESS will bring the resources online, by using the fail-over service, failing over the Quorum disk's PPRC direction, and the PPRC Cluster resource DLL failing over the other disks' PPRC direction.

3. When the ESS is repaired and the PPRC actions are re-established, based on a user setting (fail-over group fail-back policy) in MSCS, the resources will fail back to their preferred owner.
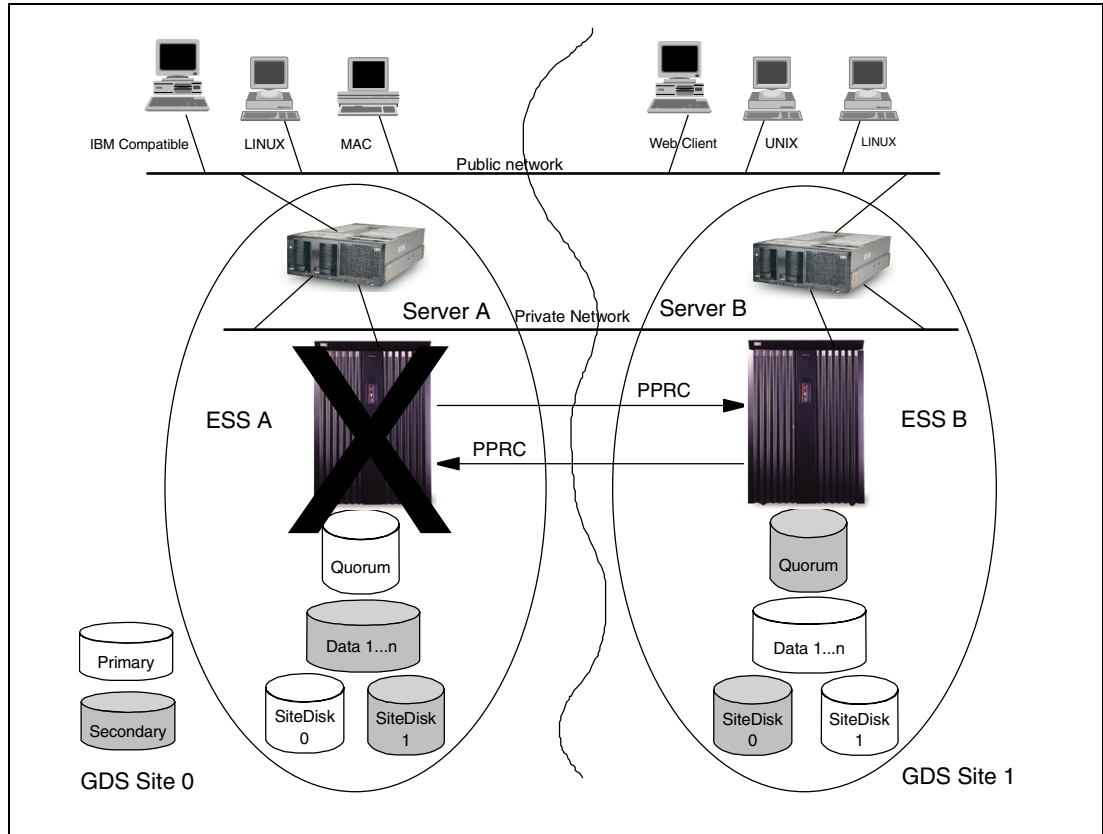
Figure 12-15   ESS fault

Scenario 5: Site fault or total communication fault

An entire site is compromised or all communication channels (public network, private network, and PPRC link) fail between the two sites (see Figure 12-16 on page 461 and Figure 12-17 on page 461).

During a catastrophic failure, such as all cluster nodes and storage at a site failing, or even during a simulated site failure, such as all communications (networking & PPRC link) failing, the fail-over policy is managed by an additional user setting, **PPRC Fault Mode** (see "Scenario 2: Storage communication fault" on page 457 for details on this setting).
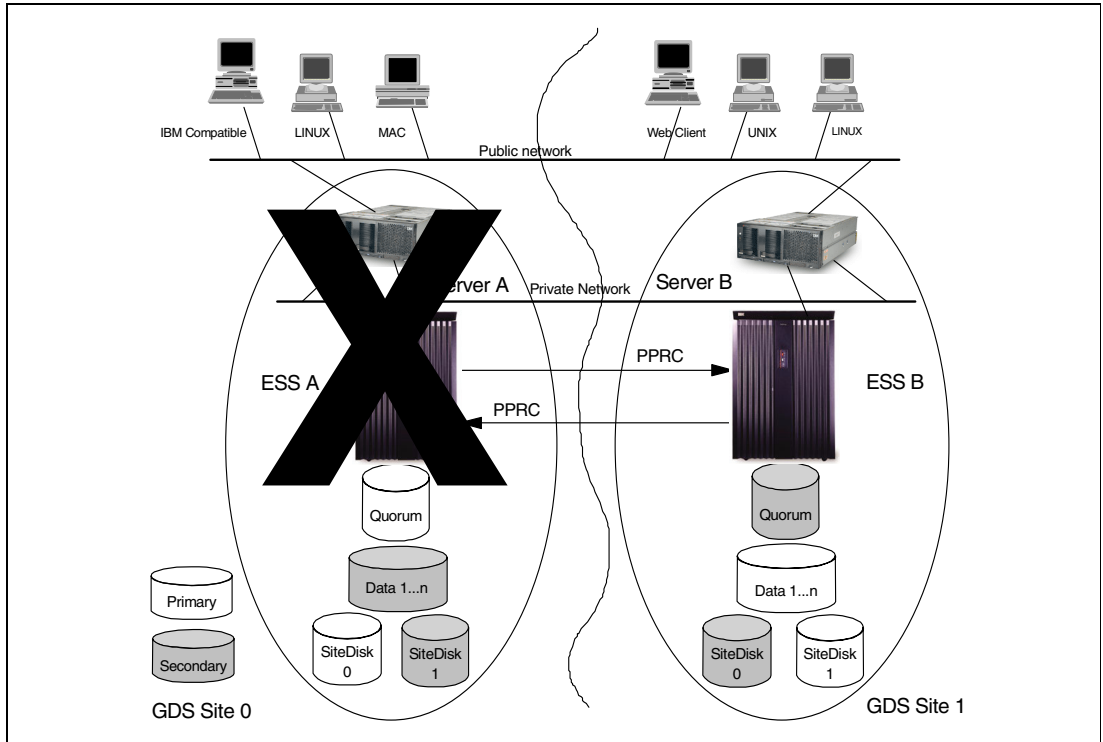
Figure 12-16   Site fault
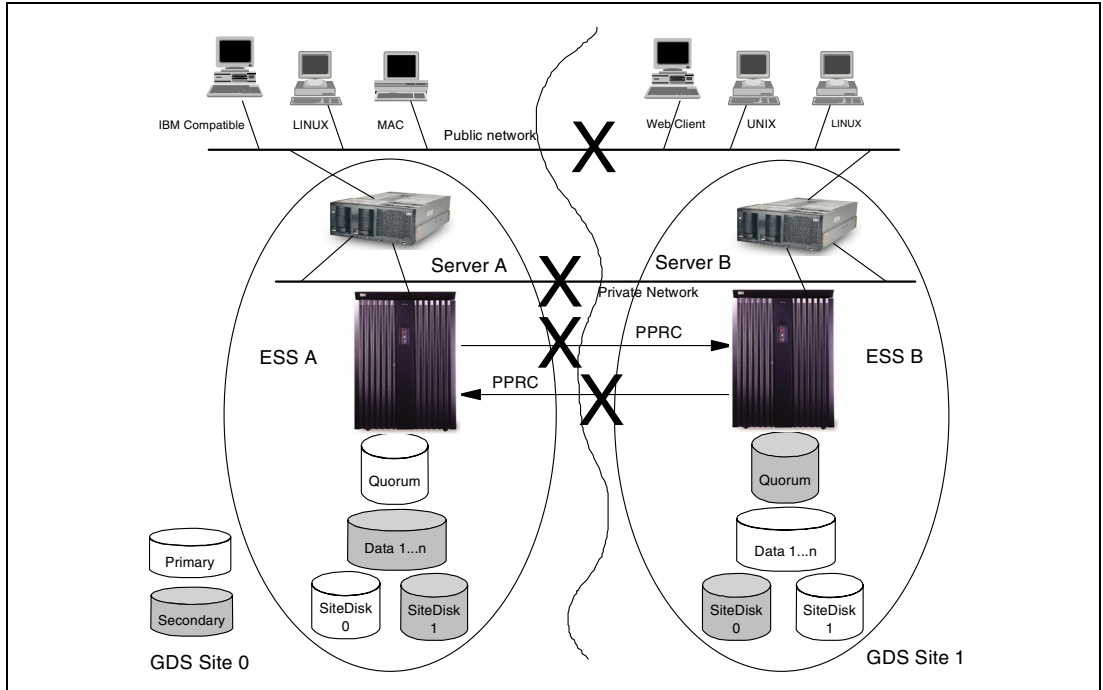
Figure 12-17   All communication fault

12.2 IBM TotalStorage Data Manager for SQL 2000

IBM TotalStorage Data Manager for SQL 2000 (TDM for SQL 2000) exploits the freeze/thaw capability in SQL Server 2000 databases on the IBM Enterprise Storage Server (ESS) to allow FlashCopy backups. TDM for SQL 2000 performs backup and restores.

12.2.1 Introduction

IBM's FlashCopy feature on the ESS can provide quicker backup and restore operations with minimal service disruption on SQL Server 2000 databases. In the case of backup, FlashCopy can reduce the performance impact on the production server from several hours to just minutes. More importantly, should the database need a restore, FlashCopy can similarly reduce the outage from hours to minutes. A set of sample scripts and the associated application are provided for single server or single site usage.

TDM for SQL 2000 provides online snapshot backup and restores of Microsoft SQL Server 2000 databases on the IBM TotalStorage Enterprise Storage Server (ESS). It uses the FlashCopy or Peer-to-Peer Remote Copy (PPRC) capability of the ESS to create online backups of the database (on both local and remote ESSs) in seconds or minutes. During a backup operation, TDM for SQL 2000 uses the Microsoft Virtual Device Interface (VDI) to suspend writes on the SQL Server 2000 database to preserve data integrity during the brief interval it takes to execute a FlashCopy. Then it uses the FlashCopy feature of the ESS to create snapshot copies of the database volumes. These backups can be used for database cloning, data mining, or for disk or tape backups. The application then registers the backup with the msdb database in SQL Server, so that the backup can be used for a database restore.

Without this application, customers would require the shutting down of their database to create a FlashCopy backup, or use online backup tools, which are CPU intensive and take a long period of time to complete.

12.2.2 Benefits

► Improved Availability
► Hot backups with no application downtime
► Fast recovery
► Local and remote copies
► No impact on users, host server, or network
► Tight Integration with SQL Server 2000
► Uses SQL Server Virtual Device Interface (VDI)
► Backups recorded in SQL Server msdb database
► Snapshot backups can be rolled forward
► Supports creation of clone databases
► Data mining/business intelligence
► Application development and testing
► Compatible with TSM, VERITAS, Legato, and NTBackup

We can use either the TDM for SQL 2000 graphical user interface (GUI) or the application's command line interface, to perform backup and restore operations. We can also use the TDM for SQL 2000 GUI to generate a script, which we can then customize for unattended backup and restore for our business needs.

Hardware and software requirements
► Intel server with a minimum of 133 MHz and 64 MB RAM

> ► At least one IBM TotalStorage Enterprise Storage Server (ESS)

> **Note:** Ensure that each ESS in our system is configured properly for the ESS Specialist, FlashCopy, and PPRC. For more information about the ESS Copy Services configuration, see the *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide,* SC26-7346*.*

> ► Microsoft Windows 2000 (Service Pack 2 or higher)

> ► Microsoft SQL Server 2000 (Service Pack 2 or higher)

> ► IBM 2105 Command Line Interface (CLI) version for ESS, microcode Version 1.5 or above

> ► Either of the following Web browsers:
>
>   – Netscape Navigator Version 4.0 or above
>
>   – Microsoft Internet Explorer Version 4.0 or above

> **Tip:** If we have difficulty using a browser to access the ESS Specialist, upgrade the browser to the highest available level.

12.2.3 Installation of TDM for SQL 2000

1. Install the CLI. Run TSDMSql_Installer.exe from the installation media.

2. Follow the instructions in the install shield wizard. The instructions will guide you through the installation.

3. Click **Finish** to complete the installation.

4. Establishing Authorization:

   TDM for SQL 2000 invokes SQL Server and ESS Copy Services applications. Therefore, access to TDM for SQL 2000 requires SQL Server authorization and may require ESS advanced copy services authorization. If a user name and password have been established for ESS Advanced Copy Services, we must specify that user name and password in the TDM for SQL 2000 Options before we can use TDM for SQL 2000. To learn how to create ESS advanced copy services user names and passwords, see *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7346. To get SQL Server authorization, log into TDM for SQL 2000, specify the Windows or SQL Server authentication in the login window. For SQL Server authentication parameters, refer to:

   http://www.microsoft.com/support

12.2.4 Using TDM for SQL 2000 GUI to backup/restore a SQL Server 2000 database

During a backup operation, TDM for SQL 2000 generates a metadata file that contains information about the backed up database.

> **Tip:** It is best to back up the metadata file. It contains information about the backup that TDM for SQL 2000 uses during restore operations. The metadata file is saved in the following format: instancename_dbnametimestamp.mtd.

TDM for SQL 2000 stores metadata file in two locations:

► The metadata folder that is specified in TDM for SQL 2000 Options.

► The FlashCopy target volume contains the primary data file for the backed-up database.

Before we begin, we should define and save the following ESS advanced copy services tasks using the ESS specialist:

► For local backup operations, create a local FlashCopy.

► For remote backup operations, suspend the PPRC task, create a remote FlashCopy, and resume the PPRC task.

