IBM eserver

iSeries in Storage Area Networks

A Guide to Implementing FC Disk and Tape with iSeries

Learn how iSeries fits in a SAN environment

Plan how to migrate from SCSI to fiber channel disk and tape

Understand the new Multipath function

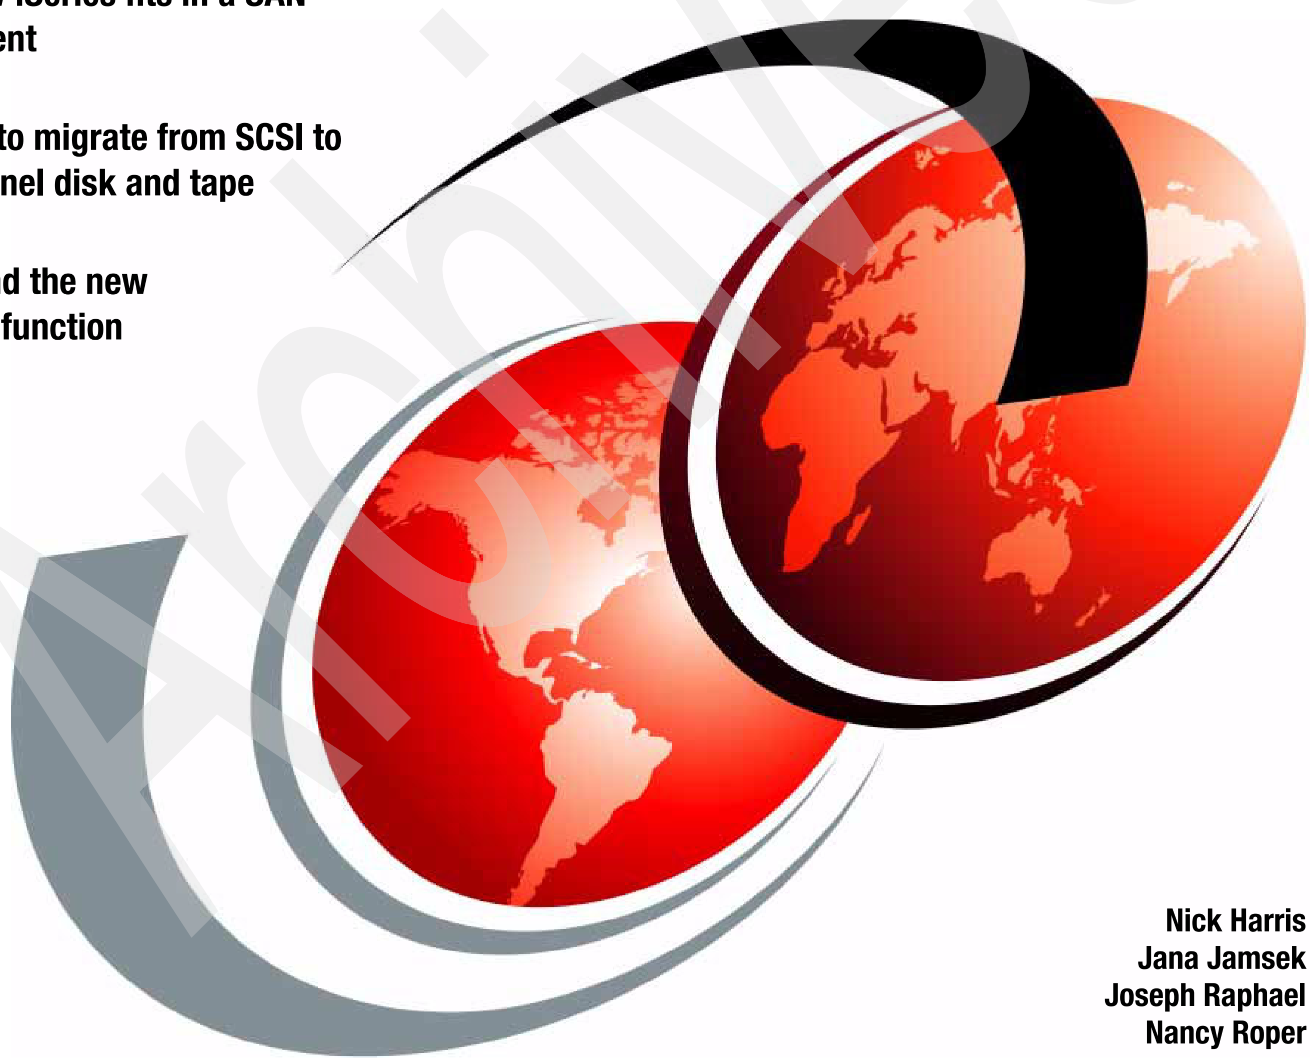

Nick Harris
Jana Jamsek
Joseph Raphael
Nancy Roper

Redbooks

ibm.com/redbooks

International Technical Support Organization

iSeries in Storage Area Networks
A Guide to Implementing FC Disk and Tape with iSeries

March 2005

SG24-6220-01

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

Second Edition (March 2005)

This edition applies to V5R1 and V5R2 of OS/400 and V5R3 of i5/OS.

Contents

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:

- eserver®
- eserver®
- ibm.com®
- iSeries™
- i5/OS™
- pSeries®
- xSeries®
- zSeries®
- AIX®
- AS/400®
- Domino®
- DB2®
- Enterprise Storage Server®
- ESCON®
- FlashCopy®
- FICON®
- Informix®
- IBM®
- Magstar®
- MVS™
- OS/390®
- OS/400®
- PowerPC®
- Redbooks™
- Redbooks (logo) ™
- RISC System/6000®
- S/390®
- Seascape®
- System/38™
- System/390®
- Tivoli®
- TotalStorage®
- Wave®
- WebSphere®

The following terms are trademarks of other companies:

Intel and Intel Inside (logos) are trademarks of Intel Corporation in the United States, other countries, or both.

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 positions the Fibre Channel (FC) support provided by IBM @server i5 and iSeries™, plus the implementation of SAN support in OS/400® Version 5 Release 1 and later. The new FC disk adapter support primarily replaces the existing SCSI #6501 disk adapter support for external disk. FC support opens the possibility of an iSeries implementation in a Storage Area Network (SAN).

Fibre Channel adds a new dimension to iSeries networks, providing extended distance for both tape and disk. Migration from SCSI to FC attached disk and tape will be discussed, with illustrations of how the migration is achieved.

New sections in this redbook cover a number of Fibre Channel switches and the new Multipath function.

IBM specialists and Business Partners will be able to use this redbook for planning and implementing SAN-based solutions on iSeries. The design and implementation of these new networks will also need the involvement of IBM Storage Subsystems Group (SSG) and associated Storage Business Partners.

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 Rochester Center.

Nick Harris is a Consulting IT Specialist for IBM and has spent the last six years in the ITSO, Rochester Center. He specializes in LPAR, IBM @server iSeries hardware and software, external disk, Integrated xSeries® Server for iSeries, and Linux. He also writes and teaches IBM classes worldwide on areas of AS/400 and iSeries system design and server consolidation. He spent 13 years in the United Kingdom AS/400 Business and has experience in S/36, S/38, AS/400, and iSeries servers. His e-mail address is mailto:niharris@us.ibm.com.

Jana Jamsek is an IT Specialist in IBM Slovenia. She works in Storage Advanced Technical Support for EMEA as a specialist for Storage and iSeries. Jana has eight years of experience in the iSeries and AS/400 area, and four years experience in Storage. She holds a master's degree in Computer Science and a degree in mathematics from University of Ljubljana, Slovenia. She co-authored the Redpaper "The LTO Ultrium Primer for IBM @server iSeries Customers".

Joseph Raphael has been with IBM since November 1997, working as a Technical Executive in IBM Global Services India. He has worked on Intel®-based platforms, network computers, high end printers, storage devices, and iSeries. Presently, he works as the Regional Specialist for Storage and iSeries platforms in Bangalore, India. His main job responsibility includes planning, implementation, and support for all the Storage and iSeries platforms, and he has been involved in various CISC to RISC migrations, OS upgrades, LPAR implementations, SAN, Tivoli® Storage Manager and Backup Recovery and Media Services implementations.

Nancy Roper is a Consulting IT Specialist working in the IBM Rochester iSeries Advanced Technical Support Group. Her areas of expertise are iSeries High Availability and Systems Management, including iSeries disk, tape, SAN, and software products, such as BRMS and the Advanced Job Scheduler. Nancy holds a bachelor’s degree of Mathematics, with joint honors in Applied Math and Computer Science, from the University of Waterloo, in Waterloo, Canada.

Authors of the first edition:

Stu Preacher is a Certified IT Consultant, specializing in AS/400 and iSeries systems management and high availability. He has worked for IBM for over 28 years. He has been working with the AS/400 and iSeries since their introduction, having previously worked on System/34 and System/38™. Currently, Stu is Group Leader for iSeries Technical Support in the United Kingdom and is responsible for iSeries External Storage in Europe. He has worked with many of the largest AS/400 and iSeries customers across Europe.

Jason Blackburn is an IT Consultant within the IBM Storage Subsystems Division at San Jose. He has three years of experience providing technical support in IBM. Jason has provided technical support for networking and server consolidation on UNIX® for IBM and its customers. He is currently engaged in providing support for open systems storage across multiple platforms and a wide customer base.

Mark Huntley is a Senior Technical Consultant for a leading United Kingdom Business Partner, DCS eIntegration. He specializes in iSeries, storage, xSeries and pSeries® solutions. He has worked on the iSeries and AS/400 servers since their introduction in 1988 as well as on the S/38. Mark worked for IBM UK for 14 years providing pre-sales support and consulting on a range of leading UK accounts. He has participated in residencies for the *IBM eServer iSeries System Handbook*, GA19-5486; *IBM eServer iSeries and AS/400e System Builder*, SG24-2155; and *IBM eServer iSeries in Storage Area Networks: Implementing Fibre Channel Disk and Tape with iSeries*, SG24-6220. His e-mail address is mailto:mark.huntley@dcs-ei.co.uk.

Friedrich Neidhart is a Technical Specialist (FTSS) for Storage on iSeries, AS/400® (pre-sales) in Germany, SSG Central Region, and has worked for IBM since 1968, starting as a Hardware Customer Engineer on System 3, System/38, and then as an AS/400 System Engineer. During the last four years Fred has focused on Storage Solutions (tapes and disk) for AS/400 and iSeries.

Mikey Sparks is an I/T Specialist with IBM Global Services in Minnesota. He has worked with IBM AS/400 customers for over 13 years in manufacturing, telecommunications, and distribution industries. He has specialized in IBM systems management products on the AS/400, such as BRMS, TSM, and LPAR.

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

Andy Kulich
Amit Dave
Steve Finnes
Virgina Kautz
Scott Maxson
Jeff Palm
Kristi Harney
IBM Rochester

Become a published author

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

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. JLU Building 107-2
  3605 Highway 52N
  Rochester, Minnesota 55901-7829

Part 1

Introduction and positioning

In Part 1 we discuss the iSeries storage subsystem and how it relates to the disk located in an Enterprise Storage Server®. The new Fibre Channel tape and disk adapters are covered along with some examples of use.

1

Introduction to iSeries and Storage Area Networks (SAN)

In this chapter we overview the latest enhancements to the iSeries SAN capabilities. We will also describe the iSeries internal storage characteristics and features. Key to any SAN based solution are IBM TotalStorage® Enterprise Storage Servers and Tape products. We review these products and how they connect to iSeries models.

1.1 Overview of enhancements

With i5/OS™ V5R3 released in July 2004 came a very significant enhancement, Multipath support. This much awaited new function provides an alternative path should there be a FC adapter failure. I/Os are routed through another FC IOA in the same multipath group. This is a low cost availability option that would typically cost an FC adapter for any currently installed. This does not alter the requirement for a load source disk that must be mounted in the iSeries CEC for a non-LPAR system and in an expansion tower for a logical partitioned system.

New eServer i5 models announced during 2004 and logical partitioned do not require an i5/OS managing partition or an i5/OS running in the CEC. But as with multipath support this does not remove the requirement for a load source disk.

In September 2003 a new #2787 FC adapter was release for iSeries. #2787 / #0626: PCI-X Fibre Channel Disk Controller, provides Fibre Channel disk attachment for external disk devices such as IBM's TotalStorage Enterprise Storage Server. #0626 is used for attachment of external storage to PowerPC® Linux partitions on iSeries. #2787 supports point-to-point, arbitrated loop and switch fabric topologies and has an LC type cable connector. OS/400 V5R2 plus PTFs, or later is required. For required PTF information, refer to informational APAR II13609.

With iSeries Models 270 and 8xx and OS/400 V5R1, two FC adapter became available:

- #2765 Fibre Channel Tape adapter
- #2766 Fibre Channel Disk adapter

The #2766 is the Fibre Channel (FC) replacement for the #6501 Disk adapter that is currently used to attach to ESS.The adapters remain the same for V5R2 and later. However, the support provided by V5R2 for the adapters is different. This redbook discusses these changes throughout the book.

SANs are the latest development in the disk and tape attachment business. They add value by consolidating the disk and tape infrastructure of multiple and heterogeneous platforms into a single set of centrally managed resources.

A SAN provides a dedicated, centrally managed information infrastructure for multiple platforms. It employs a combination of technologies, including hardware, software, and networking components. Using Fibre Channel hubs, switches, and gateways, a SAN offers the ability to connect storage devices, including disk arrays, optical disks, or tape libraries, and servers to provide centralized storage.

For the iSeries customer, SAN is an option in ways that it is not an option for a UNIX or PC server customer. You need to ask the question, “Should I include my iSeries disk in a SAN solution?” Before you answer, there are some points to consider.

This chapter discusses iSeries storage internal and the Enterprise Storage Server (ESS) external storage in general. It covers the positioning of external storage relative to internal storage. The new external storage is part of the SAN proposition. This chapter briefly covers what SAN is and where the iSeries fits into this technology.

1.2 iSeries storage architecture

The iSeries uses the concept of single-level storage. The iSeries storage architecture (inherited from its predecessor systems System/38 and AS/400) is defined by a high-level machine interface. This interface is referred to as Technology Independent Machine Interface

(TIMI). It isolates applications and much of the operating system from the actual underlying systems hardware. They are also unaware of the characteristics of any storage devices on the system because of single-level storage.

The iSeries is a multi-user system. As the number of users increase, you do not need to increase the storage. Users share applications and databases on the iSeries. As far as applications on the iSeries are concerned, there is no such thing as a disk unit, which we explain later in this chapter. This idea of applications not being aware of the underlying disk structure is similar to the SAN proposition. In fact, the iSeries already has a very SAN-like disk implementation.

At present, high-speed link loop (HSL) only offers a medium distance solution (250m), with Fibre HSL. The older System Product Division (SPD) technology offered increased distances up to 500m between the system and an expansion tower containing disk or a tape adapter

With V5R3 Cross Site Mirroring (XSM) function offers greater distances using GB ethernet connections between systems.

1.2.1 Single-level storage

Both the main memory and physical disk units are treated as a very large virtual address space, known as *single-level storage*. This may be the most significant differentiation in a SAN solution implementation on the iSeries when compared to other systems. It is this architecture that masks the complexities of adding additional disk units to the system. Essentially, adding a disk unit is treated as an extension of the main memory store and is, therefore, simply one large virtual space that a user accesses.

There is only one image of application and data that the users access rather than creating their own dedicated application, database, and free storage environment. As usage grows, you can upgrade disk storage, memory capacity, and processor power, all within a single system, rather than adding multiple servers.

Figure 1-1 shows an example of single-level storage being grown without the user knowing about the underlying disk structure.

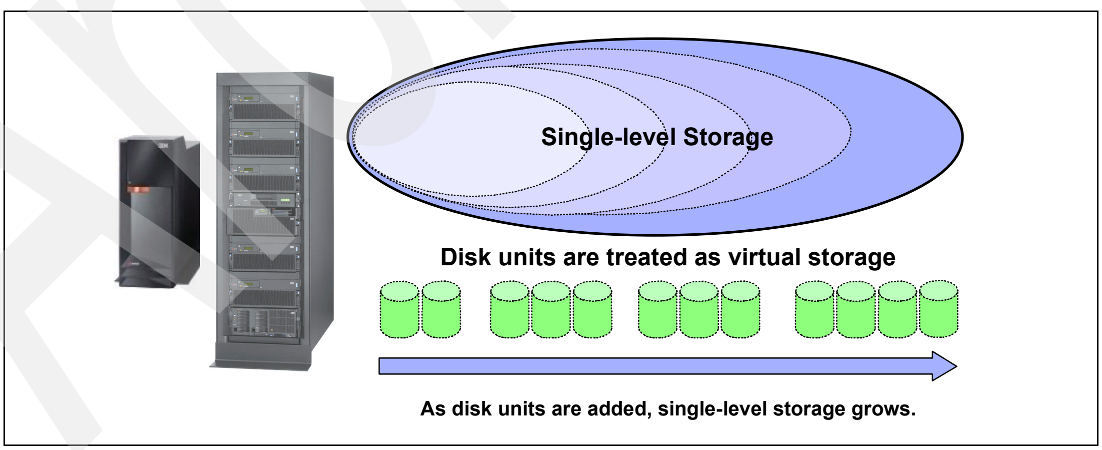

Figure 1-1 Single-level storage

1.2.2 iSeries storage management

Unlike other architectures, the iSeries has always been able to add physical disk storage. In doing so, the new disk storage is automatically treated as an extension of the virtual address space, and data is automatically spread across the entire virtual address space. This means that data striping and load balancing are already automated. Free address space (unused disk storage) is comprised from a collection of all of the disk units, rather than at each disk unit level. Storage management and caching of data into main storage is completely automated, based on expert cache algorithms. For more information, see “Expert cache” on page 7.

Essentially, expert cache has the intelligence to study the I/O patterns of each application. For random I/O, it does not cache or pre-fetch any data. For sequential I/O, it performs parallel I/O request across all of the disk units to pre-fetch data into main memory for processing.

The iSeries server keeps objects in a single address space. This address space is quite large at 64-bits (2^{64} bytes, over 18 million terabytes of addressable storage is possible). Compare this with a 32-bit addressing system where each process (job) has 4 GB of addressable memory. The operating system maps portions of this address space as necessary, either to disk units for permanent storage or to main memory for manipulation by application programs. Because a single page table maps all of the virtual addresses to physical addresses, task switching is very efficient. Single-level storage further eliminates the need for address translation, and therefore, speeds up data access.

Single-level storage enables another extremely important iSeries benefit, *object persistence*. Object persistence means that the object continues to exist in the memory forever. This mechanism is safe because objects allow only those operations, which are defined for a particular object type. For example, you cannot read a program object or execute a database file.

This mechanism allows great flexibility for storage management. Storage management on the iSeries system is automated. The iSeries server handles selecting the physical disk drive (sometimes referred to as direct access storage device (DASD)) to store the data. Storage management automatically spreads the data across the disk arms and continues to add records to files until specified threshold levels are reached. This capability is just being added as a new capability on other midrange server platforms. The iSeries has been able to scatter load data to many disk drives since its existence. This is a performance feature that has only started to find its way into other open systems.

Single-level storage is efficient too. All data is already mapped to virtual address space, while other systems have to do it every time they need to access the data. Regardless of how many application programs need to use an object, only one copy of it is required to exist. This makes the entire main storage of an iSeries server a fast cache for disk storage. The system can take advantage of increased memory without needing complex management activity. Other systems require sophisticated and expensive cache management hardware and software (on-and-out board) to achieve something that always been available on the iSeries server.