Perform a backup

During a backup operation, TDM for SQL 2000 instructs the SQL Server to suspend write operations to the disk. Read operations continue as normal. If a user attempts to perform a write transaction while the backup operation is in progress, the transaction is suspended, and the Windows 'hourglass' icon appears. The transaction continues when the backup operation is complete. The user does not have to resubmit the transaction. To perform a backup, do the following:

1. Log in to TDM for SQL 2000 using Windows NT authentication and SQL Server authentication. After logging in, the main TDM for SQL 2000 window appears showing the Backup view, as shown in Figure 12-18.
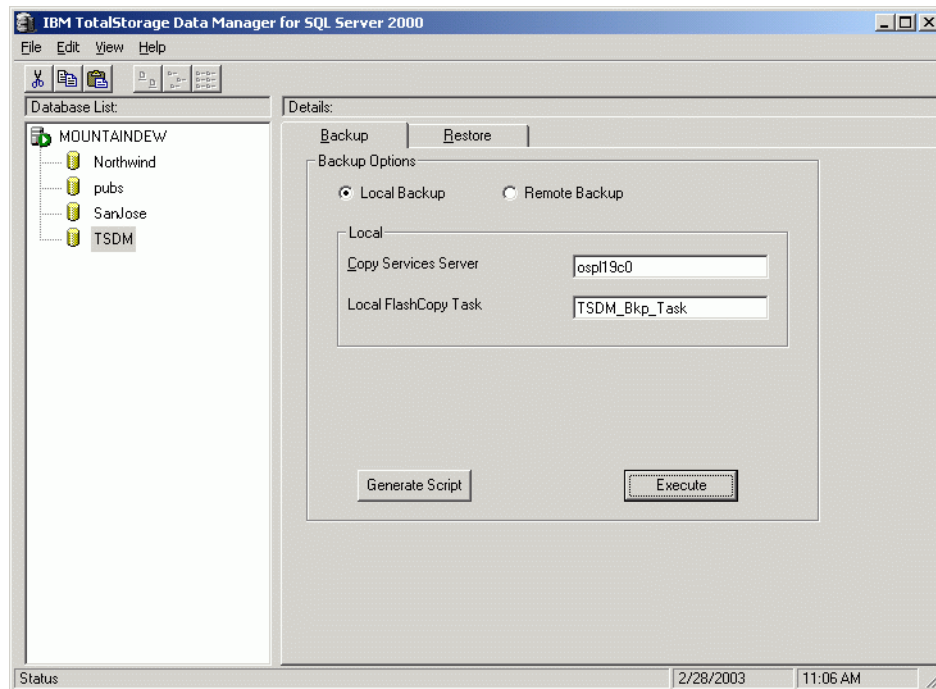
Figure 12-18   TSDMSql backup window

2. In the left navigation section of the main window, highlight the database we want to back up.

3. Select Local Backup or Remote Backup.

   – If we select Local Backup, specify the name of the Advanced Copy Services server and the predefined local FlashCopy task in the appropriate fields.

   During local backup, TDM for SQL 2000 does the following things:

   i. Temporarily suspends SQL server writes. Reads are allowed to continue.

   ii. Flushes all buffers to disk.

   iii. Invokes the FlashCopy establish function to copy the disks. This is only the logical establish and takes only a few seconds.

      iv. Writes are allowed to resume. No writes are lost, no writes have to be re-issued.

      v. The copy is processing in the background and is set to the lowest priority, so it has very little effect on performance.

  – If we select Remote Backup, specify the name of the Advanced Copy Services server and the predefined suspend PPRC task, remote FlashCopy task, and resynchronize PPRC task in the appropriate fields.

    During remote backup, TDM for SQL 2000 does the following things:

      i. Temporarily suspends SQL server writes. Reads are allowed to continue.

      ii. Suspends the PPRC relationship.

      iii. Writes are allowed to resume. No writes are lost, no writes have to be re-issued.

      iv. Creates a FlashCopy of the PPRC target volume.

      v. Resumes the PPRC relationship.

4. The function of generating a script, is used to schedule the backup at any time. Select **Generate Script**. To run the backup operation immediately, select **Execute**.

  – If we select **Generate Script**, a dialog box appears. Choose the location and name of the script, and then select **Save** in the dialog box.

  – If we select **Execute**, TDM for SQL 2000 runs the backup operation immediately.

When it is all done, a message indicates the total operation time in seconds, and whether the backup was successful. If the backup fails, an error log appears (see Figure 12-19).
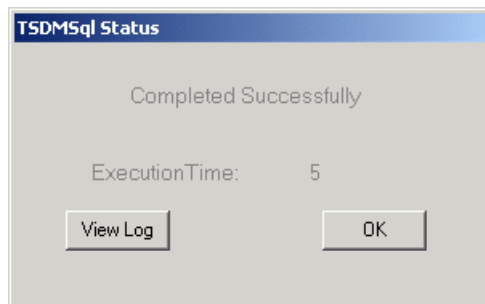
Figure 12-19   Finish window

Perform a restore

We can use the TDM for SQL 2000 GUI to perform one or more restore operations, and then specify recovery. We might have several SQL Server log restores to apply before we specify recovery.

> **Note:** We cannot use TDM for SQL 2000 to perform a restore operation from a remote server. To learn how to manually restore from a remote server, see the white paper *Storage Solutions for Microsoft SQL Server: Snapshot Backup and Recovery with the IBM TotalStorage Enterprise Storage Server*, which you can find at:
>
> http://www.storage.ibm.com/disk/ess/whitepaper.htm

Before we begin, please ensure:

► Using the ESS Specialist, define the local FlashCopy task that performs the restore operation.

► Know the name of the metadata file that TDM for SQL 2000 needs to restore the database.

Perform the following steps to restore a database using the TDM for SQL 2000 GUI:

1. Log in to TDM for SQL 2000 using Windows NT authentication or SQL Server authentication. After we log on, the main TDM for SQL 2000 window appears showing the Backup view.

2. Select the **Restore** tab to switch to the Restore view, as shown in Figure 12-20.

3. On the left navigation section of the main window, highlight the database we want to restore. TDM for SQL 2000 updates the Restore view with a list of metadata files for backup versions of the highlighted database.
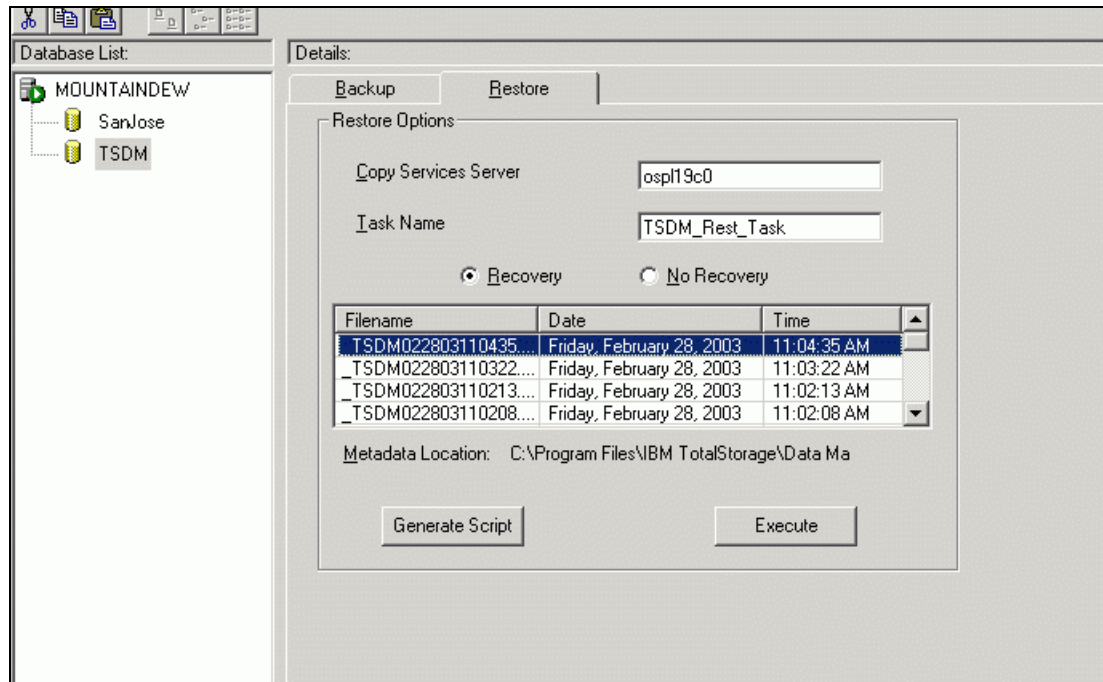
Figure 12-20   TSDMSQL restore view

4. Select the metadata file of the backup version that we want to restore.

5. Specify the name of the Advanced Copy Services server and the predefined local FlashCopy task in the appropriate fields.

6. Select **Recovery** or **No Recovery**:

   – Select **No Recovery** if we have more SQL Server log restores to apply before recovery.

   – Select **Recovery** if we do not have more SQL Server log restores to apply before recovery.

   Recovery is the default. If we select **No Recovery**, the database is offline and unavailable for use until we apply all log restores.

7. To generate a script that we can customize and use to schedule a restore at any time, select **Generate Script**. To run the restore operation immediately, select **Execute**.

   – If we select **Generate Script**, a dialog box appears. Choose the location and name of the script, then select **Save** in the dialog box.

   – If we select **Execute**, TDM for SQL 2000 runs the restore operation immediately.

When it is done, a message indicates the total operation time in seconds, and whether the restore was successful.

12.2.5 TDM for SQL 2000 command to back up/restore SQL Server 2000 database

There is also a command line version of the TDM for SQL 2000 program to perform local, remote backup and restore operations. They offer a more flexible way for a user defined application to automate backup and restore operations. They are:

▶ TDM for SQL 2000 backup
▶ TDM for SQL 2000 restore

12.2.6 Scenario

In this scenario, we have a SQL Server 2000 and we create a database named *TSDM*. Its data file is on the X: drive and the Log file is on the Y: drive; they are both located in the ESS. We run a script to insert multiple rows into a SQL Server table in TSDM. We use the Windows Performance Monitor to examine the Disk Writes/second for the Data (X:) and Log (Y:) volumes while executing a backup using TDM for SQL 2000 (see Figure 12-21).

The green line indicates that the SQL Server is busy writing into the log file located on Y:. When the backup is executed, the disk's activity suddenly drops to zero. After five seconds, the backup is complete, and writes are resumed as normal.
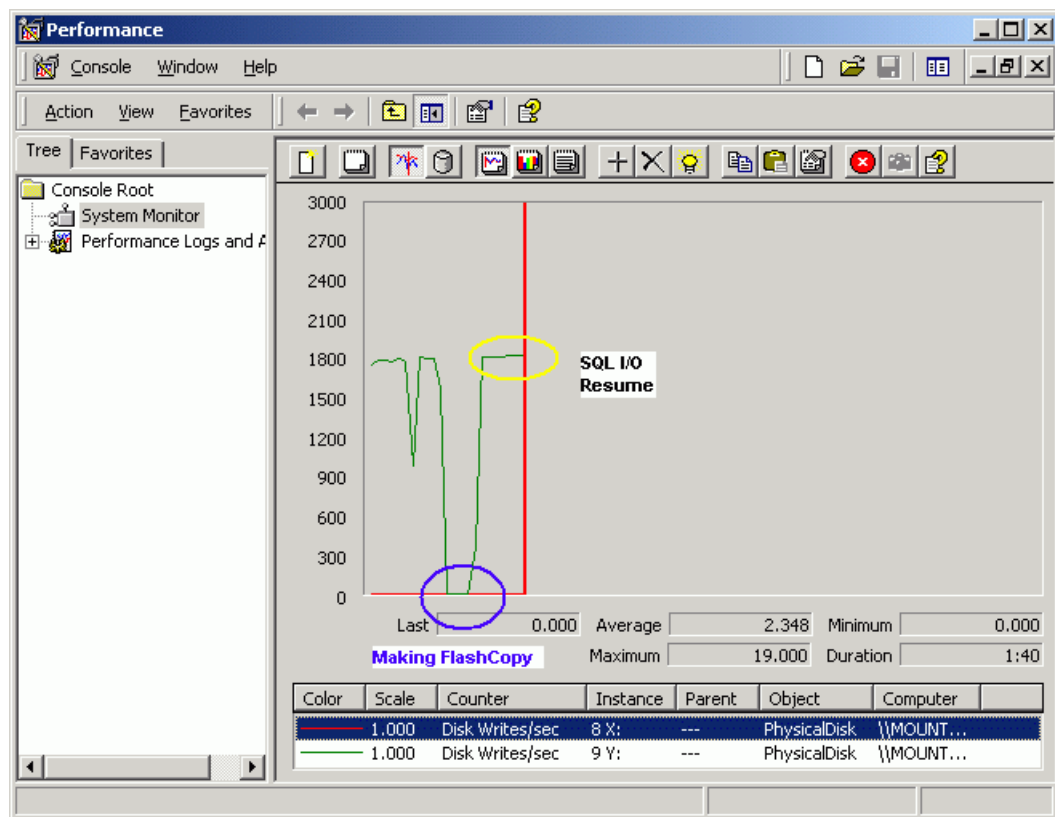
Figure 12-21 Performance monitor

After that, we can check the SQL Server event log to see if the backup operation has successfully finished.
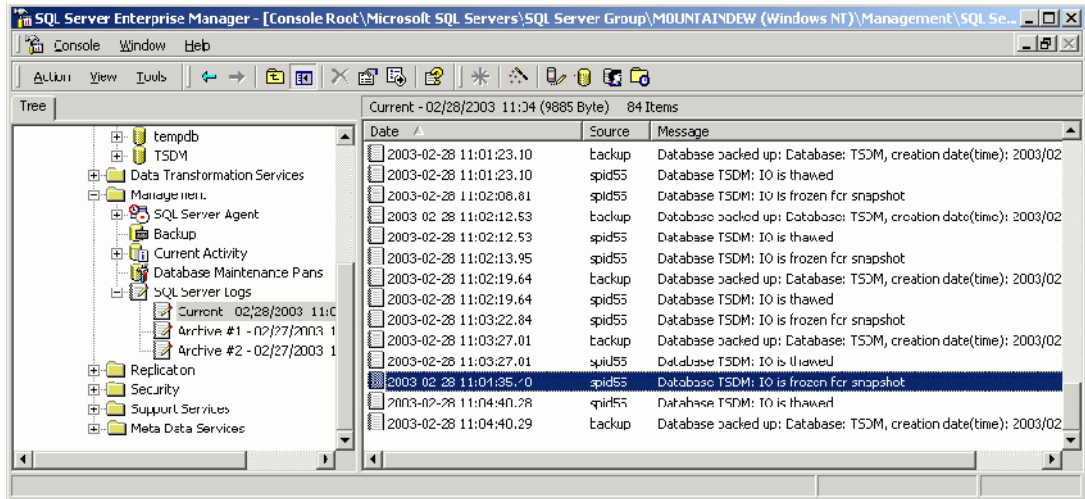
Figure 12-22   SQL log

12.3 Other Service Offerings for Windows and the ESS

The attachment rate for Windows servers on the ESS is quite high. IBM has developed some Service Offerings to allow Windows servers to take advantage of the ESS Copy Services. These Service Offerings are described in the following sections.

12.3.1 IBM TotalStorage ESS - Exchange Integration Resource Pack

This product allows the mounting and dismounting of Exchange storage groups from the command line. This enables backups using copy services or other mechanisms to be run from scripts. It also includes log management facilities.

The IBM FlashCopy feature on the ESS can provide quicker backup and restore operations with minimal service disruption on Exchange 2000 databases. In the case of backup, FlashCopy can reduce the performance impact on the production server from several hours to just minutes. More importantly, should the database need a restore, FlashCopy can similarly reduce the outage from hours to minutes. A set of sample scripts and associated ESS utility programs are provided for single server or site-wide usage. These tools will minimize the service disruption to the entire Exchange 2000 server while backup/restore is done for any particular Exchange database.

Benefits
► Improved Availability.
► Fast backups.
► Fast recovery.
► Local and remote copies.
► Reduced impact on users, host server, or network.
► Integration with Exchange Server 2000.
► Automated log management.
► Partially online backups.
► Snapshot backups can be rolled forward.
► Supports creation of clone databases.
► Application development and testing.
► Single mailbox recovery.
► Compatible with TSM, VERITAS, Legato, and NTBackup.

12.3.2 IBM TotalStorage ESSSync

This Copy Services automation utility integrates Windows NT and Windows 2000 with FlashCopy and PPRC. It serves two primary functions. First, it flushes file system buffers on the host(s) prior to a copy services action. This action is necessary to obtain a consistent copy. Secondly, it enables the target host to use the Copy Services target volumes without rebooting. It does this by temporarily making Windows unaware of the volumes. Then, at the first attempt to access the data, Windows "rediscovers" the volumes and immediately begins to use them.

Benefits

► Flushes file system buffers on Windows to ensure data integrity.

► Eliminates the need for a reboot of the target system before using the target volumes.

► Tested with MS SQL Server, DB2® UDB, Informix, Oracle, SAP, and various other applications.

► Fully scriptable.

12.3.3 IBM TotalStorage Volume Set Manager

This utility automates the process of copying Windows NT Volume Sets that span multiple ESS volumes, using FlashCopy. On Windows NT, Volume Set metadata is stored in the registry. When the customer uses the FlashCopy function to copy the ESS volumes (that make up the Volume set) to another set of ESS volumes, the target ESS volumes cannot be brought online on the target server because the configuration metadata is missing in that server's registry. This utility copies the disk registry information of a given Volume Set from the source NT machine to the target NT machine.

Benefits

► Allows for FlashCopies or PPRC copies of NT Volume Sets

► Automates the copy process

► Fully scriptable

A

ESS Application Programming Interface (API)

In this appendix, the following topics are discussed:

- ► The ESS Application Programming Interface (API)
- ► ESS API components
- ► The Client application
- ► Installing the ESS API
- ► Prerequisites

The ESS Application Programming Interface (API)

This section introduces the IBM TotalStorage Enterprise Storage Server (ESS) Application Programming Interface (API). The ESS API is to be used in conjunction with a client application.

The terms ESS API and ESS CIM Agent mean the same thing and are used synonymously.

For a more complete understanding and discussion of CIM please refer to *IBM Tivoli Storage Resource Manager: A Practical Introduction,* SG24-6886

ESS API characteristics

The use of ESS API demonstrates IBM's commitment to open standard interfaces. The ESS API is a non proprietary storage management agent that supports routine LUN management activities, such as LUN creation, mapping, and masking. The ESS API also enables ESS Copy Services configuration and use activities, such as FlashCopy and PPRC management. It supports these activities through the use of the Storage Management Initiative Specification (SMIS), as defined by the Storage Networking Industry Association (SNIA).

The ESS API helps integrate ESS configuration management support into storage resource management (SRM) applications, which allow users to benefit from existing SRM applications and infrastructures. The ESS API also enables the automation of configuration management through customer-written applications.

Common Information Model (CIM) is a management standard published by the Distributed Management Task Force (DMTF). It defines a common set of classes, associations, and object relationships for managing system components (including storage).

The following list provides an overview of the ESS API:

► Provides the means for client applications to have common access for data management applications, such as Tivoli SRM, HP OpenView, and VERITAS.

► ESS CIM Agent 1.2 implements CIM Schema 2.8 for Storage Management.

► Middleware: CIM API only, no user interface.
  – Translates a proprietary device interface to a CIM-compliant interface.
  – Enables resource management applications for ESS devices.

► Each device has its own CIM Agent associated with it.
  – ESS CIM Agent is not a stand-alone product on its own release schedule
  – Synchronized with device schedule

► Provides management of the ESS.
  – Simplifies ESS configuration and ESS Copy Services administration.

► Allows third-party software to manage the ESS using SNIA CIM protocols.

► ESS LIC level 2.1.0 provided API ESS CLI functionality.

► ESS LIC level 2.3.0 provides API support for the SMI-S Storage Virtualization Initiative (SVI) interfaces.

► A standards-based programmatic interface to the ESS.

ESS API components

A CIM agent consists of the components shown in Figure A-1. The main components are the CIM object manager (CIM/OM), the service location protocol (SLP), and the device provider. The SLP is a directory service that a client application calls to locate the CIM Object Manager. A device can be a storage server, such as the ESS.

The CIM/OM and device provider communicate through method calls made from the CIM/OM to the provider. The device provider communicates with the device through proprietary calls.
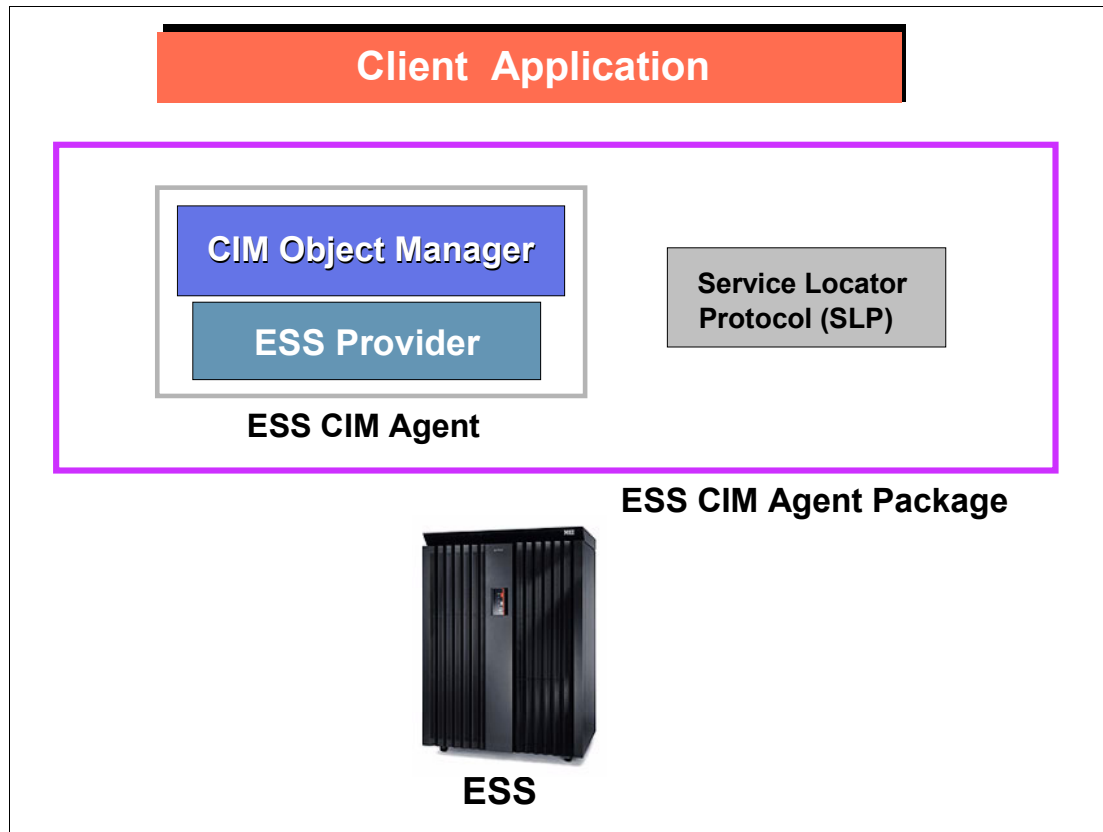
Figure A-1   Components of the API

The specific ESS API method of operation is shown in Figure A-2 on page 474. The following list describes the method of operation:

► The ESS CIM Agent registers itself with Service Locator Protocol to enable discovery by the Client application.

► A client discovers the ESS CIM Agent by calling the Service Locator Protocol service (or else it has to have built-in knowledge of the agent's network location).

► A client application makes calls to the ESS CIM Agent.

► The client application and the CIM/OM communicate through CIM Messages.

► The CIM Object Manager calls ESS Provider.

► The ESS CIM Agent (Provider) makes ESSCLI/ESSNI calls to an ESS.
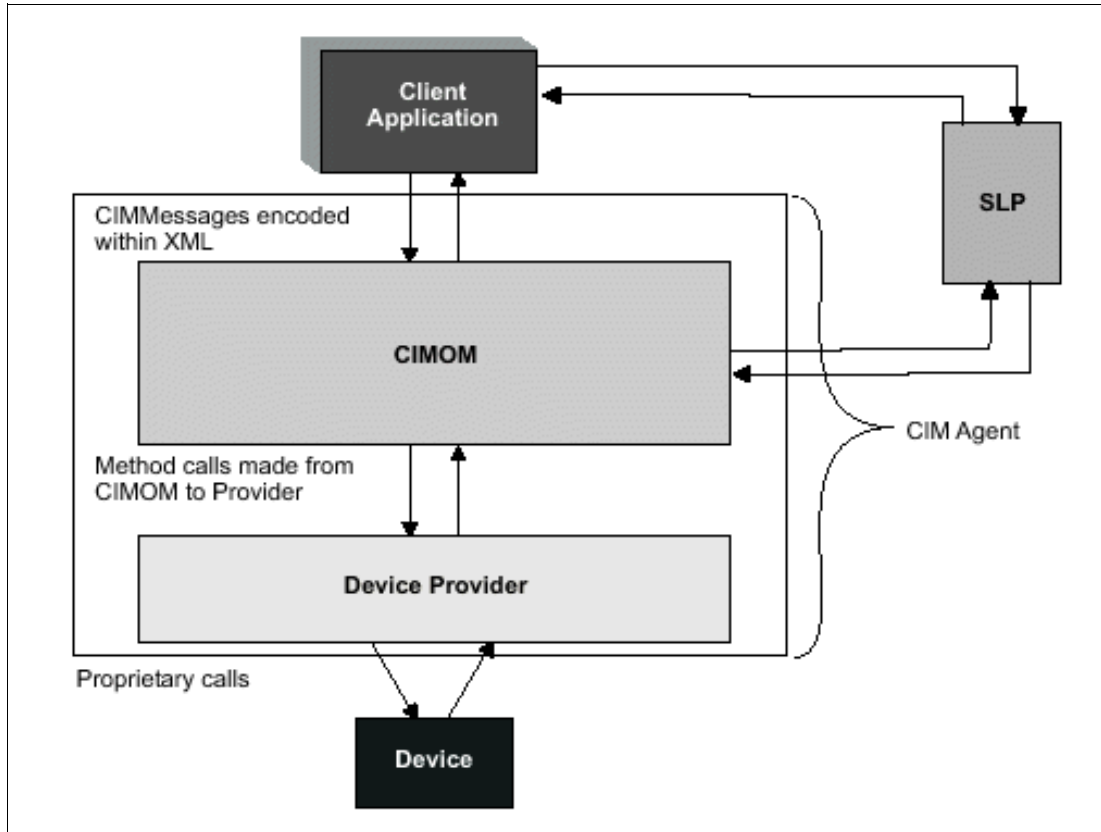
Figure A-2   API method of operation

The ESS CIM Agent subcomponents are not exposed to the client. The client application performs requests to the CIM Agent.

Summary

The following list is a summary of what the ESS API does:

► ESS API (ESS CIM Agent) allows third-party software to manage the ESS using SNIA CIM protocols.

► ESS CIM Agent provides an abstraction of control of ESS operations (including ESS Copy Services) to the data management client applications and does not explicitly expose vendor unique functionality.

► ESS CIM Agent implements a standards-based programmatic interface to the ESS.

The ESS API interface

The ESS API presents another option for ESS management by complementing the use of the ESS Web-based user interfaces (ESS Copy Services Web user interface and ESS Specialist), the ESS command line based interfaces (ESS Copy Services CLI and ESS CLI), as well as the z/OS interfaces (TSO commands, ICKDSF commands, and ANTRQST API), as Figure A-3 on page 475 illustrates.

It provides management of the ESS by simplifying ESS configuration and Copy Services administration.
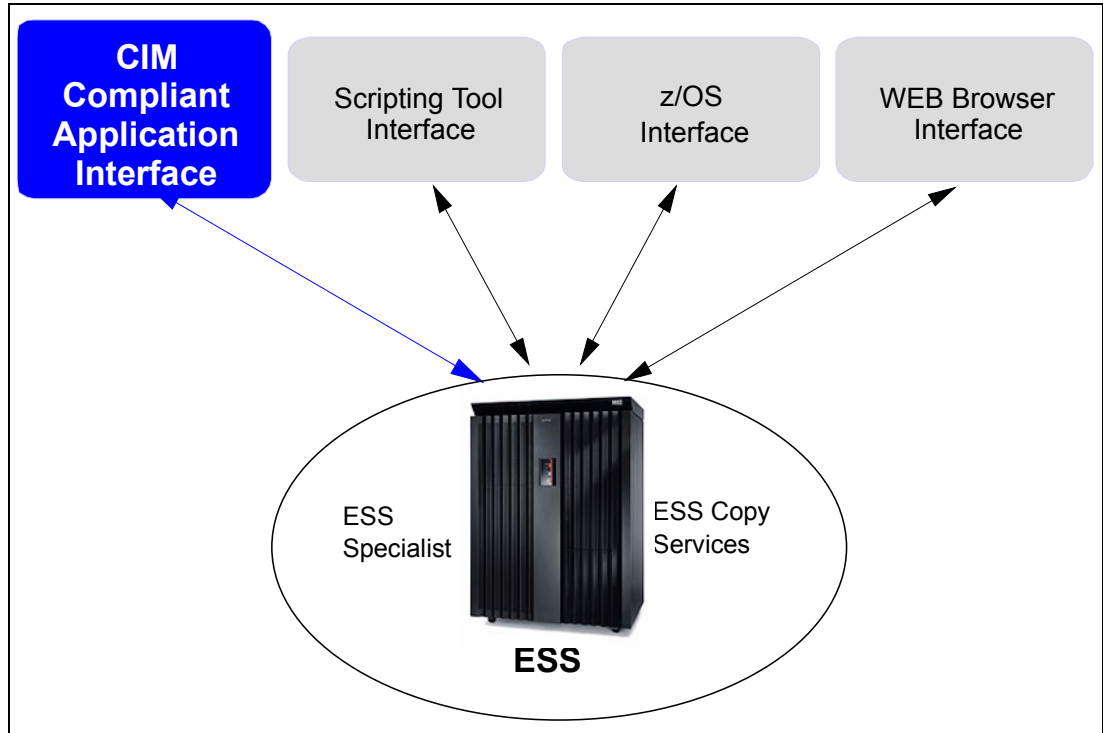
Figure A-3   ESS management interfaces

The client application

The client application can be an application for business continuance or storage administration. The following are things that can be done by a client application using the ESS CIM Agent services:

> **Note:** SNIA standard terminology is used in the following lists.