Other open systems have started to move to consolidated storage in an attempt to reduce the amount of duplicate storage for applications and data. This problem does not exist on the iSeries server with OS/400.

1.2.3 Main storage

Main storage is an integral part of single-level storage on the iSeries server. The amount of main storage varies by CPU model. Main storage is internal random access memory. It is typically divided into system memory pools.

These pools of storage facilitate tuning of the application environment on the fly. Chunks of main storage can be moved around as applications require without any intervention by operators or end-users. Because the iSeries is a multi-user, multi-application server, this is an important function, which allows great flexibility to meet changing workload demands.

In a logically partitioned (LPAR) environment at V5R1, you can move this main storage from one partition to another without performing IPL on the system. You can automate this main storage movement to account for significant changes in workload within a partition, for example, when two systems are represented by two separate logical partitions. The main storage can be moved from one partition to another during peak processing periods. For example, one system in a different time zone is not under load, but the other system needs more storage for its end-of-day processing. It is simple to move main storage for the end-of-day task and then return it when the task completes.

1.2.4 Cache

Cache is a supplemental memory system that temporarily stores frequently used instructions and data. Cache allows for quicker processing by the CPU. It augments and is an extension of main storage.

Figure 1-2 shows the various levels of cache that are available on the iSeries server. It should be noted here that the Small Computer Systems Interface (SCSI) interface speed for the internal PCI RAID-5 adapter is not 20 MB/s, which is often confused with the throughput rate of the #6501 input/output processor (IOP). In fact, the aggregate throughput for the internal PCI RAID-5 adapter can be up to 320 MB/s depending on the adapter used.

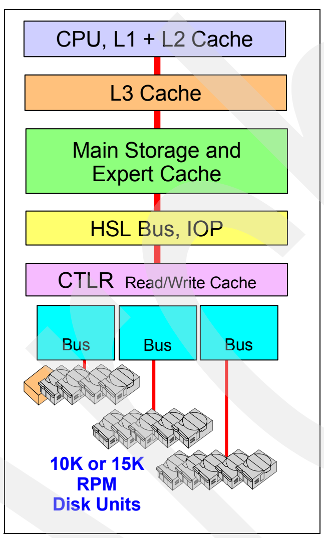

Figure 1-2 iSeries cache

Expert cache

Expert cache is a set of algorithms that execute in the main central processing unit (CPU). Expert cache uses designated pools in main storage to cache database operations. The size of the pools is adjusted dynamically or controlled by the system operator as a function of time. This allows the user to maximize the efficiency of the cache and main storage usage, as workloads change over the course of a typical day.

By caching in main storage, the system eliminates access to the storage devices and reduces associated I/O traffic on the internal system buses. Expert cache works by minimizing the affect of synchronous DASD I/Os on a job. The best candidates for performance improvement by using expert cache are jobs that are most affected by synchronous DASD I/Os.

Extended Adaptive Cache (EAC)

Extended Adaptive Cache is an advanced read cache technology that improves both the I/O subsystem and system response times. It reduces the number of physical I/O requests that are read from disk. EAC operates at the disk subsystem controller level and does not affect the iSeries system processor. Cache management is performed automatically within the input/output adapter (IOA).

Read Cache Device (RCD)

Feature Codes #4331 and #6831 (CCIN #6731) have been withdrawn from marketing. The information provided here is a reference for existing users.

The RCD is a solid state 1.6 GB cache memory assembled in a package. From the exterior, it resembles a disk unit. The RCD is optimized for use as cache memory on the iSeries server. It is packed with solid state memory chips that function at electronic speed. It works in conjunction with Extended Adaptive Cache and helps reduce physical read requests to disk units attached to the #2748 controller.

1.2.5 Direct access storage device

DASD (more commonly known as *disk drives*) are either internal to the iSeries or attached externally. Disk drives are grouped together into auxiliary storage pools (ASPs). Disk drives can be protected or unprotected. Protection can be mirroring or device parity. You can choose not to implement protection, but we do not recommend this.

Disk mirroring

iSeries native DASD mirroring provides protection from a system outage caused by a disk drive, disk controller, disk I/O processor, or even system bus failure, by optionally duplicating each of these items.

Once mirrored protection is started, data is written to two disk units of a mirrored pair for each write operation. If any disk drive fails, the system continues to use the other drive in the mirrored pair. All protection provided by DASD mirroring is visible to OS/400.

Note: If you require a high level of availability disk mirroring is not the whole solution. You should also plan to mirror your IOP/IOAs and possibly your buses.

Device parity protection

Device parity protection, more commonly known as RAID-5, is a way to protect a system from an outage caused by a single disk drive failure. Through a form of parity checking across a set of disk drives, the drives can continue to function despite a disk failure, recreating data from the failed drive as it is requested. OS/400 controls the implementation of RAID-5 across internal disk subsystems. In the case of a drive failure, the system continues to perform normally while the operator is informed of the condition.

Concurrent maintenance is provided for the latest model RAID-5 drives on the iSeries server, allowing maintenance, and even drive replacement, to be completed without requiring the system to be in a restricted state.

RAID-5 is the minimum level of DASD protection recommended.

1.2.6 iSeries Fibre Channel support

With V5R3, the supported external SAN connections via native FC adapters are the:

- #2787 PCI-X Disk FC adapter
- #2765 PCI Magnetic Media FC adapter
- #2766 PCI Disk FC adapter

These adapters are only supported in iSeries Models 270, 8xx and the new @server i5 servers. These adapters are not supported on the older PCI-based system 6xx and 7xx models.

#2787 PCI-X Fibre Channel Disk Controller

The #2787 PCI-X Fibre Channel Disk Controller provides Fibre Channel attachment capability for external disk devices. The #2787 supports point-to-point and arbitrated loop topologies and has an LC-type cable connector.

Each #2787 is shipped with a wrap connector (part number 05N6767). The #2787 supports 64-bit, 133MHz PCI-X bus speeds. The following adapter kits are required when connecting SC-type cables to the #2787:

- #0371 - LC-SC Adapter Kit (50 um) can be ordered, both on initial, model upgrades, and simple MES orders. This optional kit is used to attach SC-type fibre (50 micron) cables to a #2787. The #0371 kit contains a 2m LC-ST cable and ST-SC adapter for 50 micron fiber cable.
- #0372 - LC-SC Adapter Kit (62.5 um) can be ordered, both on initial, model upgrades, and simple MES orders. This optional kit is used to attach SC-type fiber (62.5 micron) cables to a #2787. The #0372 kit contains a 2m LC-ST cable and ST-SC adapter for 62.5 micron fiber cable.

An optics cleaning kit (part number 46G6844) and instruction sheet (part number 21P6238, form number SY27-2604) is supplied, one per system, when a #2787 is present or ordered.

The #2787 PCI-X Fibre Channel Disk Controller is the recommended replacement card for the #2766 PCI Fibre Channel Disk Controller.

When used as a Direct Attached Adapter for Linux, order the #0626 in place of the #2787.

Note: Fibre Channel cables for the #2787 PCI-X Fibre Channel Disk Controller are Customer supplied.

#2765 PCI Magnetic Media Adapter

This IOA is supported in any of the iSeries system units or expansion towers. It requires an IOP, but this processor is not dedicated. Initial estimates show the #2765 drives approximately half the capacity of the IOP. This allows for other low-speed IOAs to be driven by the same IOP. See the PCI placement guidelines in the *IBM @server iSeries and AS/400e System Builder*, SG24-2155, for more details on placement and capacity.

This new adapter is attractive to logically partitioned iSeries servers. In this configuration, it is possible to share the attached tape between partitions. Sharing is achieved by dynamic resource movement, available with V5R1. A caution with this movement is that it is performed at an IOP level. This is not a problem if the #2765 is the only IOA under the IOP. If you choose to add other IOAs under the IOP, they are also moved during sharing. The moved IOAs are then unavailable to the original partition. Any dynamic movement requires you to vary off the IOAs controlled by the IOP.

#2766 PCI Disk Adapter

The new FC disk adapter is supported in most locations in the iSeries system units and expansion towers. However, this adapter is not supported under the embedded IOP in the system CEC.

The #2766 requires a dedicated IOP in all positions and is capable of driving the IOP at a full rated speed. The iSeries configurator accurately positions and ties this adapter to an IOP.

It is possible to add more combinations of IOP/IOA and associated Logical Unit Numbers (LUNs) than the total storage supported by the particular iSeries model. You can find the total storage for a particular model in the *IBM @server iSeries and AS/400e System Builder*, SG24-2155, and it should not be exceeded.

1.2.7 Multipath support for TotalStorage Enterprise Storage Server (ESS)

At V5R3, up to eight connections can be defined from multiple input/output processors on an iSeries server to a single logical unit number (LUN) in the TotalStorage ESS. Each connection for a multipath disk unit functions independently. Multiple connections can provide uninterrupted access to the disks in an ESS even if a single path fails.

1.3 Enterprise Storage Server

The Enterprise Storage Server provides a natural solution for the consolidation of server storage through its extensive heterogeneous server connectivity. The ESS supports access to vast quantities of data through many advanced function and features.

The ESS enables customers to replace islands of isolated disks with a smaller number of shared subsystems. There are some exciting new functions for both the S/390® customer and the open systems customer. In this text, the iSeries is included in the open systems category.

For the non-mainframe customer, there are new features that previously were only available on mainframe storage. These are functions that will enable better management of data. For both sets of customers, there is capacity and performance to meet the largest requirements.

The IBM TotalStorage ESS provides unparalleled functions for all of the server family of e-business servers, and also for other 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 Seascape® architecture (see Figure 1-3 on page 11) is the key to the development of IBM's 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 Total Storage Enterprise Storage Server to evolve from the initial E models to the succeeding F models, and to the current 800 Turbo II 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.

The move to e-business 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.

To meet the unique requirements of e-business, where massive swings in the demands placed on IT infrastructure are common and continuous operation is imperative, enterprises will need very high-performance, intelligent storage technologies and systems that can support any server application in their business environment, today and into the future. The IBM TotalStorage Enterprise Storage Server has set new standards in function, performance, and scalability in these most challenging environment.

The Enterprise Storage Server also implements RAID-5 protection in a similar way to the iSeries integrated DASD. However, an iSeries operator is not notified of a disk drive failure in the ESS in the way that internal drives failures are. Instead, they are logged by the StorWatch control system. StorWatch can “phone home” to IBM Service in the event of a problem, or it can send an Simple Network Management Protocol (SNMP) trap to the iSeries if this function is enabled. Because RAID-5 is implemented within the ESS, the iSeries does not control it. While with internal drives, OS/400 is in control. This can be an advantage for ESS because you can mirror RAID-5 ESS drives, where you can’t mirror RAID-5 internal drives.

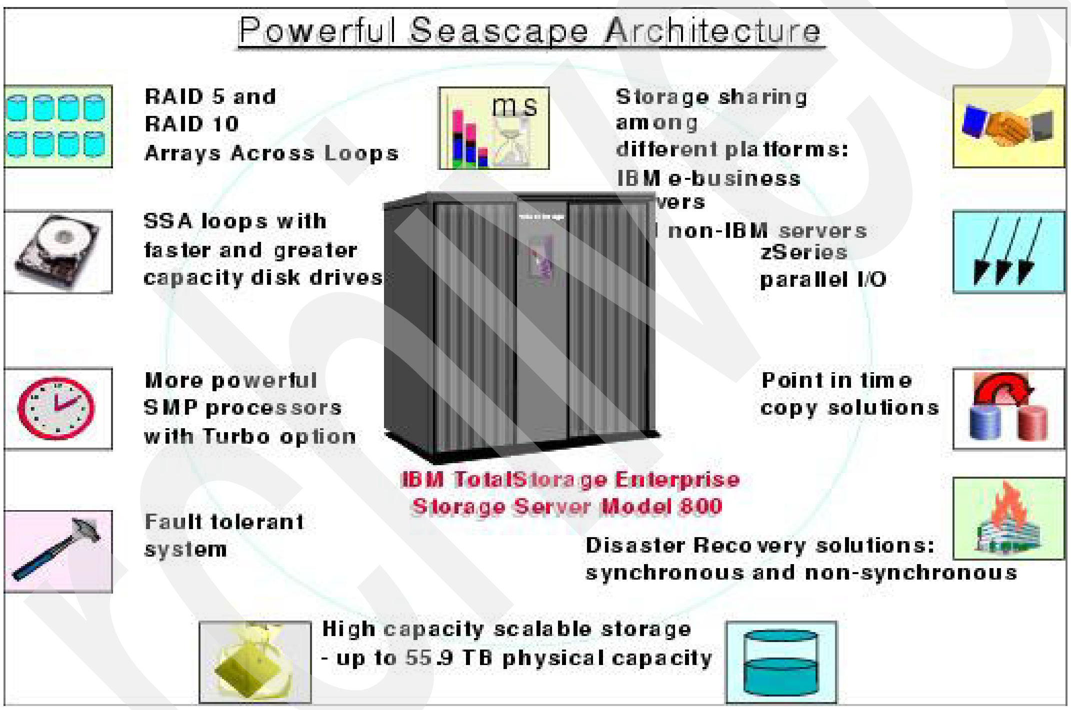

Figure 1-3 Enterprise Storage Server with Seascape

See the IBM TotalStorage Web site for more details on the latest ESS products and enhancements:

http://www.storage.ibm.com/ssg.html

1.3.1 Storage consolidation

The Enterprise Storage Server’s high performance, attachment flexibility, and large capacity enable data consolidation from different platforms onto a single high performance, high availability server (Figure 1-4).

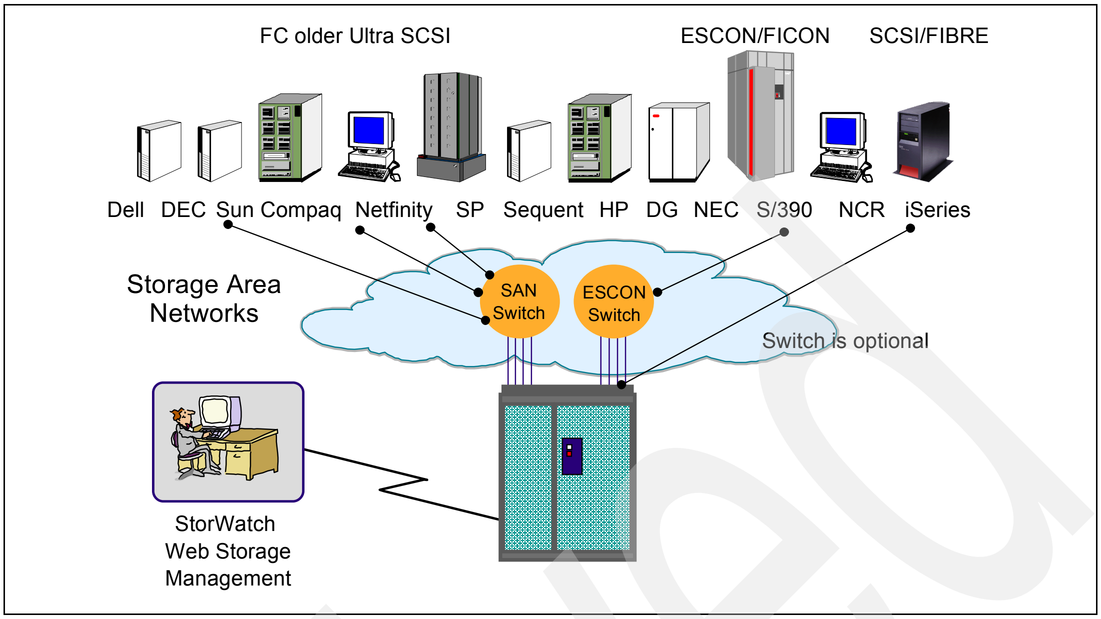

Figure 1-4 ESS storage consolidation

ESS supports all the major server platforms, from S/390 to iSeries, Windows® NT, Windows 2000 and many varieties of UNIX.

The ESS provides integrated caching and support for redundant arrays of inexpensive disks (RAID) for the disk drive modules (DDMs). The ESS Model 800 provides the following features:

- RAID-5 and either RAID-10 arrays or non-raid disk groups. Arrays can also be configured across loops
- Fast disk drives with speeds of 10,000 and 15,000 revolutions per minute (rpm).
- Fast reduced instruction-set computer (RISC) processors.
- A fault-tolerant system.
- Sharing of storage between open-systems hosts.
- Instant copy solutions with FlashCopy®.
- Disaster recovery solutions with Peer-to-Peer Remote Copy (PPRC).
- Cache option up to 64 GB.

The ESS Model 800 has two clusters each contains RISC SMP processors providing two options:

- Standard processor feature (can be upgraded to Turbo II processors)
- Turbo II processor feature.
- The ESS Model 800 supports 18.2, and 36.4, 72.8, and 145.6 GB DDMs
- The ESS Model 800 with an expansion enclosure and using 145.6 GB homogeneous DDMs, can provide the maximum data storage capacity of 55.9 TB.

The large-capacity DDMs, in conjunction with DDM intermix support, provide an almost unlimited number of configuration options for the ESS.

1.3.2 Performance

The Enterprise Storage Server is designed as a high performance storage solution and takes advantage of IBM's leading technologies. Today, where business can reach global markets through e-business, there is a need for business solutions that can deliver high levels of performance continuously every day.

The ESS can provide a solution that can handle simultaneously different workloads, Business Intelligence models, large databases for Enterprise Resource Planning (ERP), and online and Internet transactions that must run alongside each other with minimal impact. For workloads that cannot take advantage of cache, there are high performance disk arrays with fast disks and serial interconnect technology. Regardless of the type of workload, even mixed workloads, ESS delivers high performance.

> **Important:** If the workloads are mission critical, consider separating them into different clusters or on a separate ESS. In some circumstances, applications can compete for resources, causing possible degradation of performance. See the *IBM TotalStorage Enterprise Storage Server Configuration Planner*, SC26-7353, for more details.
>
> You may also consider running a full benchmark before implementing consolidation of mission-critical workloads. When implementing an external solution, you may have to accept a compromise in performance, which would not be the case with internal drives, especially if the ESS is not dedicated to a single iSeries server.

1.3.3 Disaster recovery and availability

The Enterprise Storage Server is a fault-tolerant storage subsystem, which, in general, can be maintained and upgraded concurrently with customer operation. With advanced functions, such as Copy Services, you can copy data from one set of disks within a single ESS or to another distant ESS.

As with any mixed server environment, there are special considerations. The ESS can accommodate all iSeries disks with the exception of the load source unit. A mirror image of the OS/400 load source unit can be located in the ESS for recovery purposes.

For any server that has data in main storage, which has not been written to disk, availability is limited to data residing on disk. The ESS storage pool that is provided does not take into account data that is still in iSeries memory which can be up to 1 TB with L2/L3 cache of 1.6/36MB. The memory stored there is not copied to disk unless the system is shutdown. Therefore, although ESS as a server has high reliability and redundancy features for the disks it controls, the total solution (both application and storage) cannot be considered as a truly high availability solution. This applies to iSeries servers as well as others.

To implement a high availability solution, you must also consider the application server and make that application highly reliable. Then both the application server and the storage server can be considered a total high availability solution.

1.3.4 Storage Area Network

For open systems customers who are looking at Storage Area Networks and Fibre Channel, the ESS supports a variety of Fibre Channel attachment options (Figure 1-5).

The ESS supports native Fibre Channel, which provides a basis for the future development of full SAN exploitation in such areas as disk pooling, file pooling, and copy services. However, not all of these technologies may apply to the iSeries server. Up to 16 Fibre Channel ports are available on an ESS. Each port can support point-to-point and fabric (switched) connections as well as Fibre Channel - Arbitrated Loop (FC-AL). The iSeries server currently supports both point-to-point and arbitrated loop fabric. iSeries support of switch fabrics is planned for a future release.

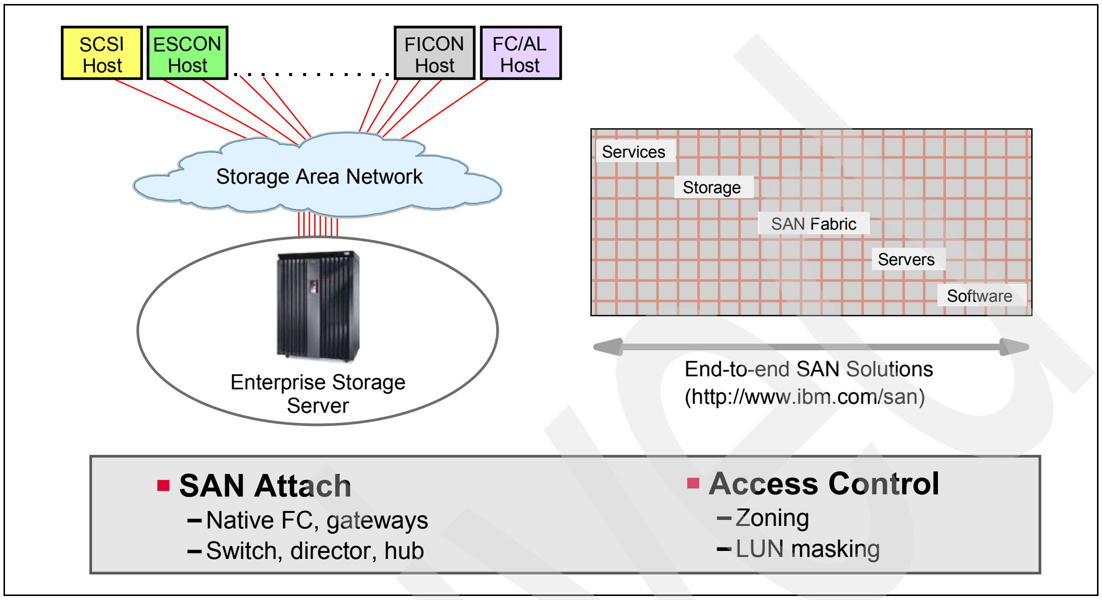

Figure 1-5 Storage Area Network support

1.4 Terminology

This section briefly defines the terminology used by both SCSI and Fibre Channel.

Disk Drive Modules (DDMs)

DDMs are the physical disk units within the ESS. Currently, these DDMs can be 9 GB, 18 GB, or 36 GB.

Logical Unit Numbers (LUNs)

In an ESS, LUNs are the virtual representations of a disk as seen from the host system. LUNs are mapped over the DDMs. In reality, a LUN may span multiple physical disks, and the size of the LUN is set when it is defined to the ESS.

In the example in Figure 1-6, the 8.58 Gb iSeries drives are really 8.58 Gb stripes across the physical DDMs. These stripes are called LUNs.

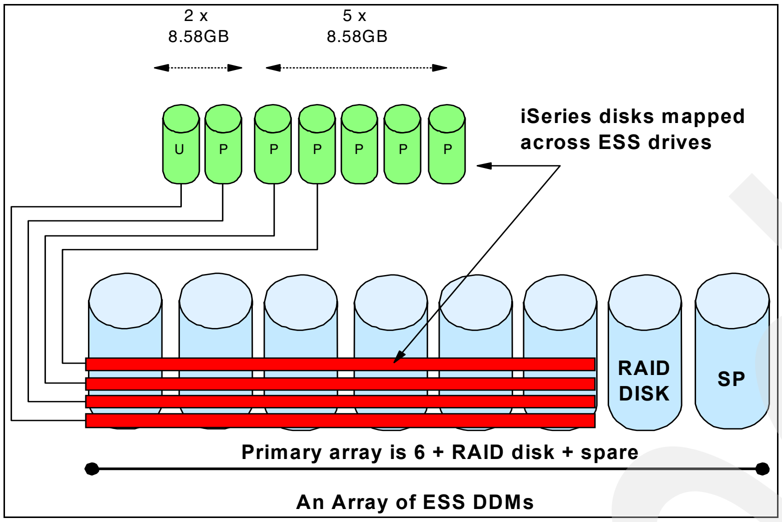

Figure 1-6 iSeries drives represented as LUNs

It is possible to have multiple LUN sizes across the same set of DDMs. It is also possible to have multiple servers attached to the same array. For iSeries installations, we recommend that you assign only one system per array. This is because the entire array needs to be reformatted if a LUN is to be assigned to a different system type.

When discussing Fibre Channel tape devices, you can regard LUNs as the individual units that make up the tape device. For example, with a 3590 E11, there are two LUNs: one for the actual tape drive and one for the Random Access Cartridge Loader.

Fibre Channel initiator

An initiator on an FC-AL is the host system. With iSeries attachment via #2766 FC adapter, there can be up to five initiators per arbitrated loop.

Fibre Channel Target

The target on an FC-AL is the device. For example, the ESS is a single target, regardless of how many LUNs or hosts are connected to it. Similarly, a tape device (for example a 3580) is also a single target. With iSeries, there can only be one target per arbitrated loop.

Fixed Block Architecture (FBA)

SCSI disks use a fixed block architecture, that is, the disk is arranged in fixed size blocks or sectors. With an FB architecture the location of any block can be calculated to retrieve that block. The concept of tracks and cylinders also exists, because on a physical disk, there are multiple blocks per track, and a cylinder is the group of tracks that exists under the disk heads at one point in time without doing a seek. In the ESS, all the tracks and cylinders are logical. They are mapped onto arrays and disks, which may have very different track and cylinder sizes.

1.4.1 Copy functions

ESS Copy Services supports several hardware copy functions (Figure 1-7) for two purposes:

- Mirroring for disaster recovery solutions
- Copy functions that provide an instant copy of data

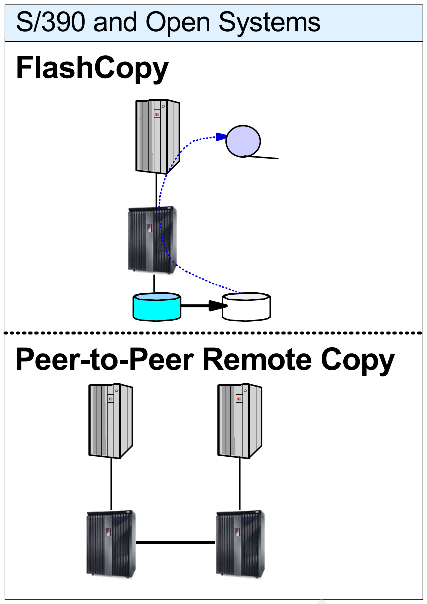

Figure 1-7 Copy functions

Peer-to-Peer Remote Copy (PPRC)

With ESS's Peer-to-Peer Remote Copy function, S/390 and open systems now have a disaster recovery solution to protect data in case one location is destroyed. PPRC provides a synchronous copy of selected volumes up to a distance of 103 km. PPRC is administrated from the Web browser interface (StorWatch) and with command line interface (CLI). At present, OS/400 does not have a command line interface into the ESS. This means the initiation tasks must be run from the StorWatch browser.

The ESCON® cable connects the ESS systems in a PPRC arrangement as shown in Figure 1-8. The ESCON connect is a synchronous function, with write operations completed on both copies (local and remote) before they are reported complete to the operating system.

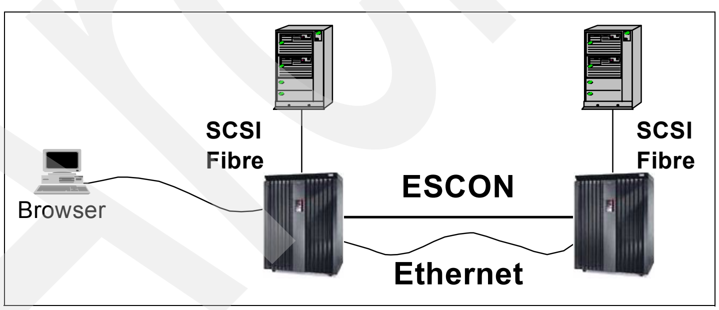

Figure 1-8 Peer-to-Peer Remote Copy function

There is a new component of PPRC called *Peer-to-Peer Remote Copy - Extended Distance* (PPRC-XD). You can learn more about PPRC-XD on the IBM TotalStorage Web site at:

http://www.storage.ibm.com/hardsoft/products/ess/index.html

FlashCopy

A significant function of IBM ESS Copy Services is FlashCopy. This function is discussed in detail in Chapter 13, "FlashCopy" on page 367.

For all environments today, taking backups of data can take a very long time. Even though high availability storage is fault tolerant and protected by RAID, backups are still required to protect data from failures through logical errors and disasters. Backups are often taken outside a primary work shift because of the impact to normal operations. Databases must be closed to create consistency and data integrity, and the online systems are normally shutdown. Increasingly businesses cannot afford the down time to take these backups.

FlashCopy provides the ability to create a copy of your volumes (LUNs) immediately, whenever you need it. FlashCopy provides an instant copy of a logical disk. For example FlashCopy can enhance asynchronous backup and create test data from production data.

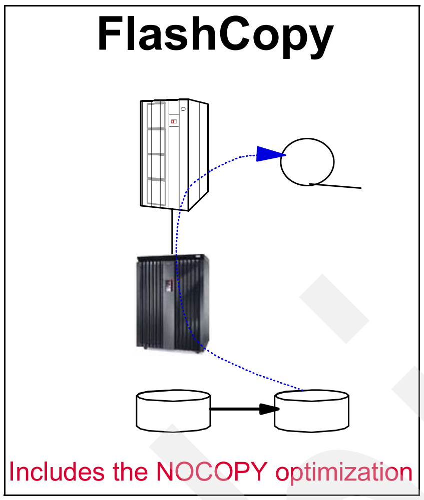

Figure 1-9 FlashCopy function

In most cases, FlashCopy still requires the system to be in a restricted or shutdown state. However, the FlashCopy is fast, so systems and applications can be restarted immediately after the command has completed. The copy of data is available immediately after the command is issued. There is no native OS/400 command interface to initiate the FlashCopy service. This function must be initiated from the ESS StorWatch Specialist.

iSeries backup and recovery offers a similar native function to FlashCopy. This function is called *save-while-active* and provides a T_0 (time-zero) point-in-time copy of the data specified. For details on Save-while-active, see "High availability versus disaster recovery" on page 396, or *Backup and Recovery Version 5*, SC41-5304.

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

PPRC-XD is a feature, available on the ESS, which leverages the proven ESS PPRC technology. Peer-to-Peer Remote Copy Extended Distance brings new flexibility to the IBM TotalStorage Enterprise Storage Server and PPRC, PPRC-XD is 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. For additional information and how to use it, refer to the redbook *IBM TotalStorage Enterprise Storage Server PPRC Extended Distance*, SG24-6568, which can be found at:

http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246568

At this time we do not know if PPRC-XD is a realistic option for iSeries servers as it probably requires an i5/OS shutdown to get a sync point. This may not be practical in many iSeries situations.

ESS PPRC Global Mirror

In 2004 IBM TotalStorage launched the Global Mirror function for ESS. Asynchronous PPRC (PPRC Global Mirror), a new function included in PPRC Version 2, is designed to enable a two-site disaster recovery and backup solution for the zSeries and open systems environments. Using asynchronous technology, PPRC Global Mirror operates over high-speed, Fibre Channel communication links and is designed to maintain a consistent and restartable copy of data at a remote site that can be located at virtually unlimited distance from the local site. The data consistency can even be managed across multiple ESS machines configured in a single session, with data currency at the remote site lagging behind the local site an average of 3 to 5 seconds for many environments.

Planned Features and Functions

- Two-site consistent asynchronous long distance copy solution for zSeries® and open systems data
- Mirrored data copy created independent from primary application write
- Consistency groups created across multiple ESSs
- Consistency groups can contain a mix of zSeries and open systems data
- Exploits PPRC Fibre Channel capability to promote optimal performance
- Designed to result in low host impact at consistency group formation time

Benefits

- Excellent RPO and RTO
- No active external controlling software required to form consistency groups
- No external automation scripts are required
- Excellent performance/Lower TCO
- Designed to offer lower costs

The sequence of events shown in Figure 1-10 describes the function of Global Mirror:

1. The system using the attached ESS experiences native performance on application, and performance transmission to remote ESS.
2. The Consistency Group formation takes about (2-3 ms). Then there is a return to native performance on application (tracking changes), the remaining consistent data is transmitted, and there is a high performance transmit to remote.
3. The remaining consistent data arrives at remote, the source system has native performance on application.
4. At the remote site there is an automatic FlashCopy to remote consistent volumes, native performance is maintained on the source application.
5. Performance transmission is resumed to remote, while incremental changes are sent and the sources system returns native performance on application.
6. Return to Step 1, If an outage occurs before next cycle completes, the consistent data is used to recover, again native performance is retained on the application.

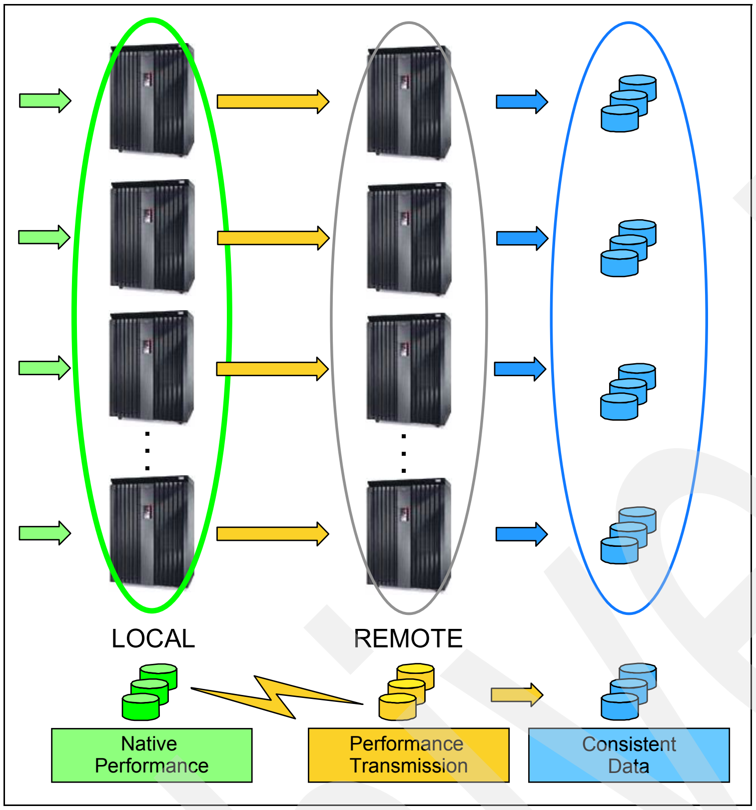

Figure 1-10 PPRC Global Mirror

At this time we are unaware of this application being tested with an iSeries. While the implementation is possible, we are unsure whether this would be a practical solution.

1.5 Storage Area Network

A SAN is a dedicated network that is separate from local area networks (LANs) and wide area networks (WANs) as shown in Figure 1-11. It generally refers to interconnected storage-related resources that are connected to one or more servers. It is characterized by high-interconnection data rates between member storage peripherals and by its highly scalable architecture. The term SAN usually refers to hardware but often includes specialized software for management, monitoring, and configuration.

A SAN differs from traditional networks, because it is constructed from storage interfaces. SAN solutions use a dedicated network behind the servers, based primarily (although, not necessarily) on Fibre Channel architecture. Fibre Channel provides a highly scalable bandwidth over long distances. And with the ability to provide full redundancy, it includes switched, parallel data paths to deliver high availability and high performance. Therefore, a SAN can bypass traditional network bottlenecks. It supports direct, high-speed transfers between servers and storage devices in the following ways:

- **Server to storage**: This is the traditional method of interaction with storage devices. The SAN advantage is that the same storage device may be accessed serially or concurrently by multiple servers.
- **Server to server**: This is high-speed, high-volume communications between servers.

- **Storage to storage**: For example, a disk array can back up its data directly to tape across the SAN, without processor intervention. Or, a device can be mirrored remotely across the SAN.

Not all of the features of a SAN discussed here are available on every platform.

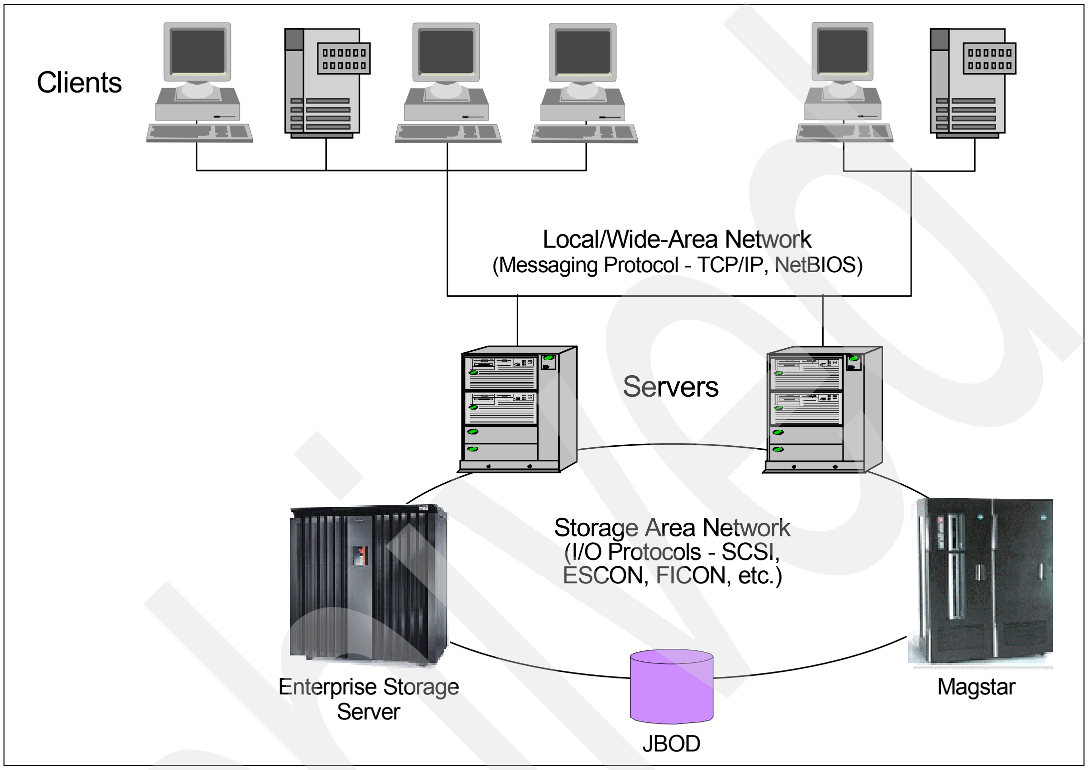

Figure 1-11 Storage Area Network

2

iSeries external storage positioning and examples

This chapter explains the positioning of external storage (ESS) in regard to the iSeries. This is necessary because the iSeries already has a very powerful and reliable integrated disk subsystem. Fibre Channel (FC) support is an additional function of OS/400 introduced at V5R1 and enhanced with V5R2. FC support offers iSeries customers even more flexibility over their system's environment.