- ► Device Management
  - – Create host connections
  - – Create LUNs
  - – Connect and disconnect LUNs from hosts
  - – Discover all hosts known to the ESS
  - – Discover the storage configuration of an ESS
  - – Physical disks
  - – Logical volumes
- ► Storage Pool/Copy Services
  - – Create new storagePool
  - – Delete storagePool and return space
  - – Create volumes in a storagePool
  - – Create a copy relationship
  - – CopyType could be Synchronous, Asynchronous, UnSynchronizedAssoc, and so on
  - – Modify a copy relationship
  - – Detach: Teminate the relationship
  - – Fracture: Suspend the relationship
  - – Resync: Re-establish the relationship
  - – Restore: Use the target to update the source

Using the ESS Web user interface copy services, tasks are created and saved and later can be executed by the ESS Copy Services CLI using commands such as `rsExecuteTask.sh`. The ESS CIM API can enable a client application to do some of the same type of ESS data management functions that a saved task might do, mentioned here in "Device Management" on page 475 and "Storage Pool/Copy Services" on page 475.

Installing the ESS API

The ESS API CIM Agent and CLI are provided with the ESS LIC at no additional charge. The CIM Agent is available to run on AIX, Linux, and Windows 2000 or later operating system environments.

Please refer to the interoperability matrix for the most current information at:

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

Please refer to publication *IBM TotalStorage Enterprise Storage Server Application Programming Interface Reference*, GC35-0489 for instructions on installing the ESS API and CLI on your operating system.

Prerequisites

In order to install the ESS API in your environment, consider the prerequisites listed in "Hardware" on page 476 and "Software" on page 477.

Hardware

Ensure that your system satisfies the prerequisites for installing the ESS CIM Agent on a Windows 2000 or later, AIX, or Linux operating system before starting the installation.

The following hardware is required:

► Personal computer, workstation, or server with Intel Pentium® III or higher processor
► CD-ROM drive
► Video graphics adapter display or better

The following space on your Intel workstation is required:

► 1 GHz PIII or PIV processor (2 GHz is better.)
► 1 GB of random-access memory (RAM) minimum, depending on your system configuration
► 100 MB of disk space
► Up to 50 MB of temporary disk space for installation purposes

The following space on your AIX server is required:

► 1 GHz processor
► 1 GB RAM
► 4 GB disk as well

Note: You might need to increase the total available disk space on your hard drives if the ESS CIM Agent and other associated products are split between more than one logical drive. Also, the ESS CIM Agent might require additional memory to operate if you configure it to manage many devices or devices with large configurations.

For the ESS, the following considerations apply:

- ► LIC level must be 2.3.0 or later.

- ► The CLI is not supported with E models but is with F and 800 models (refer to the following Web site for a complete interoperability list at http://www.storage.ibm.com/disk/ess/supserver.htm).

Software

All prerequisite software must be installed before you start the installation of the ESS CIM Agent. The ESS CIM Agent installation program will check for the existence of the ESS CLI. If the program does not detect the ESS CLI, the installation of the ESS CIM Agent will not complete successfully.

The following software is required:

- ► Operating systems:
  - – Windows 2000 or later
  - – AIX 5L Version 5.1 or later
  - – Linux Red Hat 7.2

    Note: For the latest versions of Red Hat compatibility with the ESS API, please refer to:

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

- ► The workstation/computer/server must be IP connected to ESS(s) and Storage Management Client.

- ► ESS CLI level 2.3.0 or later. This software is on the ESS CLI CD.

- ► Common Information Model (CIM) Agent. This software is on the CIM Agent for ESS CD.

- ► Transmission Control Protocol/Internet Protocol (TCP/IP).

- ► Adobe Acrobat Reader version 4.0 or later.

  You need the Adobe Acrobat Reader to read the License Agreement and product information from the ESS CIM Agent LaunchPad. You can download the Adobe Acrobat Reader from the following Web site:

  http://www.adobe.com/support/downloads/main.html

- ► Vendor storage management products, such as Tivoli SRM.

B

Copy Services feature codes

In this appendix, we summarize the feature codes used to configure and order an F model and 800 model ESS and the available options for the Copy Services functions. We include the options for Copy Services Version 2 (V2) for PPRC and FlashCopy.

479

Copy Services feature codes

The IBM TotalStorage Enterprise Storage Server (ESS) advanced functions enhance the capabilities of the ESS Model 800:

► Parallel Access Volumes (PAV) offer significant performance enhancements in the zSeries and S/390 environments by enabling simultaneous processing for multiple I/O operations to the same logical volume.

► Extended Remote Copy (XRC) is a combined hardware and software business continuance solution for the zSeries and S/390 environments providing asynchronous mirroring between two ESSs at global distances.

► Peer-to-Peer Remote Copy (PPRC) is a hardware-based business continuance solution designed to provide synchronous mirroring between two ESSs that can be located up to 103 km from each other. PPRC is described in detail in Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53. This feature includes the PPRC Extended Distance (PPRC-XD) remote copy function, for non-synchronous mirroring between two ESSs over continental distances (the distance only limited by the network and channel extenders technology capabilities). PPRC V2 for the Fxx and 800 models includes the function of Asynchronous Cascading PPRC. Asynchronous Cascading PPRC is described in detail in Chapter 4, "Peer-to-Peer Remote Copy (PPRC)" on page 53.

► FlashCopy is designed to provide a point-in-time instant copy capability for logical volumes in the ESS. FlashCopy is described in detail in Chapter 3, "FlashCopy" on page 33. FlashCopy V2 includes additional functions:

  – Data Set Level Copy (ZSeries environment only)

  – Multiple relationship FlashCopy

  – Improved Performance/reduced establish time (up to 10 times)

  – Consistency Groups

  – Inband commands over PPRC Links

  – Spanning of LSS boundaries

► FlashCopy V2 is described in detail in Chapter 3, "FlashCopy" on page 33.

► Authorization features, together with the corresponding ESS features for the advanced functions (80xx, 81xx, 82xx, 83xx), allow you to order these functions. Copy Services V2 features for PPRC and FlashCopy advanced functions for the ESS are 185x and 186x for the Fxx and 85xx or 86xx for the model 800.

► PPRC/PPRC-XD and FlashCopy copy services V2 require LIC 2.2.0.

Advanced functions

The advanced functions of the ESS are ordered as optional features, either initially when ordering the ESS, or later as a field upgrade to the ESS.

The advanced function feature codes on the Fxx models are based on the RAID-5 effective capacity. The advanced function feature codes for the model 800 are based on physical capacity.

The initial activation of an advanced function and the installation of a larger license is a concurrent activity. The removal of a license to deactivate an advanced function is a disruptive activity and requires a machine IML.

PPRC and FlashCopy for Exx and Fxx models V1 and V2

Peer-to-Peer Remote Copy (PPRC), PPRC-XD, and FlashCopy V1 are applicable to E10, E20, F10, or F20 environments. The license for PPRC and FlashCopy must be equal to or greater than the total effective RAID-5 capacity of the ESS. Table B-1 lists the Copy Services Feature Codes for V1.

Table B-1   Copy Services Feature Codes V1

Feature Code	Description	Requirement	ESS support
1830 - 1835	FlashCopy	LIC 1.3.0 or later	Models E10, E20, F10, or F20
1820 - 1828	PPRC/PPRC-XD	LIC 1.5.0 or later	Models F20
1830 - 1838	FlashCopy	LIC 1.5.0 or later	Models F20

Peer-to-Peer Remote Copy (PPRC), PPRC-XD, and FlashCopy V2 are applicable to F10 or F20 environments. The license for PPRC and FlashCopy must be equal to or greater than the total effective RAID-5 capacity of the ESS. Table B-2 lists the Copy Services Feature Codes for V2.

Table B-2   Copy Services Feature Codes V2

Feature Code	Description	Requirement	ESS support
1860- 1868	FlashCopy	LIC 2.2.0	Models F10, F20
1850 - 1858	PPRC - PPRC-XD	LIC 2.2.0	Model F10, F20

The implementation of PPRC V1 requires the PPRC feature to be installed on both primary and secondary ESS. Table B-3 lists the Feature Codes for the ESS models showing the different capacities.

Table B-3   PPRC and FlashCopy Feature Codes V1

ESS Effective Capacity	PPRC/PPRC-XD	FlashCopy
	Feature Code	Feature Code
Up to 0.5 TB	1820	1830
Up to 1 TB	1821	1831
Up to 2 TB	1822	1832
Up to 4 TB	1823	1833
Up to 8 TB	1824	1834
Up to 12 TB	1825	1835
Up to 16 TB	1826	1836
Up to 20 TB	1827	1837
Up to 25 TB	1828	1838

The implementation of PPRC V2 requires the PPRC feature to be installed on both the primary and secondary ESS. If used in an Asynchronous Cascading configuration, PPRC must be purchased for the intermediate ESS. Table B-4 on page 482 lists the Feature Codes for the ESS models showing the different ESS capacities.

Table B-4   PPRC and FlashCopy Feature Codes V2

ESS Effective Capacity	PPRC/PPRC-XD	FLASHCOPY
	Feature Code	Feature Code
Up to 0.5 TB	1850	1860
Up to 1 TB	1851	1861
Up to 2 TB	1852	1862
Up to 4 TB	1853	1863
Up to 8 TB	1854	1864
Up to 12 TB	1855	1865
Up to 16 TB	1856	1866
Up to 20 TB	1857	1867
Up to 25 TB	1858	1868

FlashCopy and PPRC, PPRC-XD V1 183x, 182x can be upgraded to FlashCopy and PPRC, PPRC-XD V2 186x,185x using feature exchange. The FlashCopy and PPRC, PPRC-XD V1 to FlashCopy, and PPRC, PPRC-XD V2 upgrade must be for a license of the same or greater capacity. A feature exchange of a lower capacity level (a numerically lower feature number) is not supported. FlashCopy and PPRC, PPRC-XD V2 (186x, 185x) to FlashCopy, PPRC, PPRC-XD V1 (183x, 182x) downgrades are not supported.

ESS advanced functions for the ESS Model 800

The ESS Function Authorization feature numbers provide a set of pricing tiers for the ESS advanced functions V1. These tiers provide increased granularity (as compared to earlier models) with pricing matched to the physical capacity of the ESS Model 800. These are shown in Table B-5.

Table B-5   ESS Function Authorization features V1

Advanced Function	IBM 2240 ESS Function Authorization	
	Machine type and model	Feature Codes
PAV	2240-PAV	8000 to 8012
XRC	2240-XRC	8100 to 8112
PPRC	2240-PRC	8200 to 8215
FlashCopy	2240-FLC	8300 to 8315

Deactivation is supported for PPRC, FlashCopy V1 8099, 8199,8299, 8399 (0 TB Disable) and can be exchanged for any 80xx, 8199, 82xx, or 83xx feature. No refund credit is provided to the customer.

The ESS Function Authorization feature numbers provide a set of pricing tiers for the ESS advanced functions V2. These tiers provide increased granularity (as compared to earlier models) with pricing matched to the physical capacity of the ESS Model 800. These are shown in Table B-6 on page 483.

Table B-6   ESS Function Authorization features V2

Advanced Functions	IBM 2240 ESS Function Authorization	
	Machine type and model	feature codes
PPRC	2240-PRC	8500 to 8515
FlashCopy	2240-FLC	8600 to 8615

The IBM 2240 ESS Function Authorization V1 feature numbers are for billing purposes only and authorize the use of ESS advanced functions at a given capacity level on a specific ESS Model 800. These are shown in Table B-7.

Table B-7   ESS Function Authorization: Capacity tiers and features V1

Physical capacity tier	PPRC/PPRC-XD	FlashCopy
	2240-PRC	2240-FLC
Up to 1 TB	8200	8300
Up to 2 TB	8201	8301
Up to 3 TB	8202	8302
Up to 4 TB	8203	8303
Up to 5 TB	8204	8304
Up to 6 TB	8205	8305
Up to 8 TB	8206	8306
Up to 10 TB	8207	8307
Up to 12 TB	8208	8308
Up to 16 TB	8209	8309
Up to 20 TB	8210	8310
Up to 25 TB	8211	8311
Up to 30 TB	8212	8312
Up to 40 TB	8213	8313
Up to 50 TB	8214	8314
up to 60 TB	8215	8315

The IBM 2240 ESS Function Authorization feature numbers V2 are for billing purposes only and authorize the use of ESS advanced functions at a given capacity level on a specific ESS Model 800. These are shown in Table B-8.

Table B-8   ESS Function Authorization: Capacity tiers and features V2

Advanced Function	PPRC/PPRC-XD	FlashCopy
	2240-PRC	2240FLC
Up to 1 TB	8500	8600
Up to 2 TB	8501	8601
Up to 3 TB	8502	8602

Advanced Function	PPRC/PPRC-XD	FlashCopy
	2240-PRC	**2240FLC**
Up to 4 TB	8503	8603
Up to 5 TB	8504	8604
Up to 6 TB	8505	8605
Up to 8 TB	8506	8606
Up to 10 TB	8507	8607
Up to 12 TB	8508	8608
Up to 16 TB	8509	8609
Up t0 20 TB	8510	8620
Up to 25 TB	8511	8611
Up to 30 TB	8512	8612
Up to 40 TB	8513	8613
Up to 50 TB	8514	8614
Up to 60 TB	8515	8615

Capacity tier calculation

The ESS advanced functions are enabled and authorized based upon the physical capacity of the ESS:

► PPRC and FlashCopy enabling and authorization must be equal to or greater than the total physical capacity of the ESS.