This section covers the following topics:

- iSeries direct attach configuration
- Internal mirrored to external configuration
- Mirrored external configuration
- iSeries and SAN-attached tape drives configuration
- iSeries and PPRC configuration

2.1 Where external storage fits

The iSeries already has its own Storage Area Network (SAN)-like integrated storage subsystem, so why do you need external storage? That's a question many people are asking even as we write this book. Where the customer chooses to place their disk subsystem function is becoming a difficult decision. With an iSeries solution, customers can choose an integrated solution or a shared solution.

Many organizations are suffering from an enormous and hardly uncontrolled explosion of open systems (Windows NT® and UNIX). Storage consolidation is an attractive proposition for IT departments under extreme business pressure to maintain continuity.

Let's look at the scale of the problem that IT departments are facing. Figure 2-1 shows the lack of scalability in volume-based systems. As the user population increases, so do the proportions for the server contents. By the time we reach six systems, we effectively lose a disk to "spare" space. We also use nearly a disk in application space that will be replicated program code. In the data storage area, we duplicate space for such objects as price tables, directories, etc. This "lost" space can account for great asset values when multiplied across hundreds of servers. Businesses cannot afford to accept this redundancy.

Once the number of servers grows, the complexity of managing them grows even faster. Such simple tasks as monitoring and increasing storage can be major business problems. External disk storage subsystems seek to address these issues and do it well.

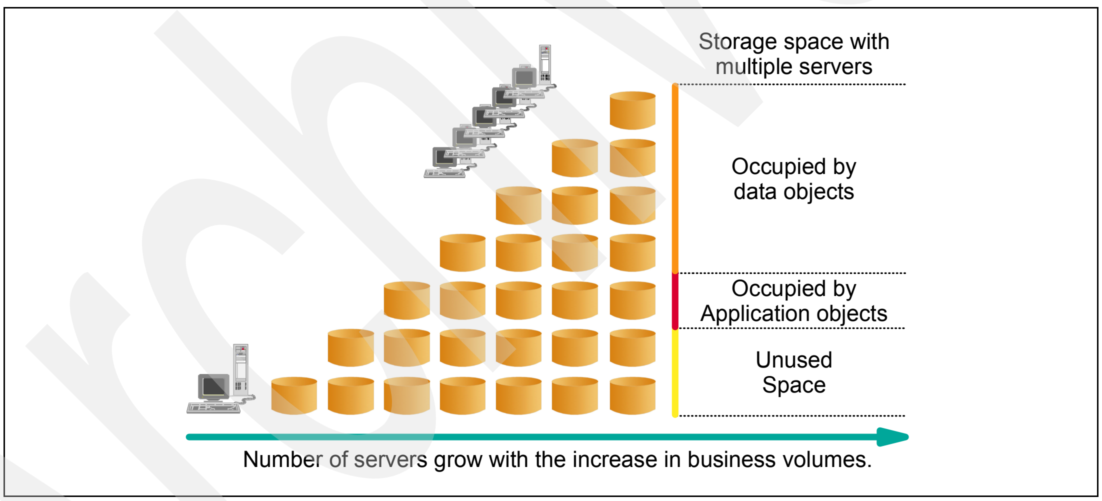

Figure 2-1 Impact of the growth of open systems

Most of the original implementation for a Web-base or e-commerce infrastructure was implemented on distributed PC servers and UNIX servers. These servers are generally single-purpose platforms with applications, which rely on data that is extracted from the line of business database server. The cost and complexities of distributed server implementations are well known, and the storage management piece of this topology is particularly noticeable.

UNIX and PC servers have basic storage management methods. Fundamentally they are manual environments that we are all familiar with as we consider our own desktop PCs. In addition to requiring investments in redundant skills and management software, the process is very labor intensive because the staff performs the same functions, multiple times across multiple servers, resulting into a higher cost of ownership.

SANs offers a powerful solution for consolidating storage resources to boost storage management, flexibility, and cost effectiveness. Serving as an indispensable part of every business e-infrastructure, SANs link multiple storage systems and servers over a secure, high-speed network that is data-centric rather than server-centric. They enable storage devices, such as tape libraries and disk arrays, to be shared. SANs can improve data availability, reduce downtime costs, decrease management and administrative costs, and improve asset utilization.

Therefore, a key premise of SAN is to enable a degree of consolidation and management centralization. The IT budget is broken down based on the costs associated with the distributed server environments:

- Operations
- Hardware
- System software
- Applications
- The central site

Most of the budget goes to operations. The question is, do we look at the management cost and complexity associated with a server consolidation strategy, with a storage consolidation strategy, or a mixture of both? The fact is that server consolidation is not a truly viable option for distributed servers. Neither UNIX or PC server architectures are conducive to application and database consolidation. Therefore, it is natural to look at the storage aspects during a consolidation to find a resolution through centralized storage management.

IBM servers offer a solution for both server consolidation and storage consolidation.

In Chapter 1, "Introduction to iSeries and Storage Area Networks (SAN)" on page 3, we describe how the iSeries deals with the same problem. Single-level storage is an extremely powerful tool for managing growth. Users, applications, and files can be added to the system with little or no impact to the business.

So what do you do? There is no easy answer, but here are some ideas:

- If you are already using the iSeries and are under no pressure to consolidate, do nothing. OS/400 and your iSeries hardware will manage your disk for you.
- Customers who have a large iSeries storage capacity will no doubt run mission-critical applications. They may already looking at a higher level of availability than consolidated storage can offer at this time. Therefore, the best option for them is to go for an iSeries cluster with OS/400 replication services.
- If there are a limited number of open systems servers, consider using the iSeries as a SAN integrated environment. UNIX servers can be accommodated in the iSeries OS/400 PASE (Portable Application Solutions Environment) environment and Windows-based servers in the Windows integration environment.
- If your iSeries part of the IT infrastructure, and there are many Open System servers to be consolidated, an ESS is a great solution. Make sure you understand whether the co-existence of these diverse applications provides acceptable availability.
- If you are using older #6501 technology to attach to ESS or other external storage, then FC support and ESS are the best option. It offers integrated storage support with improved performance, without the unsupported solutions or limitations of other external storage offerings.
- Fibre Channel tape holds many advantages over Small Computer Systems Interface (SCSI) attach. The ability to attach many systems over distance offers great business and irrecoverably enhancements.
- For those customers with no current or limited disaster recovery (DR) capabilities and who don't fit into one of the integrated storage examples, Peer-to-Peer Remote Copy (PPRC) may provide a DR solution for these business needs.

2.2 Examples of iSeries SAN environments

The following examples show various scenarios that you can implement with the iSeries server. While ESS can attach multiple platforms within the same storage server, other platforms cannot be on the same arbitrated loop as the iSeries server. The iSeries can share the same ESS with different platforms, but the workload mix must be well understood. If the iSeries is running a mission-critical application, use care before mixing applications and platform types within the same ESS cluster/server. This warning applies to any consolidation scheme whether server or storage where application response is critical.

The ESS design offers maximum performance and flexibility by opening up cache to everyone. The downside of this is that there is no prioritization of workloads, and cache works on a first come, first served basis. The differences in data blocking can affect a number of the system’s abilities to access the single ESS resource.

If you are planning to make the significant investment in consolidated storage, it is well worth benchmarking.

2.2.1 iSeries direct attach configuration

In this configuration, the iSeries has one internal drive, which is the load source unit (LSU). The remainder of the drives are located in the ESS. While this may be the most common configuration, it is not necessarily the best option for availability. An internal/external mirrored approach can offer more resiliency.

The example in Figure 2-2 shows an iSeries with a single 8.59 GB internal drive that is the LSU. The ESS contains 24 x 8.59 GB LUNs that represent the LSU and 23 additional disk units. This provides 206.16 GB of total storage. These LUNs are spread across four arrays and two SSA loops. The loops are within one cluster. No other systems have their LUNs in the same arrays as the iSeries LUNs. In this arrangement, a small amount of residual storage is unusable. The residual changes depend on the drive size and LUN size.

In this arrangement, backups are performed from the iSeries using standard commands, Backup Recovery and Media Services (BRMS), or another backup utility.

The ESS could be located some distance from the processor. Direct-attached ESS can be up to 500 m from the iSeries server depending on the fibre type. Alternatively, when using hubs, the ESS can be up to 10 km away.

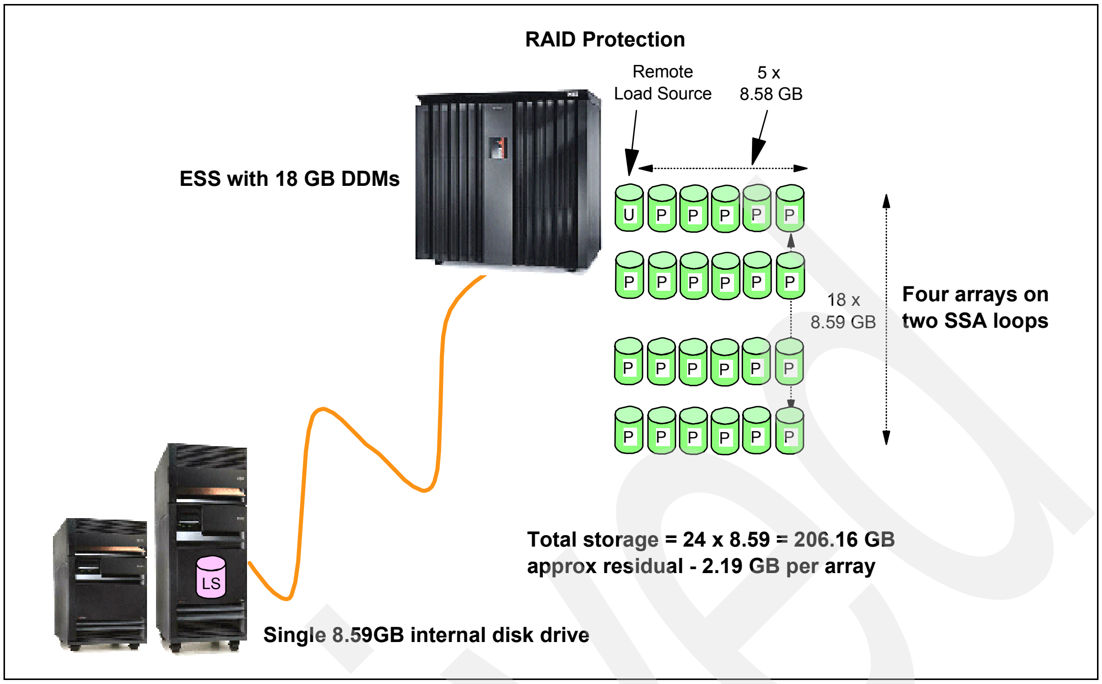

Figure 2-2 Basic ESS-attached LUNs

A degree of disaster recovery can be offered with this arrangement. If the ESS is distant and a second iSeries is available at the remote local that connects to the distant ESS, in the event of a disaster at the source iSeries, the second iSeries or logical partition (LPAR) can recover the ESS LUNs. This recovery requires an IPL, recover load source, and an abnormal IPL.

2.2.2 Internal mirrored to external configuration

The mirrored configurations offer a simple solution that meets most needs. It offers some storage server consolidation, but retains the characteristics of the iSeries. This solution can be implemented where both iSeries and ESS exist and the ESS is used to accommodate the iSeries mirrored drives (Figure 2-3).

Figure 2-3 iSeries internal disk mirrored to ESS

Implementation requires careful consideration. The ESS LUNs must be created as “unprotected”. This allows the iSeries internal drives to mirror to the ESS drives. In this scenario, the ESS drives retain their RAID protection, which is transparent to the iSeries.

The LUNs on the ESS must match the size of their corresponding iSeries drives. Any mismatch causes the iSeries mirroring task to fail. When the ESS LUNs are created, they are tagged as unprotected. Once formatted, they appear to the iSeries as non-configured drives. The drives are added to the iSeries ASP and then mirroring starts.

Figure 2-3 shows a hub, which is not necessary, but in this case, it provides additional redundancy. The FC from the iSeries to the ESS is a single point of failure, as is the ESS host bay. To provide an additional path to another host bay, an addition FC is provided from the hub. In the event that the host bay on the ESS needs to be re-IPLed, the second path from the hub to a different ESS host bay allows access to the iSeries LUNs.

If we introduce a second hub connected to the first using longwave (LW) GBICs, the ESS can be located up to 10 km away. This offers disaster recovery if a second iSeries is available at the remote location. This second iSeries needs to have an equivalent number of internal drives. Their size should match if mirroring on the second system is also required. However, for a DR situation, it could run in a degraded manner with only the internal LSU and external ESS LUNs.

2.2.3 Mirrored external configuration

This example is an extension of the previous one (internal mirrored to external) and requires careful planning. In this case, there is only one internal drive in the iSeries, and two ESS are available (Figure 2-4).

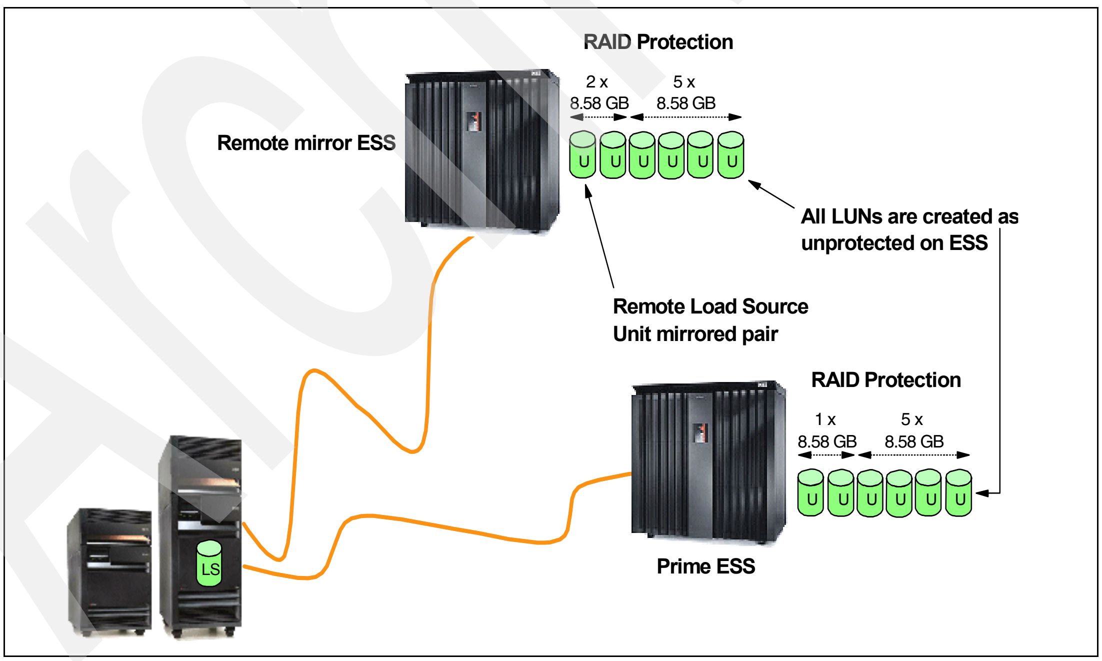

Figure 2-4 iSeries to mirrored ESS

One ESS is nominated as the “prime” machine and the other is its remote mirrored pair. All LUNs on both ESSs are created as unprotected.

Initially the system is setup with a non-mirrored internal load source unit and unprotected (from an iSeries point of view) LUNs on the primary ESS. The second ESS is then assigned to the iSeries. Its LUNs are also unprotected. At this point, the internal load source can mirror to this remote ESS, along with the other LUNs on the primary ESS.

This configuration provides a disaster recovery capability for both the iSeries and ESS. The iSeries and two ESS can be remote from each other. However, remote load source recovery is impossible from the primary ESS because it does not contain a remote load source mirror.

This potential problem can be overcome by simply creating a PPRC copy of the remote load source mirrored LUN back to the primary ESS.

2.2.4 iSeries and SAN-attached tape drives configuration

The example in Figure 2-5 will be a popular scenario for iSeries customers, who until now, have had limited distance between their tape subsystems and processor. With FC and SAN capabilities, they can now move their tapes up to 10 km away from their systems.

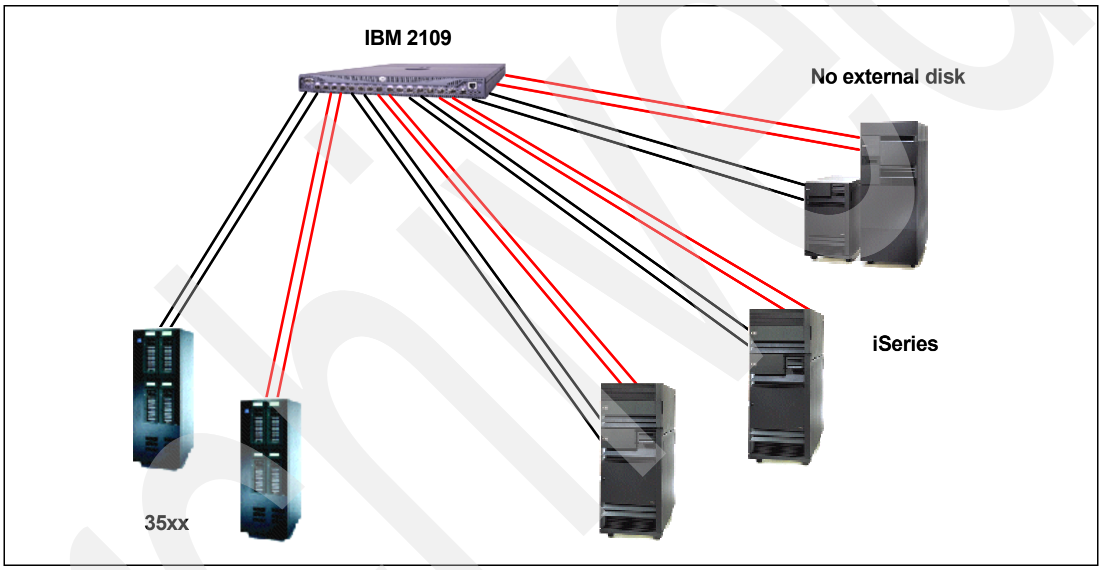

Figure 2-5 iSeries with local SAN-attached tape

Three iSeries have two #2765 adapters to allow access to both tape 35xx drives. Each tape drive is on a separate zone within the SAN.

The example in Figure 2-6 shows a similar scenario with three iSeries and multiple tape drives. However, in this case, the tape drives and one system are remote. You can view this as a disaster recovery solution. The remote iSeries offering can take over the workload of the two local machines in the event of a disaster at the local site. The remote machine can be divided into logical partitions (LPAR) to help restore both local machines' workloads at the remote site in the event of a failure.

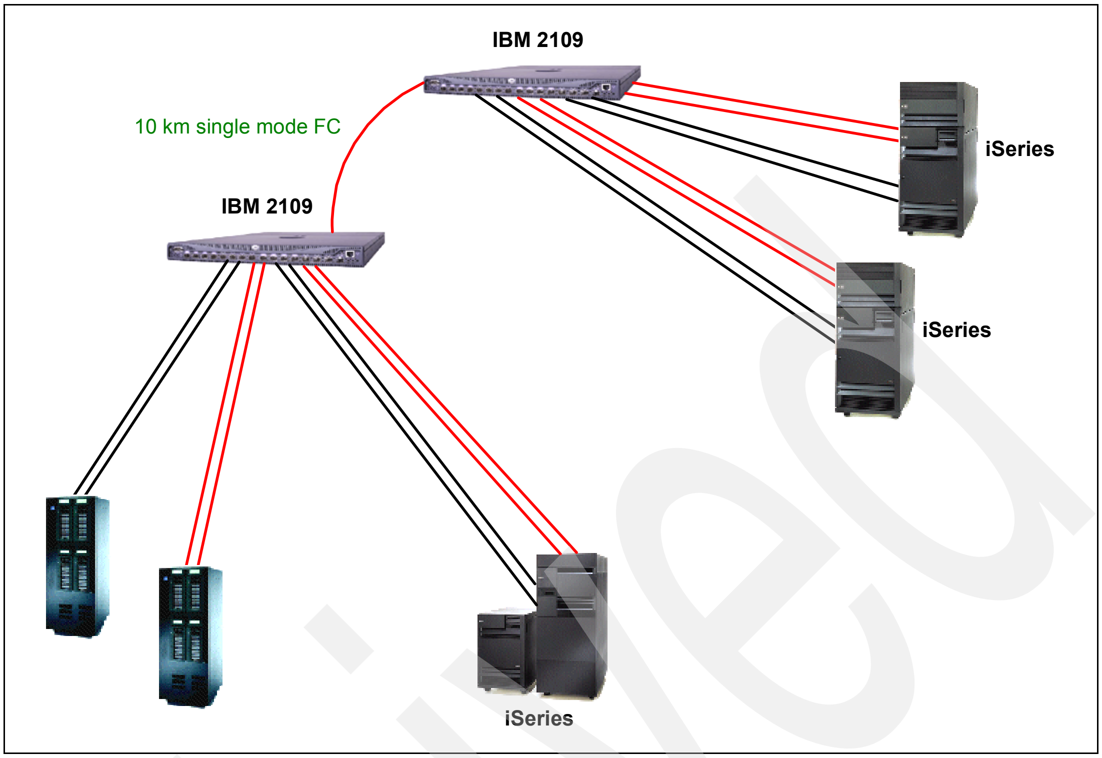

Figure 2-6 iSeries with distant tape

You can further extend this solution into a clustered environment, enhancing the DR capability to full high availability with replication software. This allows administrative switchover for routine maintenance activity on the local site.

2.2.5 iSeries and PPRC configuration

Peer-to-Peer Remote Copy is now available for SCSI and FC-attached ESS disk. Figure 2-7 shows an arrangement of the iSeries, ESS, and PPRC.

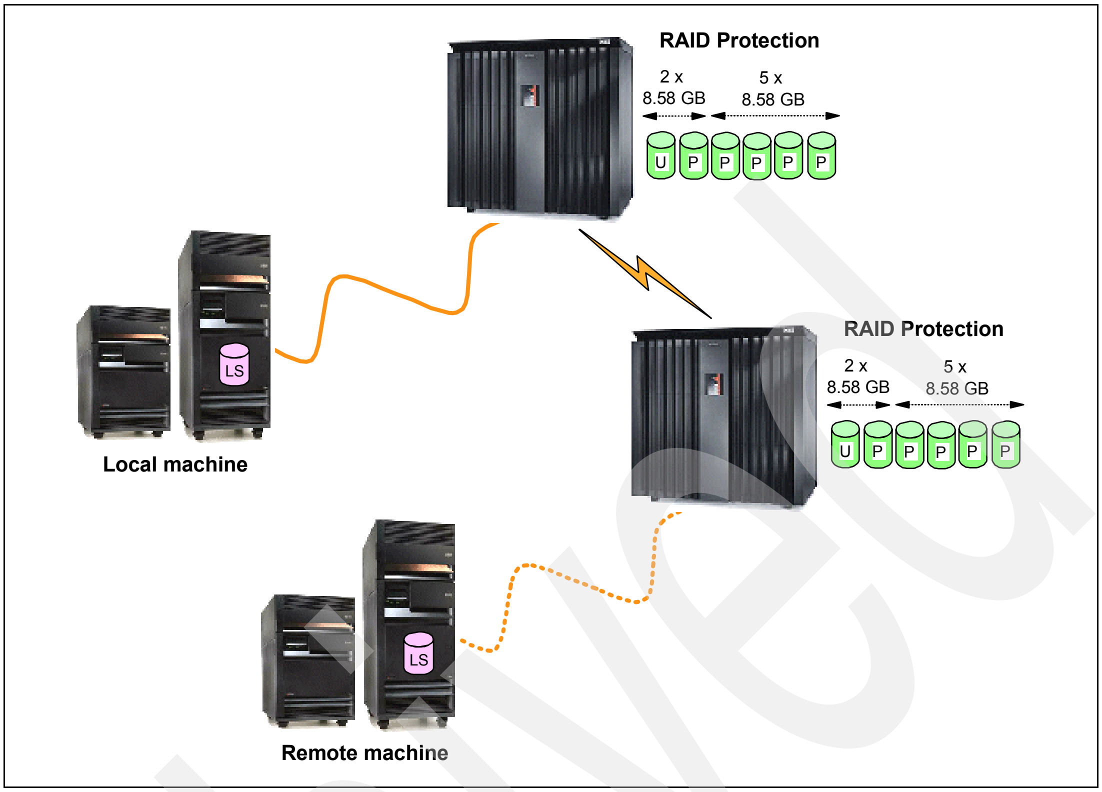

Figure 2-7 iSeries and ESS PPRC

The local iSeries has only one internal drive, and this is the load source unit. Within the local ESS, there is the remote load source mirrored pair and the remainder of the iSeries LUNs. On the remote site, there is another ESS. This storage server contains the PPRC copy of the remote load source unit and all other LUNs from the local iSeries. The backup iSeries server is powered down, but the LUNs are assigned. The backup server could be a separate server or a partition within an logically partitioned system.

In the event of a disaster at the production site, the backup iSeries recovers to the backup ESS LUNs. This recovery process includes an abnormal IPL on the iSeries. This IPL can take several hours. You can reduce this time by implementing OS/400 availability functions to protect applications and database, for example, journaling, commitment control, system managed access path protection (SMAPP).

You can learn more about PPRC in Chapter 12, “Peer-to-Peer Remote Copy” on page 327.

2.2.6 iSeries and FlashCopy configuration

FlashCopy allows a T_0 (time-zero) point-in-time copy to be made within the ESS. Figure 2-8 shows an arrangement of the iSeries, ESS, and PPRC.

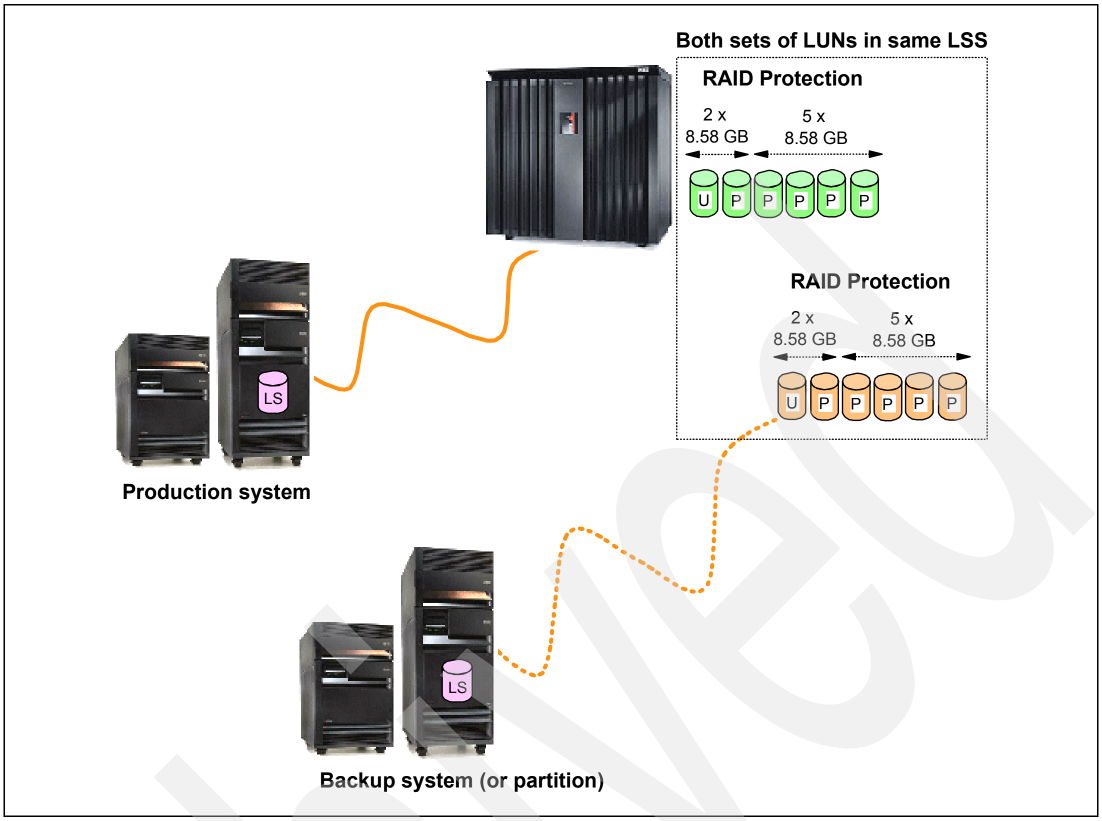

Figure 2-8 iSeries and ESS FlashCopy

The production iSeries has only one internal drive, which is the load source unit. Within the ESS, there is the remote load source mirrored pair and the remainder of the iSeries LUNs. FlashCopy creates another copy of all the LUN s (mirrored load source unit and iSeries LUNs) inside the same ESS.

The backup iSeries server is normally powered down, but the LUNs are assigned. The backup server could be a separate server or a partition within a logically partitioned system.

To do backups from the FlashCopy set of LUNs, the backup iSeries recovers to the backup ESS LUNs. This recovery process may include an abnormal IPL on the backup server, depending on whether the production server was shut down or simply taken to a restricted state. A complete shutdown of the production server allows the backup server to have a normal IPL. Taking the production server to a restricted state requires an abnormal IPL of the backup server, but this is quick and almost the same as a normal IPL. The production server only needs to be in a restricted state for a short period of time while the FlashCopy is made. this allows memory to be flushed to the ESS and achieves a “known state” for the database to achieve consistency and integrity.

You can learn more about using FlashCopy for backups in Chapter 12, “Peer-to-Peer Remote Copy” on page 327, and Chapter 13, “FlashCopy” on page 367.

2.2.7 Additional use of ESS FlashCopy

ESS FlashCopy can be combined with other ESS connectivity options. The example discussed in 2.2.6, “iSeries and FlashCopy configuration” on page 29, is really an extension of the basic use of ESS discussed in 2.2.1, “iSeries direct attach configuration” on page 24. Similarly, it is possible to use FlashCopy with the configurations discussed in 2.2.2, “Internal mirrored to external configuration” on page 25, 2.2.3, “Mirrored external configuration” on page 26, and 2.2.5, “iSeries and PPRC configuration” on page 28, to get a point-in-time copy offsite from the production system.

You can use the copies of the data taken with FlashCopy to take offline backups or for other purposes, such as populating a DataWarehouse, providing a restart point for a long-running batch suite, or data migration.

Part 2

SAN

Storage Area Network (SAN) is the latest technology for network attached disk and tape. Part 2 provides basic details for a SAN implementation related to iSeries.

3

Storage Area Networks (SAN)

This chapter discusses the SAN components:

- Concepts of SAN
- 2109 Switches
- Zoning of switches with iSeries

Where we refer to OS/400 V5R2, in this chapter, this is the earliest release. The same functions and scenarios are supported at V5R3.

3.1 Fibre Channel topologies

Fibre Channel is a network architecture that combines the “channel transport” characteristics of an I/O bus with the flexible connectivity and distance characteristics of a traditional network. Notice the European spelling of Fibre. It is intended to distinguish it from fiber-optics and fiber-optic cabling. Fiber-optics and fiber-optic cabling are physical hardware and media used to transmit data at high speed over long distances using light emitting diode (LED) and laser technology.

Because the Fibre Channel architecture has channel-like qualities, hosts and applications see storage devices attached to the SAN as though they are locally attached storage. With its network characteristics, Fibre Channel can support multiple protocols and a broad range of devices. It can be managed as a network. Fibre Channel architecture can use either optical fiber (for distance) or copper cable links (for short distance at low cost), although you predominantly see optical fiber.

Fibre Channel is a multi-layered network architecture based on a series of American National Standards Institute (ANSI) standards that define characteristics and functions for moving data across the network. These include definitions of such physical interfaces as cabling, distances, and signaling; data encoding and link controls; data delivery in terms of frames, flow control, and classes of service; common services; and protocol interfaces.

Like other networks, information is sent in structured packets or frames, and data is serialized before transmission. But unlike other networks, the Fibre Channel architecture includes a significant amount of hardware processing to deliver high performance. The speed currently achieved is 200 MB per second, (with the potential for 400 MB and higher data rates in the future). In all Fibre Channel topologies, a single transmitter sends information to a single receiver. In most multi-user implementations, routing information (source and target) must be provided. Transmission is defined in the Fibre Channel standards across three transport topologies:

- Point-to-point
- Arbitrated loop
- Switched (referred to as a fabric)

The topologies are shown in Figure 3-1.

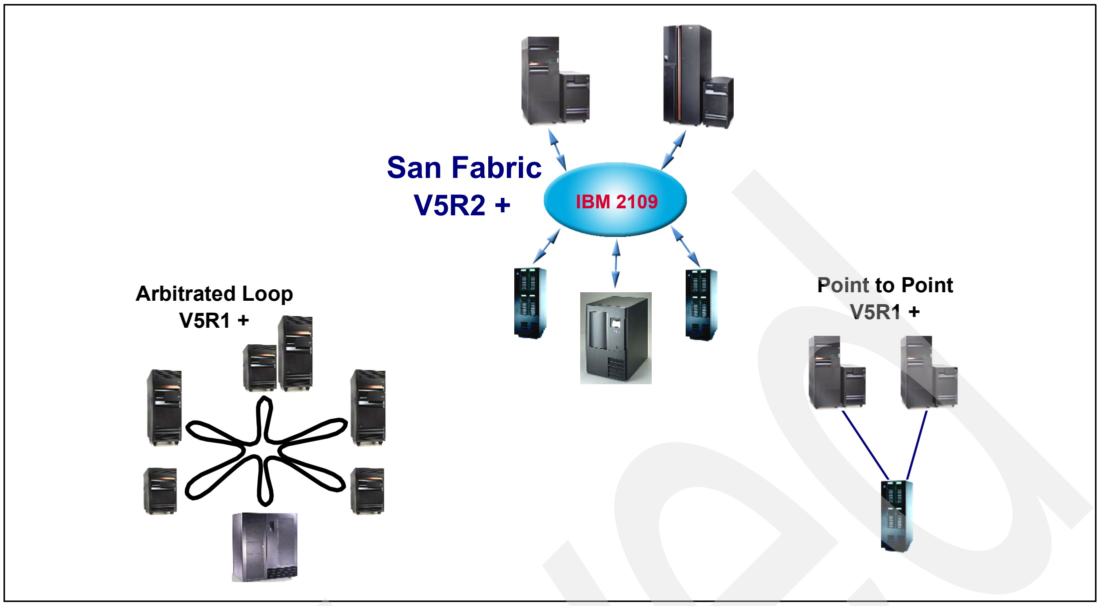

Figure 3-1 SAN fabric support

For details about Fibre Channel and Storage Area Network, refer to the Redbooks:

- *Designing an IBM Storage Area Network,* SG24-5758
- *Implementing an Open IBM SAN*, SG24-6116

Switched fabric is the preferred method of connection. Although no formal announcements have been made, it is expected that Arbitrated Loop support will have a limited life.

3.1.1 Point-to-Point

The simplest topology is a point-to-point bi-directional connection between two nodes (Figure 3-2).

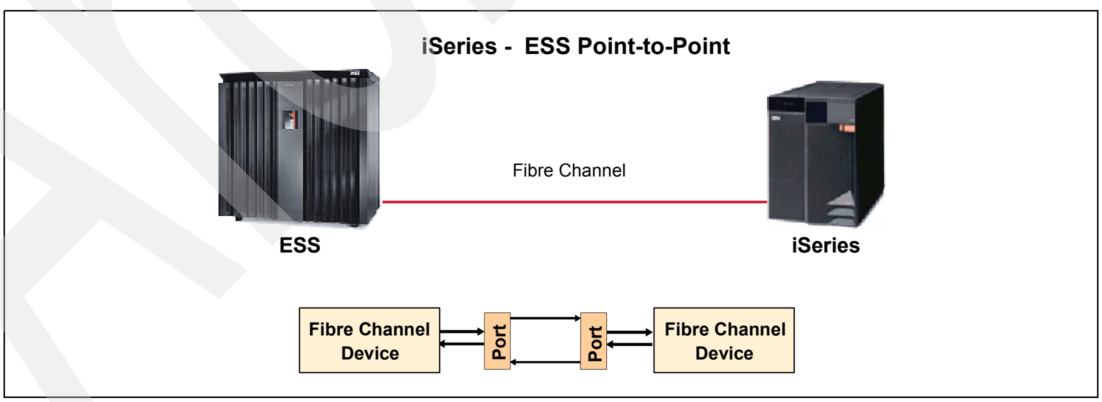

Figure 3-2 Point-to-point topology

By using a fibre cable, two Fibre Channel adapters (one host and one storage device) are interconnected. It supports the maximum bandwidth capability of Fibre Channel. It does not provide sufficient connectivity to support complex configurations but is the easiest Fibre Channel configuration to implement and administer.

3.1.2 Arbitrated loop

Arbitrated loop topology is a unidirectional ring topology, similar to token ring. Up to a maximum of 126 nodes can be interconnected via a looped interface. Information is routed around the loop and repeated by intermediate ports until it arrives at its destination. The Fibre Channel ports that support this topology must perform these routing and repeating functions in addition to all the functions required by the point-to-point ports.

Fibre Channel arbitrated loop offers relatively high bandwidth and connectivity at a low cost. For a node to transfer data, it must first arbitrate to win control of the loop. When the node has control, it is free to establish a point-to-point (virtual) connection with another node on the loop. After the point-to-point (virtual) connection is established, the two nodes consume all of the loop's bandwidth until the data transfer operation is complete.

A commonly used variation in this topology is the Fibre Channel hub. A hub incorporates the structure of the Fibre Channel arbitrated loop in a package that provides ports that are physically connected, similar to those of a switch. Figure 3-3 shows an example.