Ordering advanced functions

The advanced functions require the selection of IBM 2105 Model 800 feature numbers and the purchase of the matching IBM 2240 ESS Function Authorization feature numbers:

► The ESS Model 800 feature numbers (80xx, 81xx, 82xx, 83xx, 85xx, and 86xx) enable a given function on the ESS at a given capacity level.

► The ESS Function Authorization feature numbers (80xx, 81xx, 82xx, 83xx, 85xx, and 86xx) authorize use of the given Advanced Function at the given capacity level on the ESS machine for which it was purchased.

The ESS Model 800 feature numbers (8xxx) and the ESS Function Authorization feature numbers (8xxx) V1 are co-requisites and must always correspond to one another (these are shown in Table B-9).

Table B-9   ESS Model 800 and ESS Function Authorization features correspondence V1

ESS Advanced Function	ESS Model 800 feature	IBM 2240 ESS Function Authorization feature
Peer-to-Peer Remote Copy PPRC and PPRC-XD	2105-800 features 82xx	2240-PRC features 82xx
FlashCopy	2105-800 features 83xx	2240-FLC features 83xx

The ESS Model 800 feature numbers (8xxx) and the ESS Function Authorization feature numbers (8xxx) V2 are co-requisites and must always correspond to one another (these are shown in Table B-10).

Table B-10   ESS Model 800 and ESS Function Authorization features correspondence V2

ESS Advanced Function	ESS Model 800 features	IBM 2240 ESS Function Authorization feature
Peer-to Peer Remote Copy PPRC and PPRC-XD	2105-800 features 85xx	2240-PRC features 85xx
FlashCopy	2105-800 features 86xx	2240-FLC features 86xx

The ESS Function Authorizations (IBM 2240 FLC, PAV, PRC, and XRC) V1 must be ordered with the ESS Model 800. The ESS Model 800 order must include the specification of the matching 8xxx Advanced Function feature.

The ESS Function Authorization and ESS Model 800 capacity tiers V1 feature numbers are similar and are shown in Table B-11.

Table B-11   ESS Model 800: Capacity tiers and features V1

Physical Capacity tier	ESS Model 800	ESS Model 800
	Peer-to-Peer Remote Copy / XD	FlashCopy
Up to 1 TB	8200	8300
Up to 2 TB	8201	8301
Up to 3 TB	8202	8302
Up to 4 TB	8203	8303
Up to 5 TB	8204	8304
Up to 6 TB	8205	8305
Up to 8 TB	8206	8306
Up to 10 TB	8207	8307
Up to 12 TB	8208	8308
Up to 16 TB	8209	8309
Up to 20 TB	8210	8310
Up to 25 TB	8211	8311
Up to 30 TB	8212	8312
Up t0 40 TB	8213	8313
Up to 50 TB	8214	8314
Up to 60 TB	8215	8315

The ESS Function Authorizations (IBM 2240 FLC, PAV, PRC, and XRC) V2 must be ordered with the ESS Model 800. The ESS Model 800 order must include the specification of the matching 8xxx Advanced Function feature.

The ESS Function Authorization and ESS Model 800 capacity tiers V2 feature numbers are similar and are shown in Table B-12.

Table B-12   ESS Model 800: Capacity tiers and features V2

Physical Capacity tier	ESS Model 800	FlashCopy
	Peer to Peer Remote Copy/XD	FlashCopy
Up to 1 TB	8500	8600
Up to 2 TB	8501	8601
Up to 3 TB	8502	8602
Up to 4 TB	8503	8603
Up to 5 TB	8504	8604
Up to 6 TB	8505	8605
Up to 8 TB	8506	8606
Up to 10 TB	8507	8607
Up to 12 TB	8508	8608
Up to 16 TB	8509	8609
Up to 20 TB	8510	8610
Up to 25 TB	8511	8611
Up to 30 TB	8512	8612
Up to 40 TB	8513	8613
Up to 50 TB	8514	8614
Up to 60 TB	8515	8615

PPRC over Fibre Channel links

PPRC over Fibre Channel is supported on the ESS Model 800, and is provided with PPRC Version 2. PPRC Version 2 is an optional feature for the ESS (#85xx). PPRC over Fibre Channel also requires that the ESS is at LIC level 2.3.0.

C
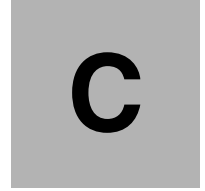
Logical subsystems

In this appendix, we explain the logical functions of a storage controller that allow one or more host I/O interfaces to access a set of devices.

The controller aggregates the devices according to the addressing mechanisms of the associated I/O interfaces. One or more logical subsystems exist on a storage controller. In general, the controller associates a given set of devices with only one logical subsystem.

Device adapters and logical subsystems

The logical subsystem (LSS) is a logical structure that is internal to the IBM TotalStorage Enterprise Storage Server and is used for the configuration of the ESS. Although it relates directly to the logical control unit (LCU) concept of the ESCON and FICON architectures, it does not directly relate to SCSI and FCP addressing.

The device adapter (DA) to LSS mapping is a fixed relationship. Each DA supports two loops, and each loop supports two CKD logical subsystems and two FB logical subsystems (one from each cluster). So a DA pair supports four CKD LSSs and four FB LSSs.

An ESS can have zero, eight, or 16 LSSs defined for open systems storage (FB logical subsystems). An ESS having zero FB logical subsystems would be one with only CKD LSSs and used by the zSeries servers.

When all eight loops have the capacity installed, then there are up to 16 CKD LSSs and up to 16 FB LSSs available to support the maximum of 48 RAID ranks or groups of eight JBOD ranks. Each LSS supports up to 256 logical devices (each logical device is mapped to a logical volume in the RAID ranks or JBOD ranks).
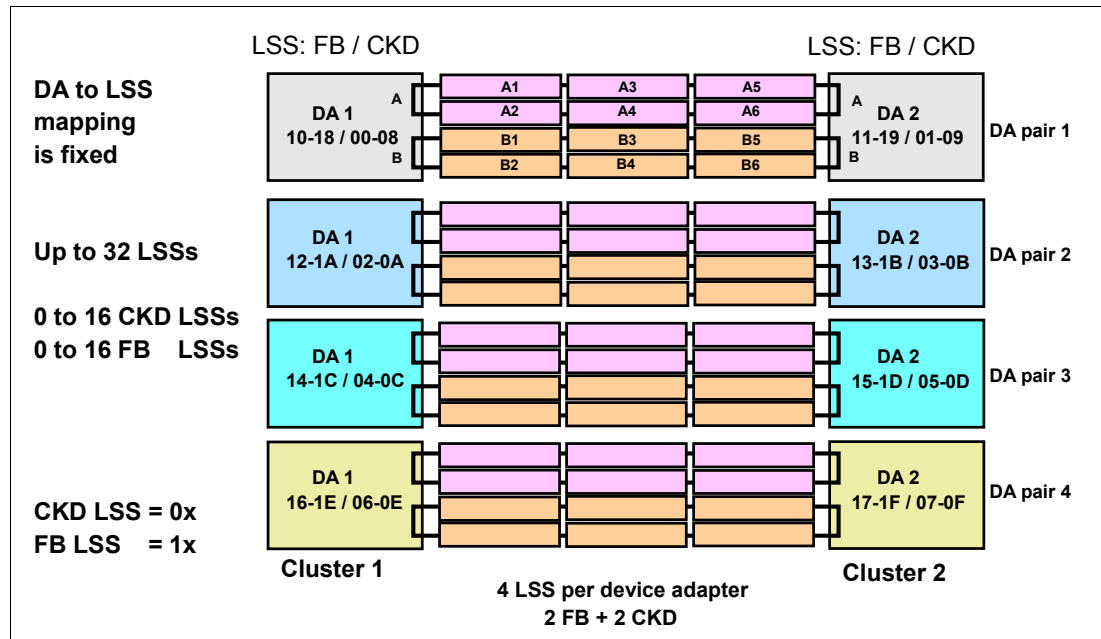
Figure C-1   Logical subsystems

The numbering of the logical subsystems indicates the type of LSS. CKD logical subsystems are numbered x'00' to x'0F' and the FB logical subsystems are numbered x'10' to x'1F' (see Figure C-1). For the CKD host view (that is, the zSeries server), a logical subsystem is also mapped one-to-one to a logical control unit.

As part of the configuration process, you can define the maximum number of logical subsystems of each type you plan to use. For open systems storage, it is the only option to choose while configuring logical subsystems (you cannot, for example, choose in which LSS a RAID rank will be). If you plan to use the ESS only for open-systems data, then you can set the number of CKD LSSs to 0. But you must also remember that going from eight to 16 LSSs is disruptive, so you should decide in advance how many you will need. Also, you should remember that there can be only 256 logical volumes in each of the LSS. So if you, for

example, use only eight FB LSS, then you have a limit of 2048 FB logical volumes for the whole ESS.

Logical subsystem (LSS)

Let us now look at how the RAID and JBOD ranks are used within a logical subsystem (LSS). Logical subsystems are related to the device adapter SSA loops, and are therefore managed by only one cluster. An LSS belongs to one DA.

Each loop supports from 16 to 48 disk drives. The minimum is 16 for RAID ranks, because we always need two spare disk drives on any loop with RAID ranks, so the minimum is two six+P+S ranks. The disk drives are configured into two to six RAID ranks, either CKD or FB, or two to six JBOD groups, or a combination of the two types. If, for example, all 48 disk drives on the loop were JBOD, we would have 48 JBOD ranks, each of which would be CKD or FB.

As part of the configuration process, each rank is assigned to one logical subsystem (automatically; it depends on the physical position of the rank in the ESS). This LSS is either CKD or FB. Two ranks, from two different loops of the same DA pair, can be associated to build up an LSS (in the open system environment, this is possible only if you configure the ESS to use eight FB LSSs, so there will be one LSS per each DA).

Example

In the example shown in Figure C-2, we have the maximum of 48 disk drives installed on both loops. We plan to map six groups of disk drives onto four LSSs.
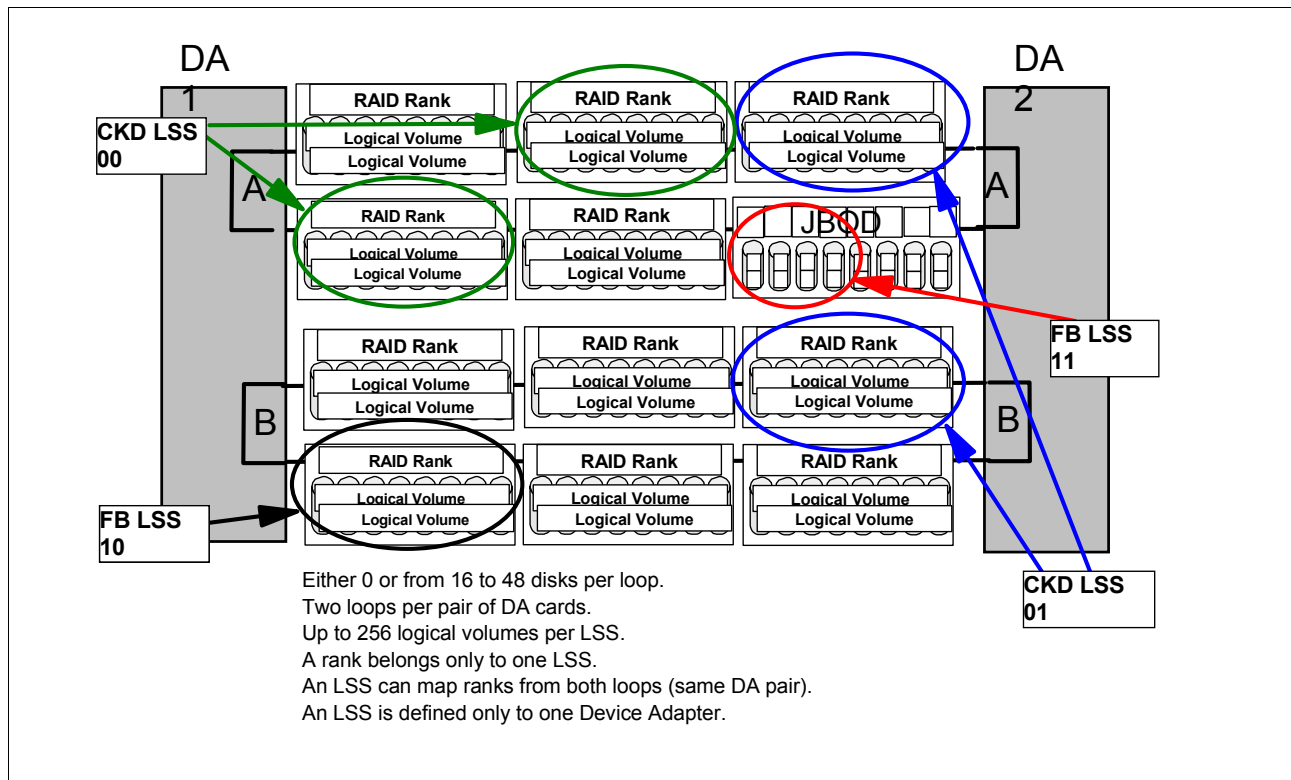
Figure C-2 Sample of logical subsystem configuration

Five RAID ranks are defined and one group of eight JBODs. We have four logical subsystems available on the DA (2 CKD and 2 FB LSSs). If we assume that this is the first DA pair in an ESS, then we can also associate the logical subsystem numbers.

As a consequence of the LSS mapping algorithms, here is the LSS definition:

1. DA1 Loop A LSS(00) CKD: Two RAID ranks with a total 16 disk drives.
2. DA1 Loop B LSS(10) FB: One RAID rank with eight disk drives.
3. DA 2 Loop A LSS(11) FB: Four JBOD disk drives (ranks) formatted for FB use.
4. DA 2 Loops A and B LSS(01) CKD: Two RAID ranks from two different loops are associated, with a total of 16 disk drives.

> **Note:** Remember, there are still six ranks to be mapped to the existing LSS.

LSS mapping of ranks on more than one loop (of the same DA pair)

For an LSS, allocating ranks from different loops (from the same DA pair) could be useful, especially for FlashCopy Version 1, so you can establish a FlashCopy pair between two different SSA loops. With FlashCopy Version 2, this is not needed anymore, as the source and the target volumes do not have to be in the same LSS. However, you must remember that all the capacity that an LSS is able to see must be mapped with no more than 256 logical volumes. 256 is the maximum number of logical volumes that can be defined for an LSS.

Host mapping to logical subsystem

For the zSeries servers, the data stored in the ESS is arranged in a count-key-data (CKD) format, and this data is retrieved with the I/O operations from the host servers. These I/O operations, for the zSeries servers, are done according to the ESCON or FICON architectures.