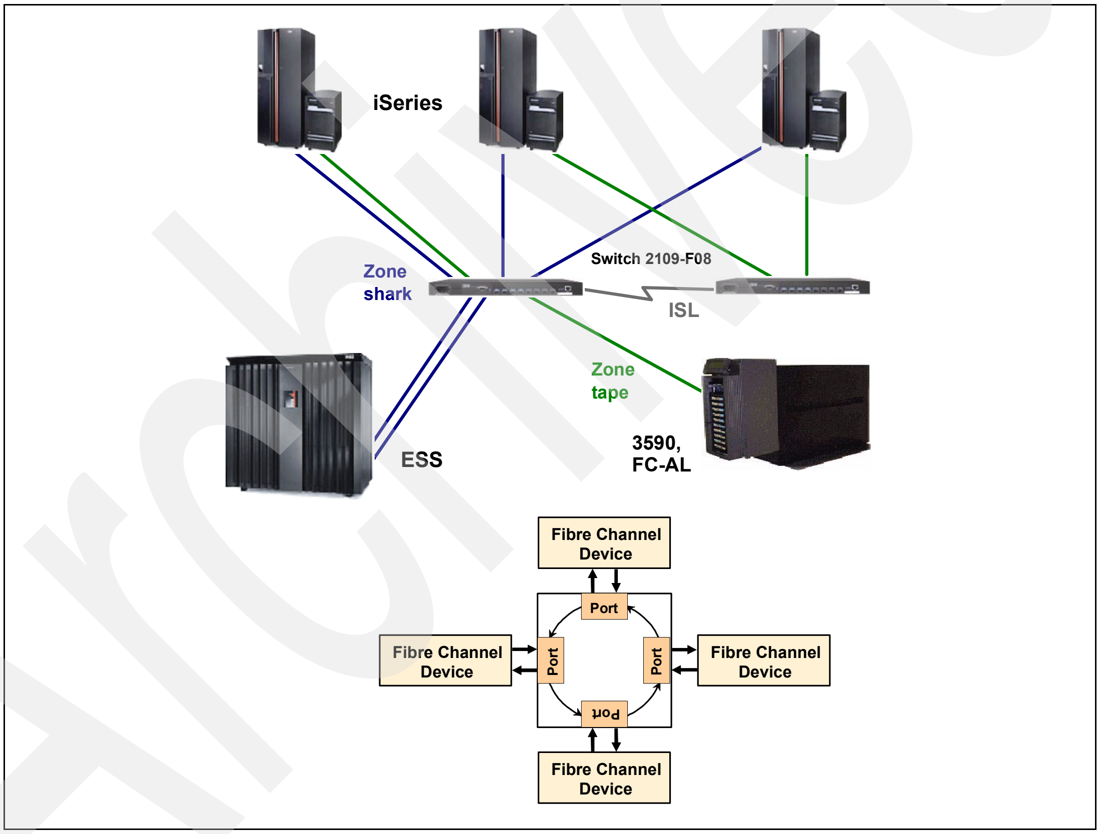

Figure 3-3 Fibre Channel arbitrated loop

There are two different kinds of loops:

- **Private loop**: The private loop does not connect with a fabric, only to other private nodes using attachment points called *Node Loop Ports* (NL_Ports). A private loop is enclosed and known only to itself. NL_Ports can only attach to other NL_Ports or to Fabric Loop Ports (FL_Ports).
- **Public loop**: A public loop requires a fabric and has at least one FL_Port connection to a fabric. A public loop extends the reach of the loop topology by attaching the loop to a fabric. FL_Ports can only attach to NL_Ports.

Private devices in public fabrics

The characteristic of a fabric allows communication between any host or storage device in the fabric. In other words, all communication is “public”. Problems may arise when a private Fibre Channel device is attached to a switch. Private Fibre Channel devices were designed to only work in private loops.

For details about how node and port addressing works, refer to *Designing an IBM* Storage Area Network, SG24-5758.

QuickLoop feature

QuickLoop creates a unique Fibre Channel topology, this allows ESS host bus adapters (HBAs) to use Fibre Channel Arbitrated Loop (FC-AL) connections without knowledge of the SAN fabric. An adapter cannot perform fabric login (commonly referred to as *private loop initiators*), to communicate with Fibre Channel Arbitrated Loop public storage devices. QuickLoop can be implemented in IBM 2109 Fibre Channel Switches.

QuickLoop allows individual switch ports to be designated as arbitrated loop ports. This allows a private loop host initiator to communicate with arbitrated loop storage devices as though they were all contained in one logical loop.

The QuickLoop feature also supports legacy devices such as those in a Private Loop Direct Attach (PLDA) environment. The QuickLoop feature allows these devices to be attached to a SAN and operate no differently than in a PLDA environment.

3.1.3 Switched fabric

Switched fabric refers to whenever a switch is used to interconnect Fibre Channel adapters. It is an intelligent switching infrastructure that delivers data from any source to any destination (see Figure 3-4). Functions on switches and storage devices, such as zoning and Logical Unit Number (LUN) masking, prevent servers from accessing storage that they are not allowed to access.

Switched fabric is currently *only* supported on the iSeries at V5R2 and later.

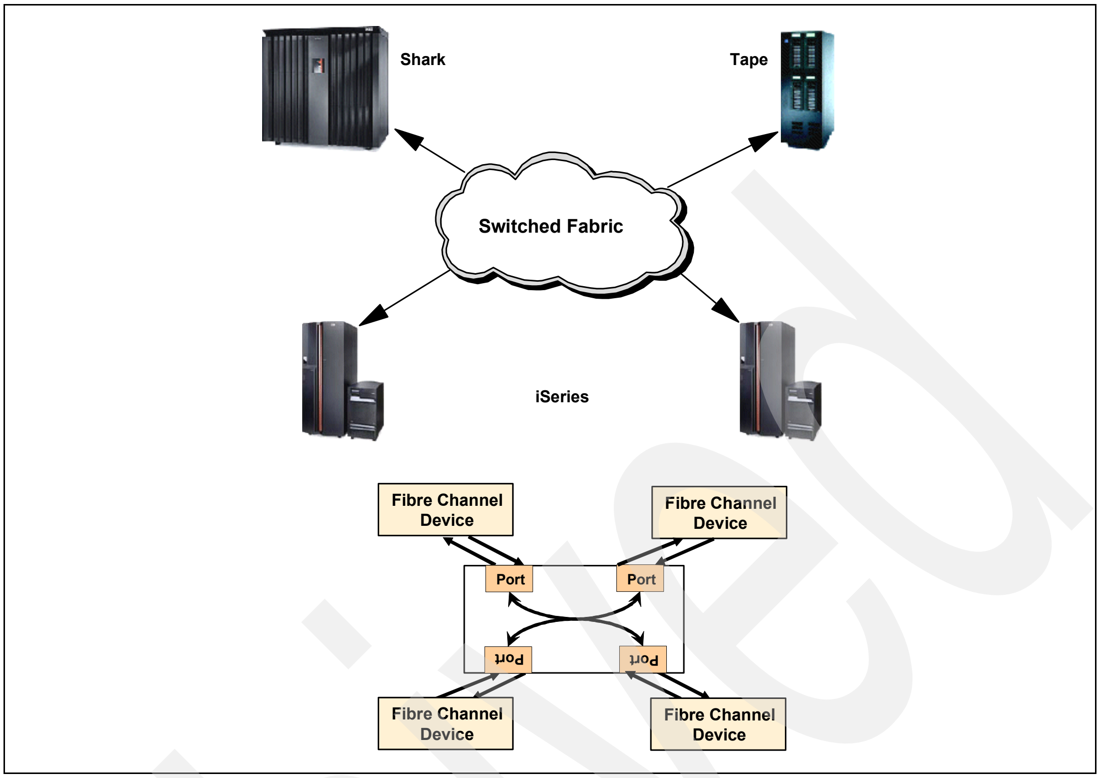

Figure 3-4 Switched fabric SAN

3.2 Making the right connections

A SAN enables disk storage capacity to be consolidated. A SAN attached with fiber presents technical issues that you must resolve:

- Some hosts take ownership of any visible LUNs on attached cables or on the SAN.
- Two or more hosts connected to the same LUN unknowingly compete for the same storage device.

Various combinations of connecting the iSeries through the switch are discussed below

Fibre Channel SANs are often zoned to block other servers from connecting to the same storage device. Zoning can be implemented in either the hardware or software. LUN masking is another means to prevent unwanted access to a server's storage devices.

3.2.1 Cascaded switches

In Version 5 Release 1, a maximum of two cascaded switches are supported. With V5R2, when using switched fabric, more than two cascaded switches are allowed. This greatly increase the flexibility of adding an iSeries to a SAN and allows for increased redundancy to be built into the SAN fabric.

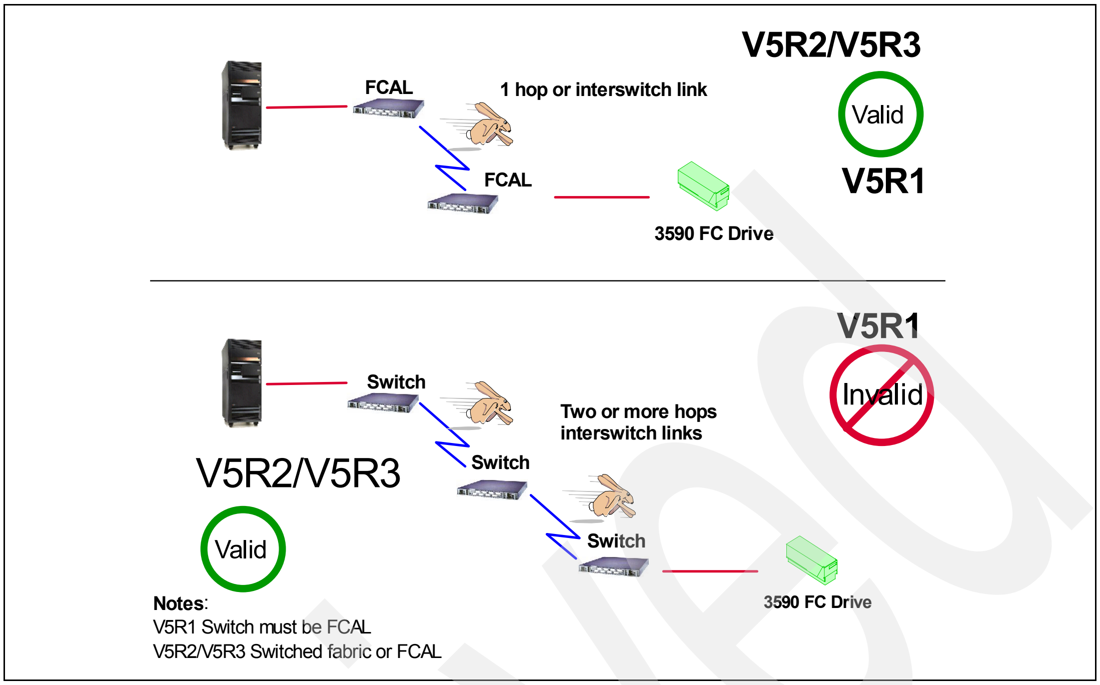

Figure 3-5 Interswitch links

There can be a distance of 10 KM between the switches with V5R2 or later support at 2 Gbps, and when using a 9 micron cable.This assumes there will be only two switches between the initiator and target. We recommend that you use no more than three interswitch links between the initiator and target for performance reasons. This is not a hard limit, but rather a guideline to ensure good performance. This is illustrated in Figure 3-6. You could increase the distance out as far as 100KM, but performance would degrade as distance is increased.

Support for more than two cascaded switches allows the iSeries to fully participate in more complex SAN fabrics. A 10 KM limit remains between switches, so you should have only one 10 KM hop between switches.

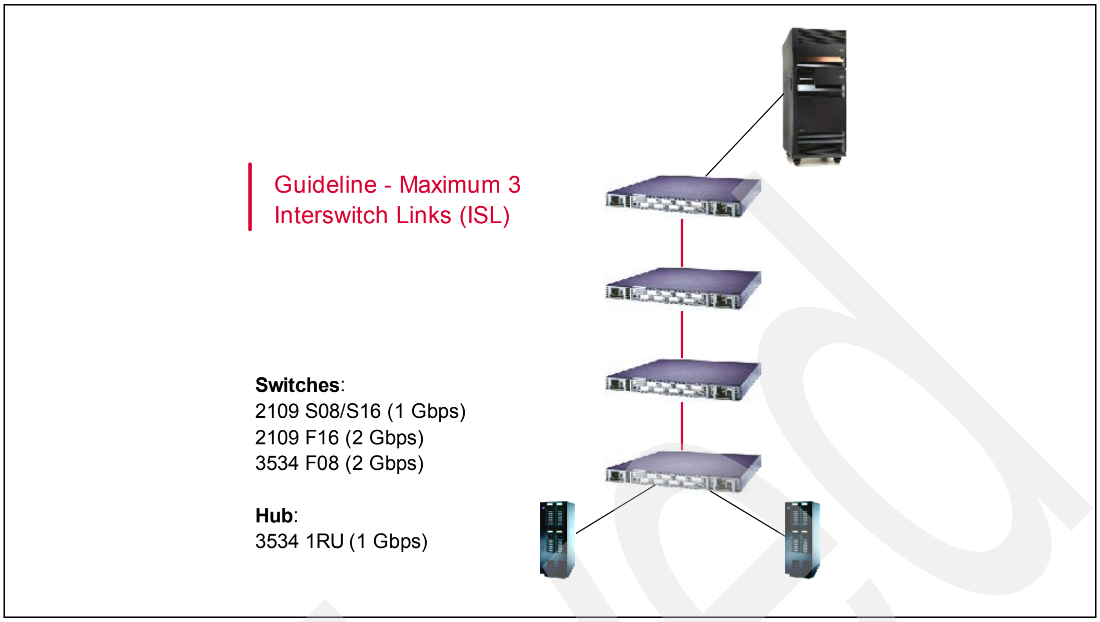

Figure 3-6 Maximum of three interswitch links

3.2.2 Fibre Channel Arbitrated Loop (FC-AL) over switched fabric

It is now possible to have an FC-AL loop over a switched fabric. This is supported for both V5R1, V5R2 and V5R3. This enables you to integrate V5R1 systems into a switched fabric environment, still using FC-AL.

In this scenario (Figure 3-7), the first switch (connected to the iSeries) is logically connected to a final switch connected to the target tape device. It is a port-to-port connection using FC-AL.

In this scenario, the switches implement a QuickLoop or FCAL over a SAN fabric. This is a function of the SAN infrastructure, not the iSeries. In fact, the iSeries has no knowledge of the SAN fabric. This allows a V5R1 system to be fully integrated into a pre-existing SAN fabric and the V5R2/V5R3 system to connect using the preferred switched fabric option.

Other rules apply, such as only one target for the V5R1 system.

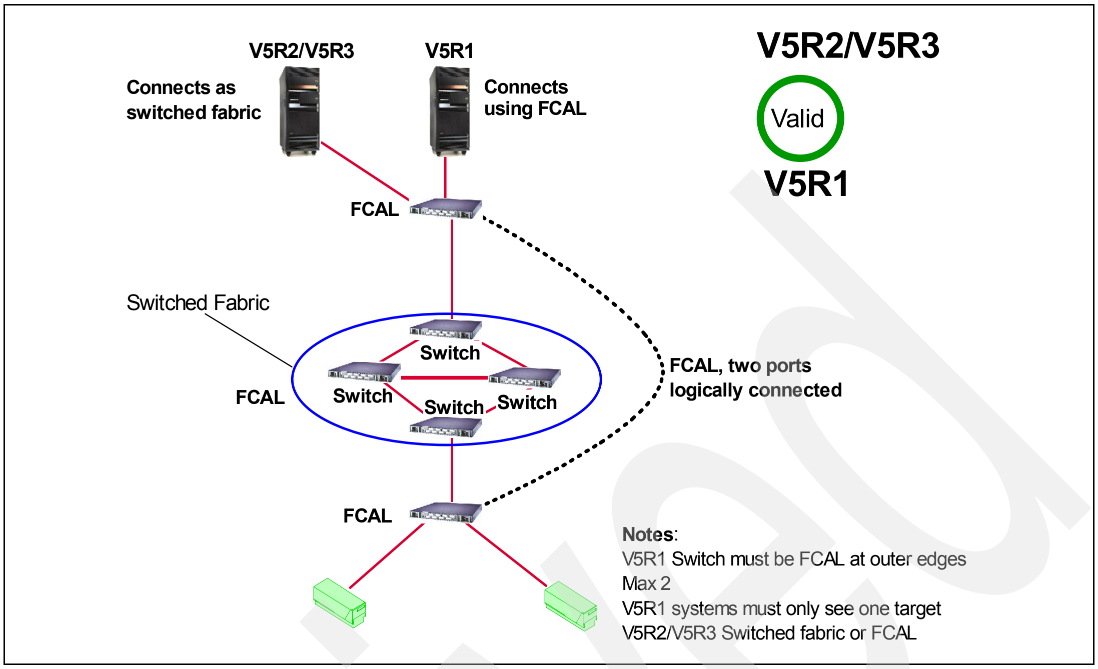

Figure 3-7   FCAL over switched fabric

3.2.3 Distance

Due to the increased speed support from 1 Gbps to 2 Gbps, the total distance supported is shorter when 2 Gbps is enabled. The new limits are shown in Figure 3-8, for short wave, and in Figure 3-9, for long wave.

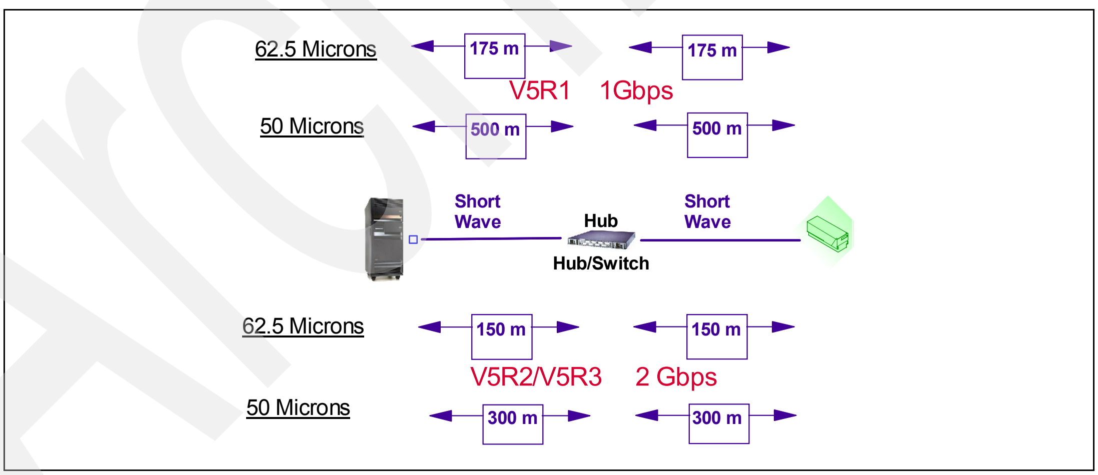

Figure 3-8   SAN distance limits: Short wave

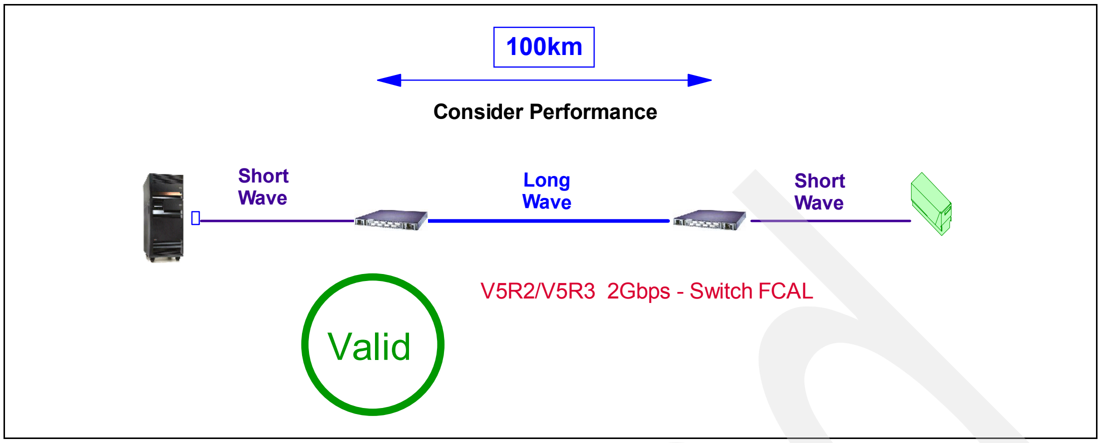

Figure 3-9 SAN distance support: Long wave

3.2.4 Multiple target support

Version 5 Release 2 and later releases include support for multiple targets from a single Fibre Channel tape or disk adapters. At Version 5 Release 1, only a single target (a single tape drive) is supported from an initiator (FC adapter). At V5R2 or later, up to 16 targets are supported for FC tape drives.