For the open system servers, the data stored in the ESS is arranged in a fixed-block (FB) format, and this data is retrieved with the I/O operations from the host servers. These I/O operations, for the open system servers, are done according to the SCSI or Fibre Channel architectures.

Additionally, the ESS has its own logical view of the stored data, based on the logical storage subsystem (LSS) definitions.

Fixed block considerations

- Each ESS can have up to 16 FB LSSs and 4096 LUNs
- SCSI:
  - Host device addressing is target/LUN on each bus.
  - Maximum of 15 targets per bus and 64 LUNs per target.
  - Target on bus is associated with a single LSS.
  - LUNs are associated to a specific logical volume on LSS.
  - Specific logical volume can be shared using a different target/LUN association.
  - LUN masking.
- Fiber Channel Protocol:
  - You must first define host N-port World Wide Port Name (WWPN) to the ESS.

– You can have either 256 or 4096 LUNs per host N-port, depending on the LUN addressing mode supported by the host operating system.

– You can restrict the host to access the ESS using only selected Host Adapters, or you can use all available Fibre Channel Host Adapters.

– Two access modes:

• Access_Restricted (use LUN masking)

• Access_Any

For detailed information about the ESS logical subsystems an logical configuration, refer to the redbook *IBM TotalStorage Enterprise Storage Server Model 800,* SG24-6424.

D

System Adapter ID (SAID)

In this appendix, we provide some additional explanation for System Adapter IDs (SAIDs).

SAID definition with ESCON links

For PPRC primary to secondary unit (channel to control unit) communication, a maximum of eight ESCON links using a modified ESCON protocol can be configured.

ESCON channels provide 160 Mbps point-to-point links. While PPRC can be bi-directional, these links are uni-directional. The primary unit ESCON port (the one in channel mode) has to be dedicated for PPRC. The ESCON port on the secondary unit can be also used for S/390 host attachment, provided an ESCON director is used and the host is connected to it.

The ESCON adapters are specified by their System Adapter ID (SAID). To alleviate some of the confusion regarding these terms, we will now explain adapter IDs: SA IDs and SA tags.

Adapter IDs are internal identifications assigned to each external port on the machine and are used for communication between the adapters and the SMP. They are also the adapter numbers we use in the state save formatter to format data for a particular adapter.

An SA Tag is a device used to find the physical location of an port in the machine. It is part of the ESCON node descriptor and is used for link-fault isolation. It is also used for identifying a path through which a PPRC path is established. The functional specification also uses the term SA ID to refer to the tag. Recently, a new definition of the SA tag was approved in the functional specification, and code will soon be promoted to switch to the new format. This will affect each and every friend chain or TSO command or Web interface that issues an establish PPRC path subsystem function.

Figure D-1 shows the how the adapter IDs and tags map to physical ESCON port locations.

Figure D-1   Physical ESCON port location

PPRC paths using ESCON links

To use PPRC, you must have ESCON host adapter ports on both the primary and secondary ESS for the peer-to-peer links. An ESCON path can either be operating as a control unit or a PPRC channel. For highest availability, IBM recommends that each ESS should have more

than one ESCON adapter for PPRC connectivity. See the SAID numbers in Figure D-2 on page 495.

Before PPRC pairs can be established, logical paths must be defined between the logical control unit images. The ESS supports up to 16 CKD logical control unit images and up to 16 FB controller images. An ESCON adapter supports up to 64 logical paths. A pair of LSSs can be connected with up to eight logical paths. You establish logical paths between control unit images of the same type over physical ESCON links

A path is used to send data between the source and target of PPRC pairs. The physical path consists of the ESCON connection between two Enterprise Storage Servers while a logical path describes the connection of the PPRC source and targets. There could be multiple logical paths established over a single physical path.

Figure D-2   SAID numbers

An ESCON port operating as a PPRC channel can support logical paths between many control-unit images whether it connects directly to another server or goes through a switch. The ESCON port has a maximum of 64 PPRC channels. A primary control-unit image can have a relationship with up to four secondary logical control-unit images. Each control-unit image pair can operate over a maximum of eight PPRC channels.

To initiate a PPRC operation through the ESSNet for a logical device, both the primary and secondary ESSs must be attached to the same ESSNet. Both the primary and secondary ESSs must also be in the same ESS Copy Services server domain. If the two ESSs are not on the same local ESSNet, you must extend the ESSNet through a local area network (LAN) to connect the two local ESSNets.

SAID usage example

In Figure D-3 on page 496, the two ESS subsystems have been configured with identical configurations, and the two ESS subsystems have been connected with two ESCON links. To establish a path from the primary ESS on the left to the secondary ESS on the right, the user must determine the connection of the ESCON cables and record which system adapter IDs to connect the controller.

Figure D-3   Determine pathing to remote ESS

Using the chart, the path between two ESS subsystem can be determined. The user can then select the appropriate SAID, connect it to the remote ESS, and access the ESS that contains volumes that will be used in the PPRC pairing.

FCP System Adapter ID

The Fiber Channel ports used for PPRC links are also specified by their *System Adapter* ID (SAID).
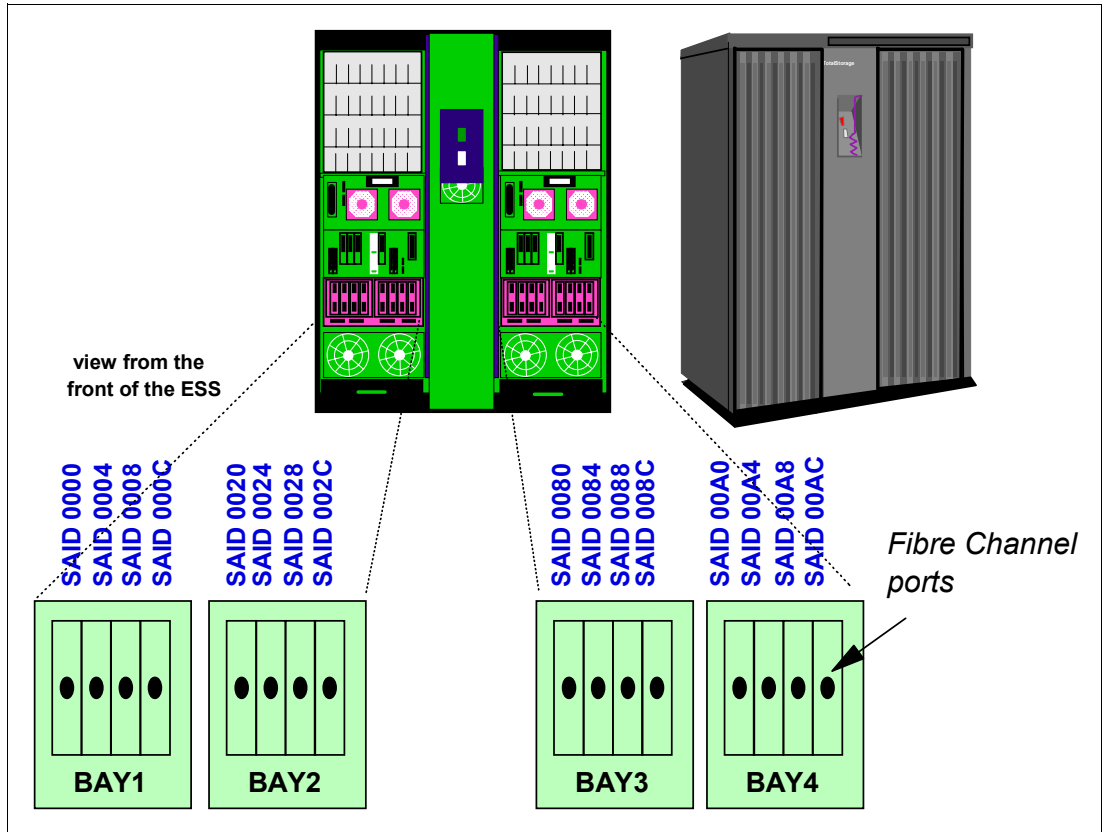
Figure D-4 System Adapter ID (SAID) identification with Fibre Channel ports

Figure D-4 illustrates how the system adapter IDs (SAIDs) are designated when the host adapter bays are seen from the front of the ESS. The first two values are zeros, and the last two values are the hexadecimal SAID byte values that correspond to the Fiber Channel port interface of the ESS.

Sample naming convention

This appendix contains an example of a task naming convention.

Task names convention considerations

The number of Copy Services tasks that are required can grow fairly large, even in a relatively small installation. It is essential to think carefully and design your naming convention before creating the tasks.

Each time you create a task, you are prompted to provide a task name and a task description.
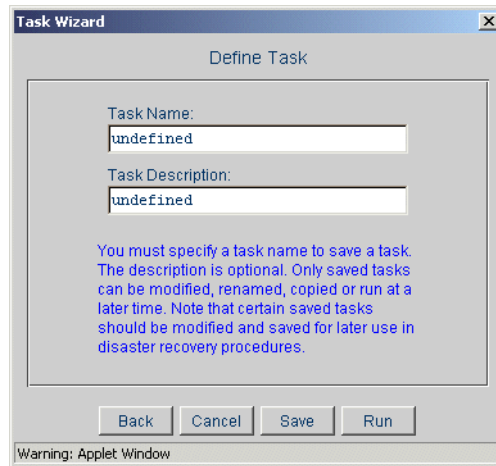
Figure E-1   Task Name and Description prompt window

Task name

You will use the task name as a reference to each task you create. A task name must be unique. The Web User Interface will not allow you to create two tasks with the same name. A task can contain up to 16 characters. Blanks and special characters are not allowed. The only special characters you can use are "-" and "_".

The task name appears in the "task" window in the graphical user interface. The log files will display the task name for each failing task. And you will use the task name to invoke the task and to list the tasks using the CLI.

When you will build your own naming convention for task names, you should consider the following:

► The task name should contain a reference to the type of task you created. There are four main task families:

– Tasks for Path management

– Tasks for PPRC and PPRC-XD management

– Tasks for FlashCopy management

– Tasks that are grouping other tasks

► The task name should also contain a reference to the option(s) you chose for the task.

► The task name should also contain a clear reference to the volumes or the LSS(s) impacted by the task.

► You may also want to include a reference to the servers, the sites, the applications, or the processes you built the task for.

Task description

You will use the task description to give a detailed description of the task. The task description accepts blanks and special characters. The task description field can accept more than 300 characters, but only the first 46 characters will be displayed in the task window in the graphical user interface.

Naming convention example

Here is the task naming convention we used during the residency. The task names in this redbook are following this convention.

> **Attention:** The task naming convention presented here is an example. We are not proposing here a standard naming convention. It may not fully satisfy all the customer's needs.

PPRC operations

We used the following syntax for PPRC, PPRC-XD task names:

<reference:4 char>_<PPRC op:2 Char>_<source volume:4 Char><target volume:4 Char>

For example: AIX1_PF_s302t512

Where:

<reference>	Host, cluster, site, application or process reference (four characters).
	We used four characters to give the task a reference to a system, a site, an application, or a process. This reference will help when trying to identify the task purpose in the task description.
	Example: AIX1, CLUA, BKP2.
<PPRC op>	Operation code (two characters).
	The operation code will identify the type of task and some options. These two characters will follow the rules described below.
<source volume>	Source volume identification (four characters).
	The four characters will contain a clear identification of the primary volume for PPRC or the source volume for FlashCopy tasks. We use the first three characters of the volume serial number given by the ESS (for example: volume 706).
	We also add a letter in front of the volume number that makes the task name more readable. This letter can also be a reference to an ESS for PPRC.
	Example 1: s706 (source volume number 706). Example 2: A501 (PPRC primary volume number 501 in ESS A).
<target volume>	Source volume identification (four characters).
	The four characters will contain a clear identification of the secondary

volume for PPRC or the source volume for FlashCopy tasks. We use the first three characters of the volume serial number given by the ESS (for example: volume 423).

We also add a letter in front of the volume number that makes the task name more readable. This letter can also be a reference to an ESS for PPRC task.

Example 1: t423 (target volume number 423).
Example 2: C811 (PPRC primary volume number 811 in ESS A).
Example 3: I342 (PPRC intermediate volume number 342 in an Asynchronous Cascading PPRC).

Additional information

This naming convention uses the first four characters to help the administrators to quickly identify the reasons for the task by associating the task to a host, a cluster, a site, and so on.

In relatively small implementations with a lot of hosts, it will be appropriate to use these characters to refer to hosts or the cluster. On the other hand, it will be appropriate to use these four characters to refer to the ESS and/or sites in environments where a large number of ESSs serve storage to the same host or cluster.

FlashCopy operations

We used the following syntax for FlashCopy task names:

<reference:4 Char>_<FlashCopy op:2 Char>_<ESS ref:1 Char><source volume:3 Char><target volume: 3 Char>

Example: AIX1_PF_s302t512.

Where:

<reference>	Host, cluster, site, application or process reference (four characters).
	We used four characters to give the task a reference to a system, a site, an application, or a process. This reference will help when trying to identify purpose of the task in the task description.
	Example: AIX1, CLUA, BKP2.
<flashcopy op>	Operation code (three characters).
	The operation code will identify the type of task and some options. These two characters will follow the rules described below.
<ESS ref>	Command Management (one character).
	This character will refer to the ESS impacted by the command.
<source volume>	Source volume identification (four characters).
	The four characters will contain a clear identification of the primary volume for PPRC or the source volume for FlashCopy tasks. We use the first three characters of the volume serial number given by the ESS (for example: volume 706).
	Example 1: 706 (source volume number 706).

<target volume>	Source volume identification (four characters).

The four characters will contain a clear identification of the secondary volume for PPRC or the source volume for FlashCopy tasks. We use the first three characters of the volume serial number given by the ESS (for example: volume 423)

We also add a letter in front of the volume number that makes the task name more readable. This letter can also be a reference to an ESS for PPRC task.

Example 1: t423 (target volume number 423).
Example 2: C811 (PPRC primary volume number 811 in ESS A).
Example 3: I342 (PPRC intermediate volume number 342 in an Asynchronous Cascaded PPRC).

Path operations

We used the following syntax for path operations:

<source LSS: 3 Char><target LSS:3 Char>_<path op:2 Char>_<links:6 Char>

Where:

<source LSS>   Source LSS (three characters).

We use three characters to clearly identify the source LSS for the Path. We use here one character to identify the ESS and two characters to identify the LSS.

Example 1: A12 (LSS number 12 in ESS A).
Example 2: L03 (LSS number 03 in Local Site ESS).
Example 3: I0C (LSS number OF in Intermediate ESS).