Multiple target support with tape

With tape, a single tape FC adapter can support multiple tape devices. For example, at V5R1, if you had four iSeries with four FC tape devices, 16 FC tape adapters would be required to enable all iSeries to address all tape devices, since each FC adapter is only allowed to see one target. At V5R2/V5R3, this number is reduced to four. That is one FC tape adapter per iSeries since each FC adapter may see multiple targets as shown in Figure 3-10

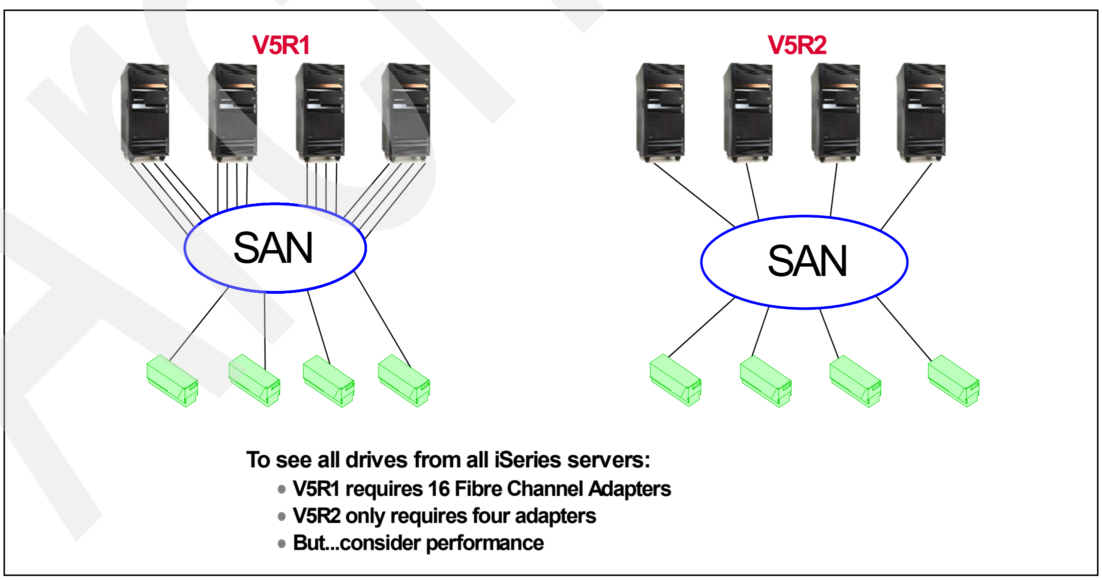

Figure 3-10 Multiple target support

A maximum of 16 tape LUNs may be addressed from a single iSeries tape FC adapter. A LUN is a tape device or media changer. For example, 3583 LTO tape library, which may have six drives and a media changer, reports as seven LUNs.

While the previous example dramatically improves the cost and hardware required to support multiple devices from multiple iSeries, you must also consider:

- Performance: Plan on a sustained speed of 100 MB/s (1 Gbps) 2Gbps with V5R2 and later. For Models 820 and 270, plan on 80 MB/s.
- Management of tape devices

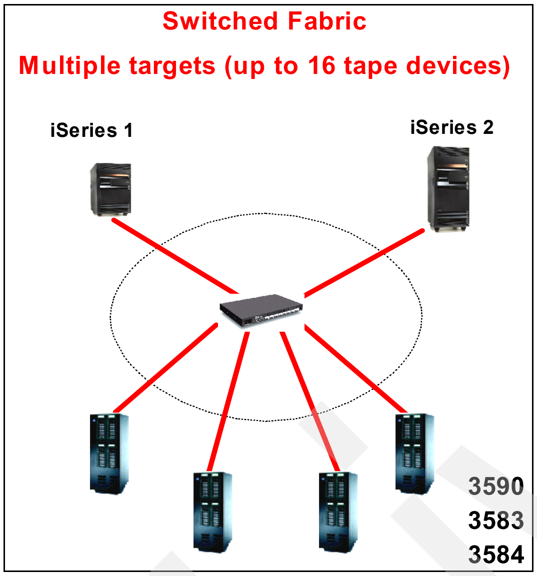

Figure 3-11 16 target tape LUNs

To share the devices, only a standard “vary on” and “vary off” command is required. When the iSeries varies on a tape, it uses a “reserve” lock on the tape device to prevent other systems using the drive. However, in a heterogeneous setup, other operating systems also share the drive.

A benefit of the iSeries is its ability to “reserve/release” a tape drive. This prevents other systems writing to the middle of a tape in use by an iSeries, and it ensures tape data integrity for iSeries users. The same may not be true for other operating systems. Therefore, it is entirely possible for an iSeries to start using a tape drive in use by another system. You need to check with each vendor to establish whether they support “reserve/release” of a tape drive.

Tape management software, such as BRMS and Tivoli, use the “reserve” lock. For other packages, check with your software vendor.

Multiple target support with ESS

The Enterprise Storage Server remains the only external disk supported by Fibre Channel adapters. With multiple target support in V5R2 or later, many ESSs may be supported from a single FC adapter initiator. With V5R1, the limit is one target (ESS).

The total number of LUNs addressable remains at 32 for both V5R1 and V5R2 and is the gating factor for the number of ESS targets reachable. For example, a single iSeries FC disk adapter can have 16 LUNs, each on two ESSs. You could even have one FC adapter with one LUN on each of 32 ESSs. At present we cannot see a reason for such an implementation.

The internal limit on the number of ESS targets supported at V5R2 or later is quite a large number, however the LUNs per FC adapter limit of 32 is always reached first.

3.2.5 Maintenance requirements

As SANs become increasingly central to everyday business operations, maximizing uptime in these mission-critical environment poses increasingly difficult challenges. Scheduled and unscheduled downtimes for maintenance should be planned into a SAN. A SAN having multiple components would need a scheduled downtime under various circumstances. We have discussed a few points below which will help in planning and maintaining a SAN:

Microcode update

Microcode updates on the SAN switches are to be considered when planning a SAN. Microcode updates on some switches affect the ports connected to the switch. Though most SAN Directors allow for a concurrent microcode update, the other switches do not have this feature.

Monitoring and Managing

All the switches provide software that helps to monitor the health of the switches. These tools can be used to diagnose the hardware failures on the switch. They can also be used to check the port statistic, performance check, port operational states, link incident alerts, threshold alerts, displaying node properties, and can also help in taking a report of statics. There is software applications like SANavigator which helps in managing the entire Fabric, and improve SAN efficiencies, improve security.

Configuration

Configuration of the SAN after implementation is a record that needs to be updated and recorded constantly. The configuration list should also consist of details like the HBA model, microcode level on the HBA, Fabric, port information of the switches, port information of the hosts, microcode information of the switches. These documents would help in faster resolution of problems due to unscheduled downtimes.

3.3 Zoning

The purpose of zoning is to make a barrier between two different groups, such as Windows NT servers and iSeries servers. Only members of the same zone can communicate within that zone. All other attempts from the outside are prevented. See the example in Figure 3-12.

Zoning can be implemented in the hardware or software and gives the user the flexibility to manage a Storage Area Network to meet different closed user group objectives. It enables resource partioning for management and access control.

> **Note:** To attach tapes or ESS volumes to the iSeries within a loop, you must create a zone for ESS volumes and for each tape.

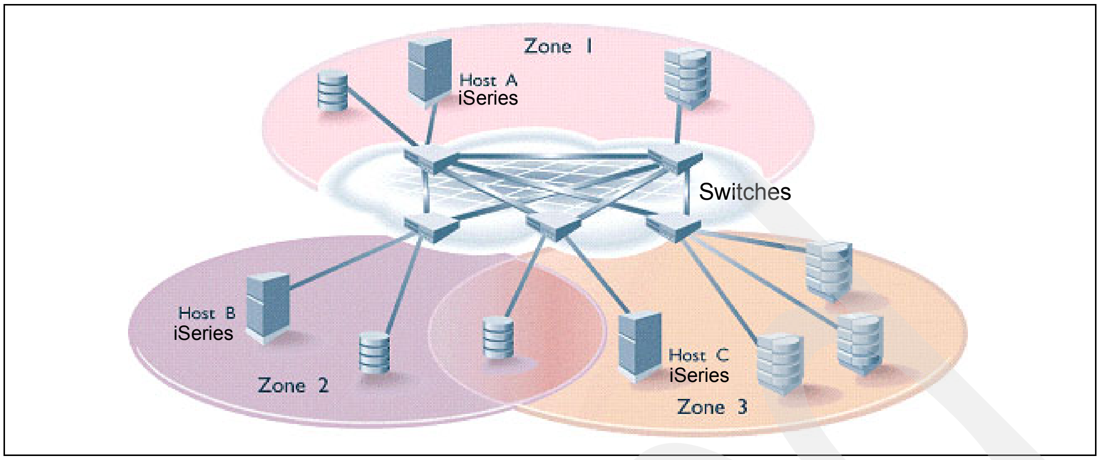

Figure 3-12 Zoning example

Zoning concepts

This section discusses zoning concepts and components:

- **Zone** (definition): A zone is a set of devices that access one another. All devices that are connected to fabric can be configured into one or more zones. Devices in the same zone can see each other. Every zone name begins with a letter followed by any number of letters, digits, and the underscore character. Every zone has a member list, consisting of one or more members. You can see an example of zoning in Figure 3-14.
- **Zone member**: Zone members are devices that are specified by:
  - Physical fabric port number
  - Node world wide name (not recommended)
  - Port world wide name

  The type of zone members defined in a zone can be mixed.

  Note not every switch supports all three member types. Refer to each switch user guide for more information.
- **Physical fabric port number**: A pair of decimal numbers where the first number is the switch domain number and the second is the port number on the switch. An example is 1, 4, which specifies port 4 on switch domain id 1. When a zone member is specified this way, any and all devices connected to that port are in the zone. If the port is an arbitrated loop, all loop devices are in the zone.
- **World wide name (WWN) notation (node and port)**: Specified as an 8 hex number. An example is `10:00:00:60:50:40:30:2b`. Zoning has no field knowledge within a world wide name. The 8 bytes are compared with the node and port names presented by a device in a login frame.
- Zone alias: Simplify the use of port entries numbers and world wide names. A zone alias is a character style name for one or more port numbers or world wide names.
- **Zone set configuration**: A set of zones. Zoning can be disabled, or one zone set configuration may be in effect. Multiple zone set configurations can be made, but only one zone set configuration is in effect at a time.
- **Hard Zones**: In a hard zone, all zone members are specified as switch ports; any number of ports in a fabric can be configured to the zone. When a zone member is specified by port number, only the individual device port specified is included in zone. Hard zoning provides the greatest security possible. Hard zoning is generally used where security must be rigidly enforced.

- **Soft Zones**: In a soft zone, at least one member is specified by WWN. A device is included in a zone if either the node WWN or port WWN specified matches an entry in the name server table. When a device logs in, it queries the name server for devices within the fabric. If zoning is in effect, only the devices in the same zones are returned. Other devices are hidden from the name server query reply. Soft zoning is generally used where flexibility is important and security can be ensured by the co-operating hosts.

> **Note:** For maximum security, we recommend *hardware zoning*. But since the standards are evolving and the industry is following them, it is likely that in the future, *software zoning* will be the preferred solution.

3.3.1 IBM 2109 Switches

IBM 2109 switches are auto-sensing Fibre Channel switches. They provide high reliability, availability, and port density. The 2109 supports the following ports, E_ports, FL_ports and F_ports, and NL ports via QuickLoop. The 2109 can operate at both 1Gbps and 2Gbps. The switch with its various components are shown in Figure 3-13.

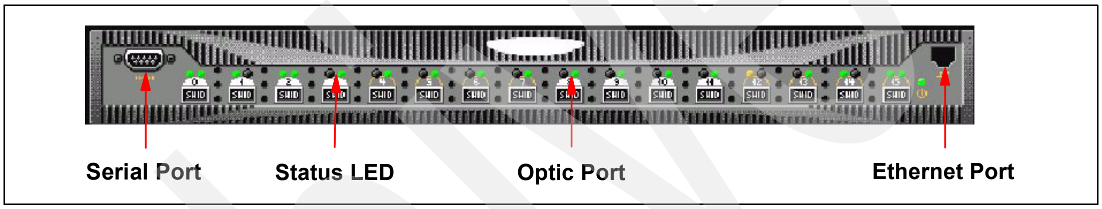

Figure 3-13 2109 switch

The 2109 has LEDs as status indicators, each port has two status indicators. The first LED for the port is a two-color (green and yellow) LED and indicates the status for the port. Green indicates normal status, and yellow indicates an error. The second LED is a single-color (green) LED and indicates the link speed for the port. Green indicates 2Gbps. If the LED is not lit (dark), it indicates 1Gbps. The are also LEDs above the power supply indicating the status of the power supply (green) indicates a normal and (yellow) indicates failure.

There are three different port types: Serial port, optic port and the Ethernet port

The serial port is used to set the internet protocol (IP) address when setting up a switch, reinitialze a switch, or running diagnostics. Serial port is not used during normal operation.

Optical ports are used to connect the fabric to the Host Bus Adapters (HBA) and devices.

The Ethernet port allows for Telnet and Web access for remote monitoring and testing. This connection uses no fabric connection.We can also use StorWatch Fibre Channel Switch Specialist once the Ethernet port is set up.

The StorWatch Fibre Channel Switch Specialist is used to Monitor and manage the 2109 switch. The initial screen of the StorWatch Fibre Channel Switch Specialist is shown in Figure 3-14.

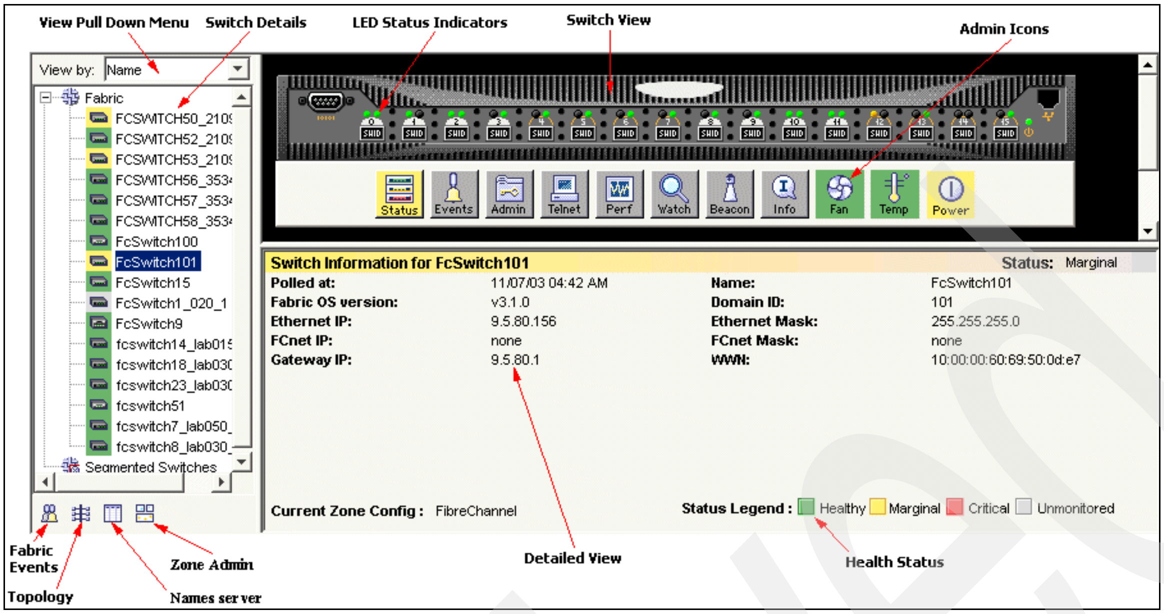

Figure 3-14 StorWatch View

Monitoring 2109 Switches

Fabric Manager includes features such as call home, security management, ISL checking, Fabric Checking, Fabric Merge check, comparing configuration, topology and fabric view, Fabric backup and troubleshooting:

- Call home- Continuously monitors the status of switches and sends a “call home” E-mail message to user-defined E-mail addresses when a triggering condition occurs. Some of these conditions are when the switch status changes, switch reboot, switch unreachable.
- Security Management -Fabric Manager provides a GUI interface to manage security once you enable security through the command line. Security policies need to be configured using the Fabric Manager Security Management. Details on security management can be found in “Fabric Manager application and user’s guide” at:

  `http://www-1.ibm.com/servers/storage/support/san/`
- ISL Checking- Helps to monitor any changes to ISL topology
- Fabric Checking- Allows to monitor fabric and registers events when switches are added or removed for the fabric.
- Fabric Merge Check- Checks if two fabrics will be allowed to merge.Compares various configuration elements of two fabrics before connecting.
- Comparing configuration- Helps validate and ensure consistent configuration settings among the switches, propagate configuration settings to switches, so that a selection of configuration settings that can be propagated
- Topology view- Provides complete graphical representation and detailed information of fabric
- Fabric backup- Creates a backup file of the entire fabric, includes the configuration file of every switch in the fabric, license key information, list of switches, ISL stamp, zone definitions, firmware versions and name server information
- Troubleshooting- Helps create a application log for the purpose of troubleshooting.

Zoning using StorWatch Fibre Channel Switch Specialist

There are multiple methods to implement zoning:

- With Telnet, using either out-of-band or in-band communication, by logging into the IBM SAN Fibre Channel Switch
- With the StorWatch Fibre Channel Switch Specialist

We discuss the most commonly used method, which is StorWatch Fibre Channel Switch Specialist.

This section explain how to implement zoning using the StorWatch Fibre Channel Switch Specialist:

1. Start a Java™-enabled Web browser and enter a switch name or IP address in the Location/Address field.
2. You see a fabric view similar to the example in Figure 3-15 for the managed hub and the example in Figure 3-15 on page 50 for Switch 2109. From the fabric view of Switch 2109, click the **Zone Admin** icon at the bottom left-hand side of the screen.

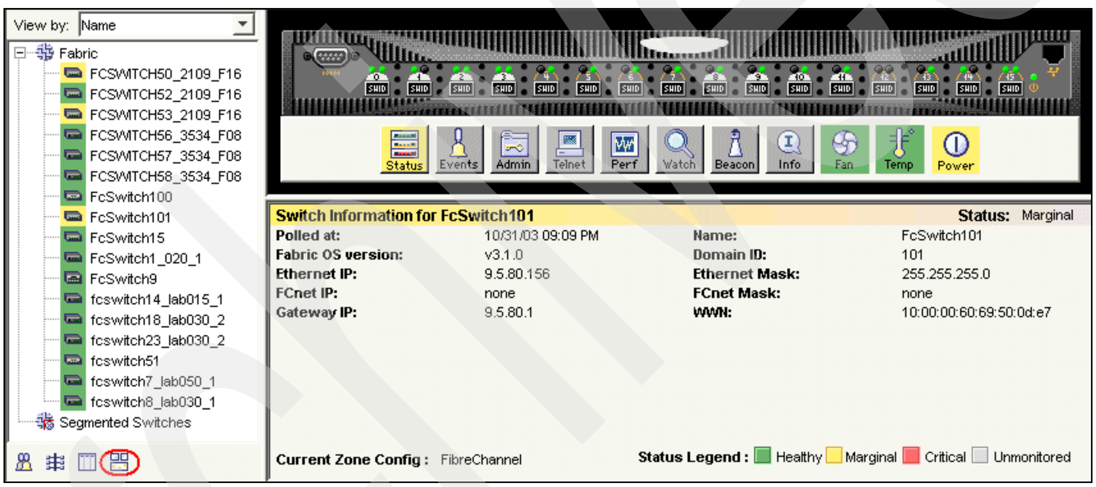

Figure 3-15 Fabric View

3. When we select **Zone Admin**, a prompt displays requesting User Name and Password, as shown in Figure 3-16.Enter the user name and password and select **OK**.