<target LSS>   Target LSS (three characters).

We use three characters to clearly identify the target LSS for the Path. We use here one character to identify the ESS and two characters to identify the LSS.

Example 1: B12 (LSS number 12 in ESS B).
Example 2: R04 (LSS number 04 in Remote Site ESS).
Example 2: I0A (LSS number OA in Intermediate ESS).

<path op>   Copy operation code (two characters).

The operation code will identify the type of path task and some options. Those two characters will follow the rules described below.

<links>   Links used (number of links) or SAID numbers (up to six characters).

Within six characters, we give a description of the physical configuration of this task. It could be the number of links, or the SAID numbers.

Example 1: 4links.
Example 2: 0104 (refers to SAID 8001 and 8004 on the source LSS).

LSS operations

We used the following syntax for tasks defined at the LSS level. Be aware that those tasks may affect several volume pairs.

<reference:4 Char>_<LSS op:3 Char>_<LSS1:3 Char><LSS2:3 Char>

Where:

<reference>   Host, cluster, site, application or process reference (five characters).

We used four characters to give the task a reference to a system, a site, an application or a process. This reference will help when trying to identify purpose of the task in the task description.

Example: AIX1, CLUA, BKP2

<LSS op>   Operation code (three characters).

The operation code will identify the type of task and the associated options. These three characters will follow the rules described below for LSS operations.

<LSS1>   Source LSS (three characters).

We use three characters to clearly identify the source LSS for the task. We use here one character to identify the ESS and two characters to identify the LSS.

Example 1: A12 (LSS number 12 in ESS A).
Example 2: L03 (LSS number 03 in Local Site ESS).
Example 3: I0C (LSS number OF in Intermediate ESS).

<LSS2>   Target LSS (three characters).

We use three characters to clearly identify the target LSS for the task. We use here one character to identify the ESS and two characters to identify the LSS.

Example 1: B12 (LSS number 12 in ESS B).
Example 2: R04 (LSS number 04 in Remote Site ESS).
Example 3: I0A (LSS number OA in Intermediate ESS).

> **Note:** If the task you define is a Consistency Created for FlashCopy consistency group, then <LSS2> is useless, since this option is available only if one LSS is selected.

Task Created using multiple volume selection

We used the following syntax for tasks running on multiple volumes:

<reference:4 Char>_<op:3 or 4Char>_<Details:6 Char>

Where:

<reference>   Host, cluster, site, application, or process reference (4 characters).

We use 4 characters to give the task a reference to a system, a site, an application, or a process. This reference will help when trying to

identify the purpose of the task in the task description.

Example: AIX1, CLUA, BKP2.

\<op>   Operation code (two or three characters).

The operation code will identify the type of task and the associated options. The operation code will start with a G, then the next characters will follow the rules described below for PPRC or FlashCopy operations.

\<details>   Details (six characters).

We use six characters to give more details on the grouped task. For example, it might be used to refer to the number of volumes involved.

Example: 4vol.

Group operations

We used the following syntax for tasks grouping several tasks. You can group tasks in the task panel.

We we also use this convention for tasks involving several volumes. You can create such tasks either by using the multiple volume selection mode in the volume panel or building tasks at the LSS level in the LSS panel.

<group reference:4 Char>_GR_<group op:2or3 Char>_<details:4 Char>

Where:

\<group reference>   Host, cluster, site, application, or process reference (four characters).

We use four characters to give the group of tasks a reference to a system, a site, an application, or a process. This reference will help when trying to identify the purpose of the task in the task description.

Example: CLU1, DIS1.

\<group op>   Operation code (two or three characters).

This operation code will identify the type of task you have grouped and the associated options. This <group op> corresponds to the operations code for PPRC, PPRC-XD, FlashCopy, Paths, or tasks defined at the LSS level as described above.

\<details>   Details (four characters).

We use four characters to give more details on the grouped task. For example, it might be used to refer to the ESS(s) or site(s).

Example: AtoB.

Operation codes

We propose here a convention based on a two or three character operation code that will clearly identify the task actions and the options associated with them.

Operation codes for PPRC and PPRC-XD

Table E-1   Operation codes for PPRC and PPRC-XD

1st Char.	2nd Char.	Description
P		PPRC synchronous.
X		PPRC-XD.
	F	Establish PPRC task with "Full copy" option.
	M	Establish PPRC with "Full copy Option" option and Suspend the PPRC Pair after establish is complete. (M will stand for "Migration".)
	O	Establish PPRC with "copy Out-of -Sync. cylinders only" option.
	D	Establish PPRC with "copy Out-of -Sync. cylinders only" option and Suspend pair after Establish is complete. (D will stand for Differential migration.)
	N	Establish PPRC with "No copy" option.
	R	Establish PPRC task with Failover option.
	K	Establish PPRC task with Failback option.
	B	Establish PPRC task with Failback option and suspend the Pair after establish is complete.
	H	Suspend PPRC Pair. (H stands for Hibernate.)
	C	Asynchronous Cascading PPRC.
T		Terminate PPRC Pair.
	S	Terminate PPRC Pair from source.
	T	Terminate PPRC Pair from target.

FlashCopy operations

Table E-2   Operation codes for FlashCopy operations

1st Char.	2nd Char.	3rd Char	Description
C			Establish FlashCopy with copy option
N			Establish FlashCopy with nocopy option
	S		Establish FlashCopy with Start change recording Option
	A		Accelerate destage mode option
	F		Freeze copy option selected
	T		Inhibits write to target (for S/390 only)
	I		Incremental option and Start change recording selected
	P		Persistent FlashCopy selected

1st Char.	2nd Char.	3rd Char	Description
	R		Reverse restore and Start change recording option selected
	_		No other options
W			Withdraw the FlashCopy relation
	W		Withdraw to target
	B		Start background FlashCopy
		I	Inband command option
		_	No other options

Path management

Table E-3   Operation code for path management

1st Char.	2nd Char.	Description
E		Path management.
	N	Establish Path with Do Not remove existing path option.
	F	Establish path with Force removal option.
	C	Establish path with PPRC consistency group option and Do not Remove existing Path option.
	A	Establish the path with the PPRC consistency group option and the Force Removal option. (A might stand for "All" options.)
R		This first character will refer to remove path tasks.
	N	Remove path with Do not remove existing path option.
	F	Remove path with Force removal option.

Multiple volume selection, and LSS operations

Table E-4   Operation codes for task defined at the LSS level

1st Char	2nd & 3rd Char	Description
G		This first character will refer to tasks defined at the LSS level.
	<op>	Any operation code as referred above (PPRC, PPRC-XD, FlashCopy, or Path).
	LP	Consistency created for the LSS(s) for PPRC consistency group.
	LF	Consistency created for FlashCopy consistency groups.
	FP	Freeze PPRC volumes.

Examples

AIX1_PF_A401Z701   This task will establish a PPRC relationship, with full copy option, between the primary volume number 401 in the ESS referenced as ESS A and the secondary volume number 701 in the ESS

referenced as ESS Z. The volumes are dedicated to applications owned by the system called AIX1.

PROD_XN_I301R301 This task name describes a task that will establish a PPRC-XD relationship, with the No copy option between the volume number 301 in the ESS in the Intermediate site to the secondary volume number 301 in the Remote ESS. The volumes are dedicated to the production environment. This task is part of an Asynchronous Cascading PPRC.

A12B13_EF_4links This task will establish a path between LSS number 12 on the ESS A and the LSS number 13 in ESS B. The task will establish the path after removing the existing path (if any). The path will use four physical links.

PROD_GWW__A12_4v This task will withdraw the FlashCopy relationships from the target. The comment helps you to understand that it involves four volumes in the LSS12 in the ESS A.

F

AIX and FlashCopy for backups

In this appendix, we discuss the preparations that are needed on an AIX system before FlashCopy can be used as part of an integrated backup solution. This is based on the outline given in 8.1, "AIX specifics" on page 328.

509

Definition of a typical backup solution

There are different understandings about what a "typical" backup solution is. The scenario that we tested is described in detail together with the sample scripts that we used. While these scripts worked in our environment, there is no guarantee that they will work in other situations. They do, however, provide a basis on which you can develop your own backup scripts.

One RS/6000 server running the AIX operating system was attached to an IBM 2105 ESS with a storage pool (LUNs) assigned to it. This server is known as the production server, on which the production applications would be run. The LUNs of this server are the source LUNs with respect to FlashCopy. The target LUNs are located on the same ESS, but are assigned to a different RS/6000 server running the AIX operating system, which is called the Backup server.

We recommend that all I/O to the FlashCopy source volumes be quiesced and that all buffers are flushed before the copy is made. In the case of a file system, this would mean issuing the `sync` command or unmounting the file system. Databases usually have their own commands to ensure that all pending writes are flushed to disk.

Preparation after the first FlashCopy

After all the production volumes have been FlashCopied, you have to make the FlashCopy target volumes available to the Backup server. Ensure that they are all varied online, and that no volume group is created over them. As this is the first FlashCopy, there is no information about the Volume Group structure in the ODM of the Backup server. This means that you have to import the volume group to the Backup server. This is done by using the procedure described below. Since the FlashCopy copies the entire LUN, it will also copy the physical volume identifier (PVID) of the source volume. The PVID is an unique identifier that AIX uses to identify its physical disks.

1. Identify the PVID of at least one hdisk from each volume group on the production server that you intend to FlashCopy.

   To determine the PVID of an hdisk in volume group *samplevg*, issue the command `lsvg -p samplevg`. This will give the output shown in Table F-1.

Table F-1   Contents of samplevg

PV_NAME	PV STATE	TOTAL PPs	FREE PPs	FREE DISTRIBUTION
hdisk26	active	953	2	00..00..00..00..02
hdisk27	active	953	0	00..00..00..00..00
hdisk28	active	953	0	00..00..00..00..00
hdisk29	active	953	0	00..00..00..00..00
hdisk30	active	953	0	00..00..00..00..01
hdisk31	active	953	0	00..00..00..00..00
hdisk32	active	953	0	00..00..00..00..00
hdisk33	active	953	0	00..00..00..00..00

2. Query one of the hdisks by issuing the **lspv** command. For example, if we wish to use hdisk26, the command that is issued would be **lspv hdisk26**. The output from this command is shown in Figure F-1.

```
PHYSICAL VOLUME:    hdisk26                    VOLUME GROUP:      samplevg
PV IDENTIFIER:      000c3c7d5c13c183   VG IDENTIFIER    000c3c7d5c13dd48
PV STATE:           active
STALE PARTITIONS:   0                          ALLOCATABLE:       yes
PP SIZE:            8 megabyte(s)              LOGICAL VOLUMES:   1
TOTAL PPs:          953 (7624 megabytes)       VG DESCRIPTORS:    1
FREE PPs:           2 (16 megabytes)
USED PPs:           951 (7608 megabytes)
FREE DISTRIBUTION:  00..00..00..00..02
USED DISTRIBUTION:  191..191..190..190..189
```

Figure F-1   lspv of hdisk26

The PV Identifier is the required PVID. Alternatively, the command **lspv | grep hdisk26** can be issued, which will produce the following output:

```
hdisk26          000c3c7d5c13c183      samplevg
```

The second column is the PVID. On the Backup server you now have to determine which hdisk correspondents to this PVID. If the target volume has not been configured as a hdisk into AIX using **cfgmgr**, then issue the commands:

```
cfgmgr -v -l <host_bus_adapter>
lspv | grep 000c3c7d5c13c183
```

If you have configured the target volume onto the Backup server as an hdisk prior to performing the FlashCopy, there will be no PVID on the hdisk, if it is a brand new configured LUN in the ESS. So the **lspv** output would be:

```
hdisk16 none                None
```

If the **lspv** command output after the FlashCopy has been established does not show its PVID, then you have to identify the hdisk that is the FlashCopy target. The way to do this is to obtain the serial number of the volume from the ESS Copy Services Web user interface (WUI) and then issue the **lscfg** command on each hdisk on the Backup server, as shown in Figure F-2.

```
root@sanf50:/ > lscfg -v -l hdisk16
  DEVICE            LOCATION          DESCRIPTION

  hdisk5            10-68-01          IBM FC 2105F20

        Manufacturer...............IBM
        Machine Type and Model......2105F20
        Serial Number..............40018767
        EC Level...................2.80
        Device Specific.(Z0)........10
        Device Specific.(Z1)........0004
        Device Specific.(Z2)........0013
        Device Specific.(Z3)........33701
        Device Specific.(Z4)........05
        Device Specific.(Z5)........00
```

Figure F-2   lscfg on hdisk16

The serial number indicates the location of the volume on the ESS, making it easy to identify the FlashCopy target. The first digit indicates the LSS where the volume resides; 4 in this case means LSS 14 (prefix a 1 to the first number). The next two numbers indicate

the volume number within the LSS and the last five numbers denote the serial number of the ESS unit.

Now that you know the hdisk that is the FlashCopy target, it is necessary to read the PVID into the ODM in order to import the volume group. This is accomplished by running:

```
chdev -l hdisk16 -a pv=yes
```

This will read the PVID off of the hdisk, if it has one assigned to it, but since this is a FlashCopy target, it will have the same PVID of the source volume. So this time, the `lspv` should provide us with what is shown in Figure F-3.

```
root@sanf50:/ >lspv | grep 000c3c7d5c13c183
hdisk16        000c3c7d5c13c183    none
```

Figure F-3   Identifying FlashCopy target

The `none` in the third column indicates that this hdisk does not belong to any known volume group on the Backup server.

3. Import the volume group on the Backup server. This is done by issuing the command:

```
importvg -y samplevg hdisk16
```

All of these steps should be repeated for every volume group from the production server.

This procedure should be used the first time that you FlashCopy the LUNs from a volume group on one server and mount them on a second server.

Further FlashCopy invocations

Before invoking FlashCopy, make sure that all volume groups on the target volumes of the Backup server are varied off. It is not necessary to remove the hdisk definitions from the Backup server, at least if you keep using the same set of volumes.

FlashCopy invocation

Run the AIX configuration manager (`cfgmgr`) on the Backup server, since there may be new target LUNs added that have not yet been configured. If you did not export the volume groups (`exportvg`) on the Backup server before the FlashCopy invocation, it may not be necessary to export and import them back; a simple `varyonvg` might be sufficient.

A `varyonvg` is much faster than an `export/importvg`. In the example above, `importvg -y samplevg hdisk26` will take about 32 seconds. A `varyonvg` will take only 0.5 sec. The time will vary depending on how many hdisks are in one volume group and also on the number of logical volumes per volume group.

When is an importvg required

An `importvg` is required when:

1. You do not have an entry of the volume group in the ODM. This means you either FlashCopied the corresponding LUNs of a volume group for the first time or you did an explicit `exportvg` before invoking FlashCopy.

2. You added or removed some hdisks, logical volumes, or file systems to (from) the volume group on the production server before invoking FlashCopy (in other words, you changed the volume group structure).

To determine if an **importvg** is necessary, or just a **varyonvg** is sufficient, we have to see if there is an entry in the ODM for the volume group. This is done by issuing the command:

```
lsvg | grep vgname
```

This will tell you if there is an entry for the volume group named *vgname* in the ODM.

A question that needs to be answered is: Is there any change of the volume group structure? This can be obtained from the Volume group Descriptor Area (VGDA). Any time you change the structure of a volume group, the VGDA gets changed. This is reflected in the changed time stamp of the VGDA.

> **Note:** Simply adding, changing, or deleting a file from a file system of the volume group does *not* require an export/importvg, as this does not change the VGDA.

You can query the time stamp of a VGDA from a hdisk by issuing:

```
# lquerypv -h /dev/hdiskX 11000 10
```

In this command, hdiskX is a member hdisk of a volume group.

This will produce an output similar to the following:

```
00011000   39E4E87D 1FAC826C 000C3C7D 5C13DD48  |9..}...l..<}\..H|
```

The second column is the time stamp.

You can query the time stamp of the ODM copy of the VGDA for a particular volume group by issuing:

```
odmget -q "attribute like timestamp and name like vgname" CuAt
```

In this command, vgname is the volume-group name.

Here is a sample output for the **samplevg** command:

```
# odmget -q "attribute like timestamp and name like samplevg" CuAt
  CuAt:
  name = "samplevg"
  attribute = "timestamp"
  value = "39e4e87d1fac826c"
  type = "R"
  generic = "DU"
  rep = "s"
  nls_index = 0
```

Here, the first eight digits in the field named `value` are the time stamp. As you can see, both values are the same, and a **varyonvg** will be sufficient.

Backup scripts

In this section we describe the primary and secondary scripts.

Primary script

The script primary.sh expects a list of the volume groups, which should be backed up via FlashCopy. It should be installed on the production server.

It produces an output file (whose name you have to specify) with the following format:

- ► vgname1 pvid1
- ► vgname2 pvid2
- ► And so on

So the script actually carries out the steps described in "Preparation after the first FlashCopy" on page 510. The format of the input file should be:

► vgname1
► vgname2
► And so on

Note that empty lines are not allowed, and that this script should always be executed before the script secondary.sh:

```
############################################################
# #
# primary.sh #
# #
# Script for supporting a FlashCopy Backup solution  (AIX 4.3.X) #
#                    #
# To be run on the Primary (Productional) Server            #
# #
# Run this script prior to running secondary.sh !!! #
# #
# Author: Vladimir Atanaskovik #
# #
# (c) by IBM #
# #
# Initial Coding: 10/19/2000    Vladimir Atanaskovik #
####################################################################################


# Enter here the File containing the Volume Groups
FileVolGrups=
# Enter here the Output File
# Default /tmp/primary_list.lst
OutputFile=/tmp/primary_list.lst


if [ -z "$FileVolGrups" ]
 then
 echo ERROR! Logical Volume Group Input File not specified!
 exit -1
fi

if [ -f $FileVolGrups ]
 then
 :
 else
 echo ERROR! Could not find specified Logical Volume Group Input File!
 exit -1
fi

if [ -z "$OutputFile" ]
 then
 echo ERROR! Output File not specified!
 exit -1
fi

if [ -f $OutputFile ]
 then
 rm $OutputFile
fi
```

```
VolGrups=$(cat $FileVolGrups)

if [ -z "$VolGrups" ]
 then
 echo ERROR! No volume Groups specified!
 exit -1
fi

echo $0 running .......

for Ix in $VolGrups
 do
 PhVol=$(lsvg -p $Ix|awk 'NR==3 {print $1}')
 PVId=$(lspv $PhVol|grep IDENTIFIER|awk '{print $3}')
 echo $Ix $PVId >> $OutputFile
done
```

Secondary script

The script secondary.sh is located on the Backup server. Prior to running this script, execute primary.sh. You need to enter the name (or IP address) of the production server, and the name (with full path) of the output file produced with primary.sh. Since this script will remote copy (**rcp**) the output file from the primary.sh script, make sure that the following things have been taken care of:

► The user who executes the secondary.sh script also has an account with the same name on the production server.

► Edit the /etc/hosts.equiv and ~/.rhosts file (in the home directory of the user) on the production server, so that an **rcp** from the Backup server is possible.

The script secondary.sh does all steps described in the sections: "Preparation after the first FlashCopy" on page 510 and "Further FlashCopy invocations" on page 512:

```
################################################################################
#
#
# secondary.sh
#
#
#
# Script for supporting a FlashCopy Backup solution (AIX 4.3.X)
#
#                    #
# To be run on the Backup Server               #
#
#
# Previous to this script run primary.sh on the primary (productional) server!!!      #
#
#
# Author: Vladimir Atanaskovik (vladimir@de.ibm.com)
#
#
#
# (c) by IBM
#
#
#
```

```ksh
# Initial Coding: 10/19/2000   Vladimir Atanaskovik
#
#              10/30/2000   VGDA Time Stamp Query and comparison added
#
#               10/31/2000   Minor Changes
#
################################################################################

#!ksh


# Specify the Primary host
PriHost=
# Specify the path and name of the output file created with primary.sh on the primary
server
# Default /tmp/primary_list.lst
PriFile=/tmp/primary_list.lst
# Specify local file name
LocalFile=./LocalFile.lst
# Specify Temp File to be used
TempFile="$$"_Secondary.tmp
# Specify 1-use vpaths  0-do not use vpaths
# Use Vpath only if IBM Subsystem device Driver is installed
let UseVpath=0

if [ -z "$PriHost" ]
      then
      echo ERROR! Primary Host not specified!
      exit -1
fi


if [ -z "$PriFile" ]
      then
      echo ERROR! Primary host vg/pvid list file not specified!
      exit -1
fi

if [ -z "$LocalFile" ]
      then
      echo ERROR! Local File for primary host vg/pvid list not specified!
      exit -1
fi

if [ -z "$TempFile" ]
      then
      echo ERROR! Temp File not specified!
      exit -1
fi

echo $0 running ......

rcp "$PriHost":"$PriFile" $LocalFile

RetValue=$?

if [ $RetValue -ne 0 ]
 then
 echo ERROR! Could not rcp Config File!
 exit -1
```

```
fi

VolGr=$(cat $LocalFile|cut -d ' ' -f1)

if [ -z "$VolGr" ]
 then
 echo ERROR! Invalid Format of Input File!
 exit -1
fi


PhVolId=$(cat $LocalFile|cut -d ' ' -f2)

if [ -z "$PhVolId" ]
 then
 echo ERROR! Invalid Format of Input File!
 exit -1
fi

lspv > $TempFile

let Jxvg=1

for Ixpv in $PhVolId
 do
 Hdsk=$(grep $Ixpv $TempFile|awk 'NR==1{print $1}')

 if [ -z "$Hdsk" ]
  then
  echo ERROR! Could not find hdisk with PVID $Ixpv!
  rm $TempFile
  exit -1
 fi


 Vgname=$(echo $VolGr | cut -d ' ' -f $Jxvg)

 NeedImport=0


 ### Is there already an ODM Entry for the Volume Group

 DoesExist=$(lsvg|grep $Vgname)


 if [ -n "$DoesExist" ]

  ### Yes, ODM Entry for the volume group does exist
  then

  # Check if the volume group is already varyed on
  IsVaryOn=$(lsvg -o|grep $Vgname)

  if [ -n "$IsVaryOn" ]
   then
   echo ERROR! Volume Group $Vgname already Varied-On! Possible Data Corruption!
   echo Varry Off $Vgname and re-establish FlashCopy!
   rm $TempFile
   exit -1
  fi
```

```
# Check VGDA Timeststamp
# on the hdisk

HdVGDA=$(lquerypv -h /dev/$Hdsk 11000 10|awk '{print $2}'|sed -e 'y/ABCDEF/abcdef/')

# and in the ODM
OdmVGDA=$(odmget -q "attribute like timestamp and name like $Vgname" CuAt|grep
value|awk '{print $3}'|cut -c 2-9)

## Does VGDA Timestamp on the hdisk match VGDA Timestamp in the ODM
if [ "$HdVGDA" != "$OdmVGDA" ]

 ## No - then importvg
 then
 exportvg $Vgname
 # update TempFile since Vgname has been deleted
        lspv > $TempFile
 NeedImport=1

 ## Yes - then just varyon
 else
 varyonvg $Vgname
 RetValue=$?
 if [ $RetValue -ne 0 ]
  then
     echo ERROR! Could Not Vary-On Volume Group $Vgname
  rm $TempFile
  exit -1
  fi


 fi


        ### No, ODM Entry for the volume group does not exist
 else
 NeedImport=1
fi

if [ NeedImport -eq 1 ]
 then

 TmpVlmgrp=$(grep $Ixpv $TempFile|awk 'NR==1{print $3}')

 if [ "$TmpVlmgrp" != "None"  ]
  then
  echo ERROR! Can Not Import Volume Group $Vgname
  echo Hdisk $Hdsk already assigned to another volume group!
  rm $TempFile
  exit -1
  fi

 importvg -y $Vgname $Hdsk
 RetValue=$?

 if [ $RetValue -ne 0 ]
  then
     echo ERROR! Could Not Import Volume Group $Vgname
  rm $TempFile
```

```
    exit -1
  fi

  if [ $UseVpath -eq 1 ]
    then
    hd2vp $Vgname
  fi


 fi


 let Jxvg=$Jxvg+1
done

rm $TempFile
exit 0
```

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 523. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *Clustering and IASPs for Higher Availability on the IBM @server iSeries Server*, SG24-5194
- ▶ *IBM @server iSeries Independent ASPs: A Guide to Moving Applications to IASPs*, SG24-6802
- ▶ *IBM @server iSeries in Storage Area Networks: Implementing Fibre Channel Disk and Tape with iSeries,* SG24-6220
- ▶ *IBM Tivoli Storage Resource Manager: A Practical Introduction*, SG24-6886
- ▶ *IBM TotalStorage Enterprise Storage Server Model 800,* SG24-6424
- ▶ *IBM TotalStorage Enterprise Storage Server PPRC Extended Distance*, SG24-6568
- ▶ *IBM TotalStorage Solutions for Disaster Recovery*, SG24-6547
- ▶ *Implementing ESS Copy Services with IBM @server zSeries*, SG24-5680

Other publications

These publications are also relevant as further information sources:

- ▶ *HACMP Remote Copy: ESS PPRC Guide*, SC23-4863
- ▶ *IBM TotalStorage Enterprise Storage Server Application Programming Interface Reference*, GC35-0489
- ▶ *IBM TotalStorage Enterprise Storage Server Command-Line Interface User's Guide*, SC26-7494
- ▶ *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface User's Guide*, SC26-7449
- ▶ *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide*, SC26-7446
- ▶ *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7444
- ▶ *IBM TotalStorage Enterprise Storage Server User's Guide,* SC26-7445
- ▶ *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7448
- ▶ "Advanced functions for storage subsystems: Supporting continuous availability" by Azagury, et al, in *IBM Systems Journal*, Vol. 42 (2003) No. 2 (June), found at:

  http://www.research.ibm.com/journal/sj42-2.html

521

- *Storage Solutions for Microsoft SQL Server: Snapshot Backup and Recovery with the IBM TotalStorage Enterprise Storage Server*, found at:

  http://www.storage.ibm.com/disk/ess/whitepaper.htm

Online resources

These Web sites and URLs are also relevant as further information sources:

- Adobe: Downloads

  http://www.adobe.com/support/downloads/main.html

- ESS Whitepapers

  http://www.storage.ibm.com/hardsoft/products/ess/whitepaper.htm

- ESS supported servers

  http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm

- IBM @server iSeries Copy Services for IBM TotalStorage Enterprise Storage Server brochure

  http://www-1.ibm.com/servers/eserver/iseries/service/itc/pdf/Copy-Services-ESS.pdf

- IBM @server iSeries support Web site

  http://www-912.ibm.com/

- IBM: High Availability and Clusters

  http://www-1.ibm.com/servers/eserver/iseries/ha

- IBM Redbooks

  http://www.redbooks.ibm.com

- IBM TotalStorage Enterprise Storage Server interoperability matrix

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

- Microsoft Download Center

  http://www.microsoft.com/downloads

- Microsoft Help and Support

  http://www.microsoft.com/support

- Microsoft Knowledge Base articles

  http://support.microsoft.com/support/

- Microsoft Knowledge Base Article 131658: "Use Ftedit.exe to Recover Fault Tolerant Disk Configuration"

  http://support.microsoft.com/?kbid=131658

- Microsoft Knowledge Base Article 149927: "FTedit.exe: What You Can and Cannot Use It For"

  http://support.microsoft.com/?kbid=149927

- Microsoft Knowledge Base Article 304736: How to Extend the Partition of a Cluster Shared Disk

  http://support.microsoft.com/?kbid=304736

- Microsoft TechNet: DiskPart

  http://www.microsoft.com/technet/prodtechnol/winxppro/proddocs/DiskPart.asp

► Sysinternals Freeware: Inside the Disk Key

http://www.sysinternals.com/ntw2k/info/diskkey.shtml

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

IBM

Redbooks

IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments

Characteristics of the powerful ESS Copy Services functions

FlashCopy and PPRC fundamentals

Guidelines for ESS Copy Services implementations

This IBM Redbook describes the copy functions available with the IBM TotalStorage Enterprise Storage Server (ESS). The powerful ESS Copy Services functions are explained in detail, and their respective characteristics are thoroughly covered. This redbook also gives information on how to manage the various ESS Copy Services functions, and finally discusses their implementations.

Because this redbook provides a broad understanding of the ESS Copy Services functions, as well as going into detail about the management interfaces and the implementation considerations, we recommend it for IT professionals who are planning the implementation of any of the ESS Copy Services functions in an open-systems environment.

This fourth edition of the redbook has been updated with PPRC over Fibre Channel, enabling the use of Fibre Channel Protocol (FCP) as a communications link between the primary and secondary ESS, enhancements to the ESS Application Programming Interface (ESS API), supporting ESS Copy Services configuration and use, Asynchronous Cascading PPRC enhancements, exploiting PPRC Failover and Failback modes, and Copy Services domain support for up to eight ESSs per domain.