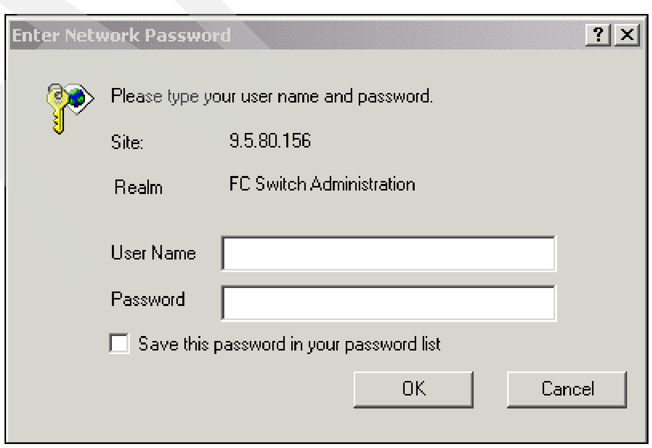

Figure 3-16 Network password entry scheme

4. The zone menu is shown in Figure 3-17.

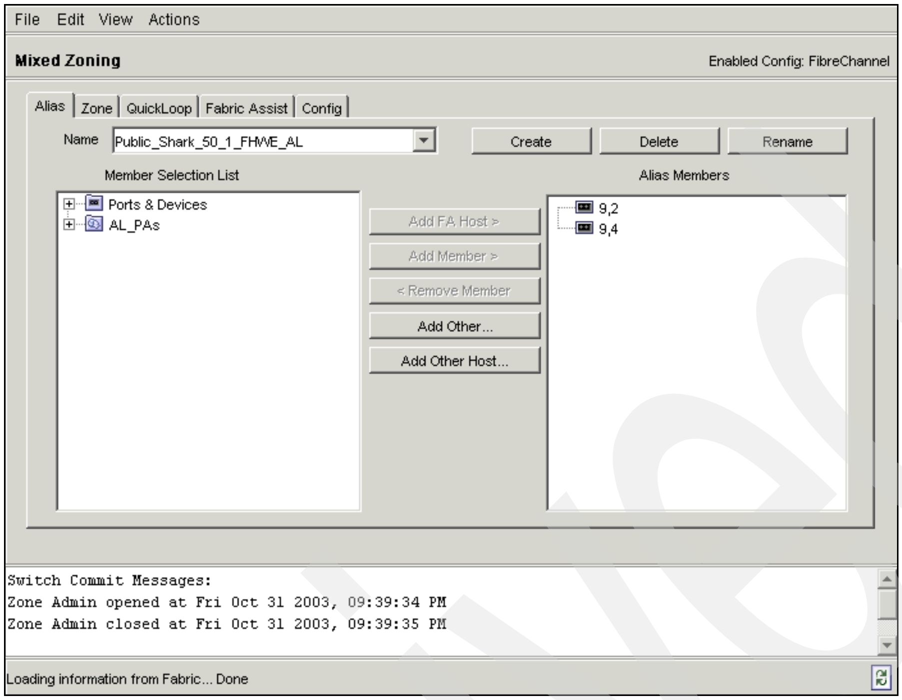

Figure 3-17 Zone admin initial view

5. Use the **Zone** tab to create, manage zones. A zone can have one or multiple members, and can include switches, ports, WWNs, aliases and QuickLoop. Use the Create tab to create a new zone as shown in Figure 3-18.

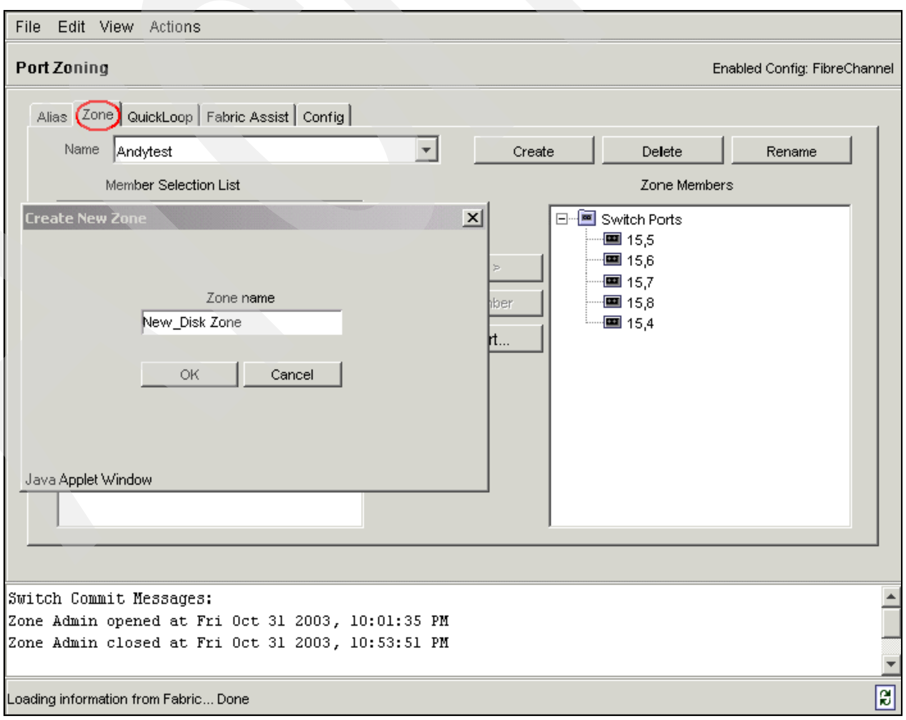

Figure 3-18 Create New Zone

6. We add the ports to the zone by selecting the port and clicking on the **Add Member** Tab as show in Figure 3-19. We have two iSeries hosts connected to port 2 and port 3 and ESS connected to port 0 and port 1.

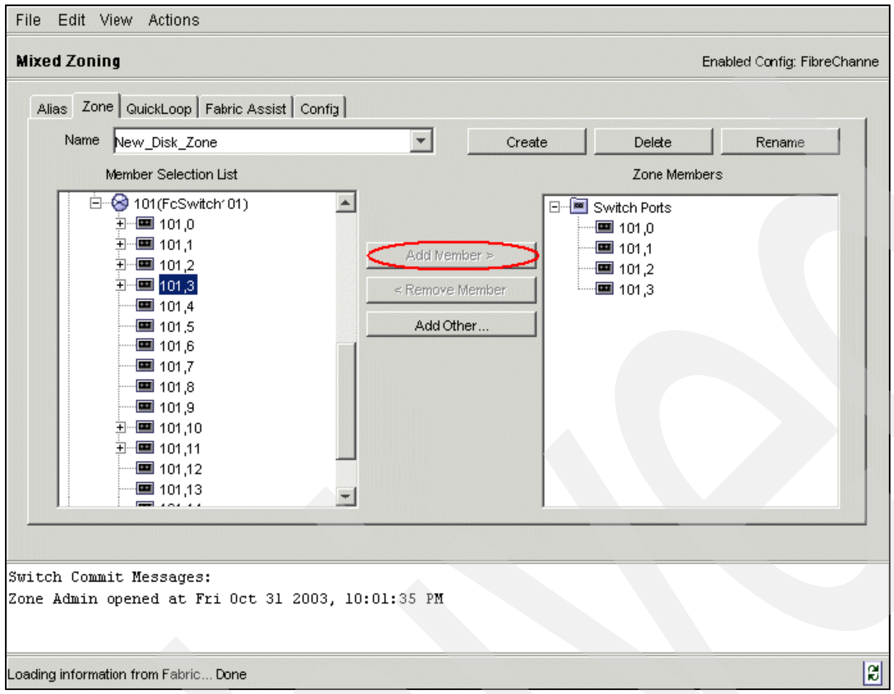

Figure 3-19 Zone for Disks

7. Tape and disks should not be in the same zone. Create another zone to add the tape devices that are connected to port 10 and port 11. This is shown in Figure 3-20.

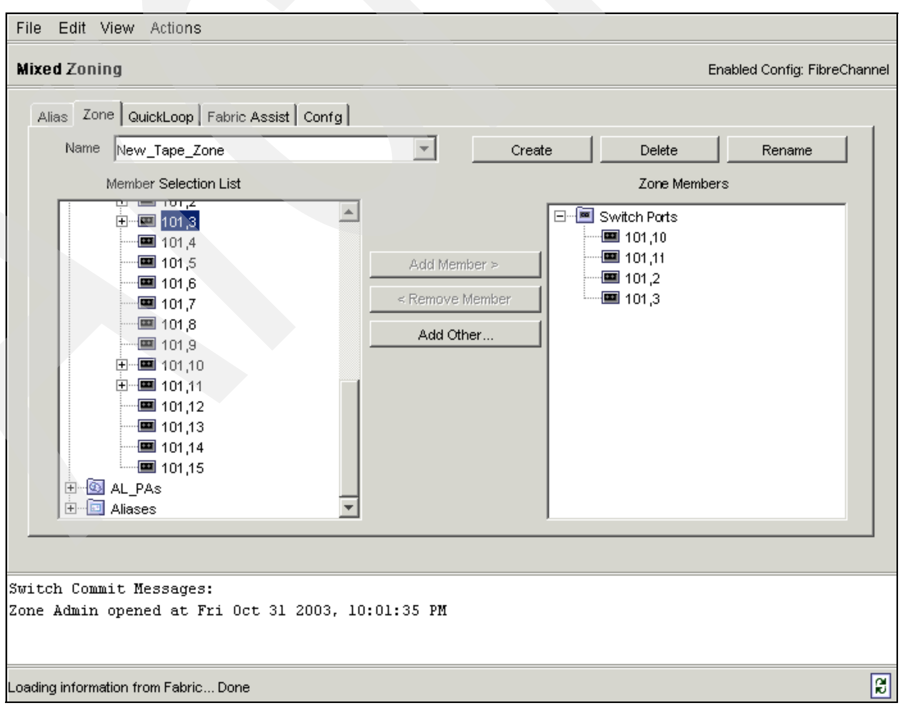

Figure 3-20 Zone for Tape devices

8. Use the **Config** tab to go to the config menu as shown in Figure 3-21. Create a config to add the disk and the tape zones into the config by using the **Create** tab. Once the config is created, add the members to the config by using the **Add Member** button.

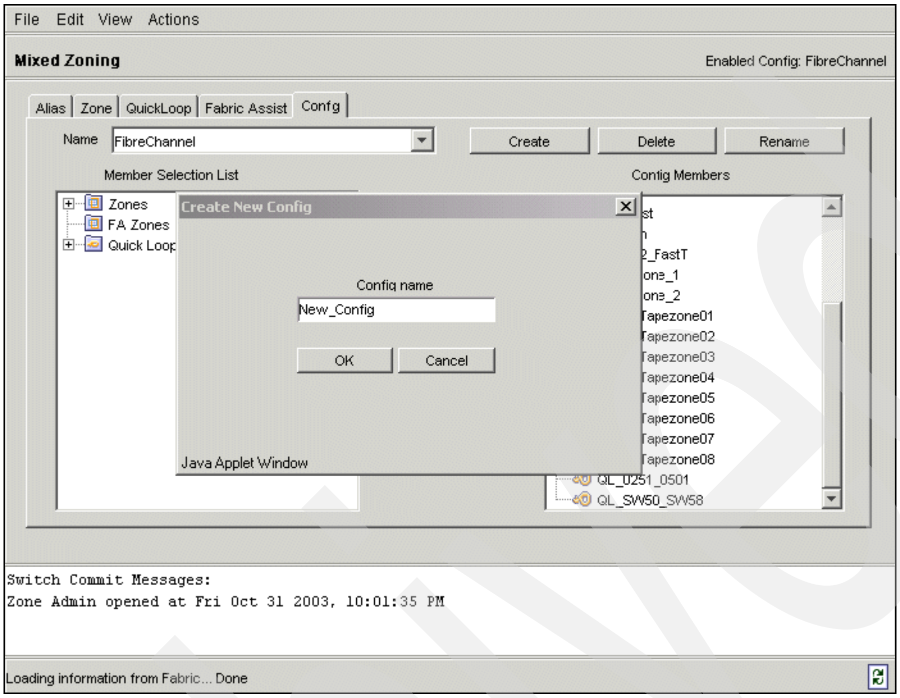

Figure 3-21 Create Config

9. Click the **Actions** tab and click **Enable config** to enable the configuration as shown in Figure 4-20.

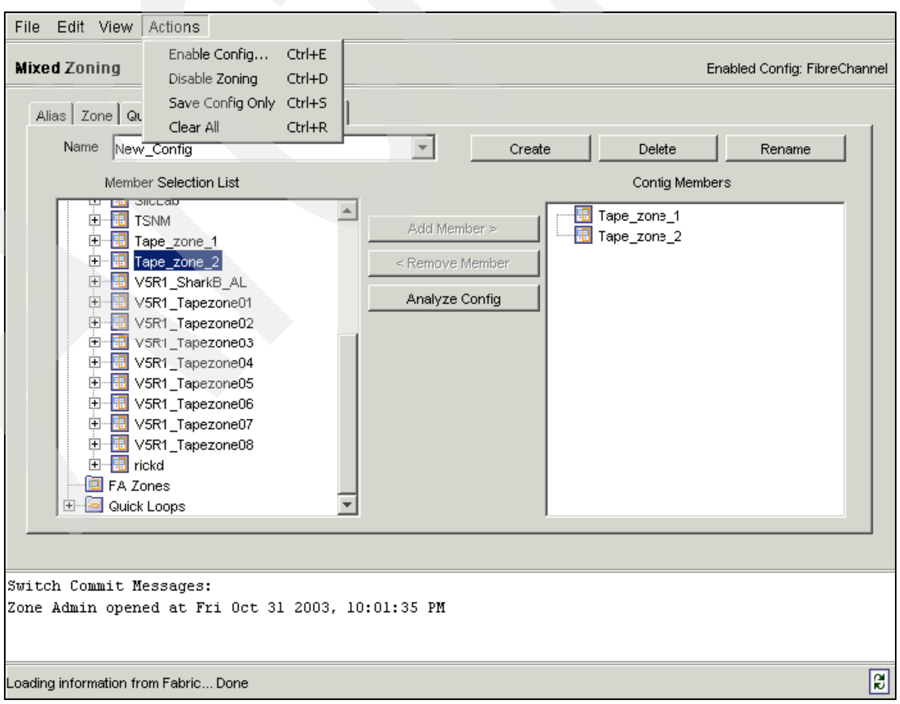

3.3.2 McData Switches

McData Switches are built on Flexport Technology, which allows you to purchase the port capacity based on requirement. Most of the switches allow for a Non-distriputive microcode load and hot plug features and connect at the speed of 2Gbps.

McData Switches are supported on iSeries from OS/400 version V5R2 or later. At the time of writing the set of switches supported by iSeries are 2031-232 (McDATA Sphereon 3232 Fabric Switch), 2032-140 (McDATA Intrepid 6140 Director), 2032-064 (McDATA Intrepid 6064 Director) and 2031-224 (McDATA Sphereon 4500 Fabric Switch). The supported zoning configuration for the above switches are shown in Table 3-1.

Table 3-1 Supported Zoning Config

Product	2032-140	2031-224	2031-232 2032-064
Number of End Ports	1024	1024	1024
Unique Zone Members	1024	1024	1024
Members per Zone	1024	1024	1024
Zones	1024	1024	1024
Maximum Devices Supported	1024	1024	1024
Maximum Number of a Switch Type in a Fabric	24	24	24
Maximum Number of ISLs per Switch	70	24	32

For the latest supported switches on iSeries check interoperability matrix for McDATA switches, see:

`http://www.storage.ibm.com/ibmsan/products/2032/index.html`

For latest zoning limits, refer to Supported Fabrics Configuration Document located at:

`http://www.mcdata.com/resources/tdoc/index.html`

Note: McData switches are supported on iSeries starting at V5R2

Documentation on installation of EFC Server is available at:

`http://www.mcdata.com/resources/tdoc/archive.html`

McDATA Enterprise Fabric Connectivity Manager

The Enterprise Fabric Connectivity Manager (EFCM) is a comprehensive storage network management application used to configure and manage McDATA Storage Area Networks. EFCM help in SAN management, optimizing storage resources and minimizing storage networking risks.

- **Topology/Physical Map:** Displays the SAN topology, including discovered and monitored devices and connections.
- **Master Log:** Displays all events that have occurred on the SAN.
- **Menu:** consists of commands you can perform on the SAN.

- **Minimap:** Displays a smaller view of the entire SAN.
- **Status Bar**: Displays data regarding the Server, connection, device, and fabric.
- **Toolbar**: Provides buttons that enable quick access to dialog boxes and functions.
- **Toolbox:** Provides tools for viewing the Physical map.
- **View tab**: Displays the Master log, mimimap, topology and product list.
- **Product List**: Displays an inventory of all discovered devices and ports.

Monitoring McDATA Switches

EFCM helps streamline troubleshooting processes. EFCM provides proactive alerts with call-home and E-mail notification capabilities. with detailed, real-time logging, and diagnostic and fault isolation data.

The events that take place in the SAN can be viewed through the Master Log, a specific log can be viewed using the option from **Monitor** menu's Logs submenu. The available logs include

- Audit Log- Displays a history of user actions performed through the application
- Event Log- Displays errors related to SNMP traps and Client-Server communications.
- Session Log- Displays the users who have logged in and out of the server.
- Product State Log- Displays operational status changes of managed products.
- Fabric log- Displays the events related to the selected fabric. The events include logged include, ISL removed, switch added, switch removed, fabric renamed, fabric persisted, fabric status changed and device removed

The other functions include:

- E-mail notification
- Call home notification
- Ethernet events monitor
- Report generation

Zoning using McDATA Enterprise Fabric Connectivity Manager

This section explains how to implement zoning using the McDATA Enterprise Fabric Connectivity Manager.

1. Start the EFCM and enter the Server IP address, user name and password in the fields provided.
2. Select the switch to be zoned on the right panel of the page and click the **Configure** menu and click the **Zoning** tab.
3. A new zone can be created using the **NewZone** button and renamed. The following rules need to be followed for naming:
   - Names are case insensitive
   - Names cannot begin with "SANav_" or "SMP_". This prefix is reserved
   - 60 character limit
   - No duplicate names are allowed
4. Select the **Zoning method** drop-down menu to select zoning method. There are two methods:
   - World Wide Name (WWN) method
   - Domain port method

   We use the WWN method for zoning in the below example.

5. From the **potential zone members** column, select the members and add by clicking the right arrow.
6. Create a new zone set by clicking the **zone set** tab under **zone sets** menu and rename the zone set as per the rules given on page 55.
7. Add the zone created previously to the new zone set.
8. To activate a zone set, select the zone set and click the **Activate** tab. In the activate Zone set menu, select the zone set created and click **ok**.

3.3.3 CNT (Inrange) Switches

CNT (Inrange) switches use a server-client based method to configure and manage the switch. The IN-VSM Enterprise client is used to configure zoning. There are multiple methods of zoning in Inrange switches.

- Port WWN
- Domain ID: Port number
- FC physical Address
- Mixed

Refer to the CNT Web site for further information about Inrange switches:

http://www.cnt.com

Inrange switches operate at 2Gbps and have the 128 port Fibre Channel switched fabric with N+1 redundancy for all active components. Inrange uses In-VSN Enterprise Manager software to centralize control for multiple Fibre Channel Directors across the enterprise.

Installing IN-VSM Enterprise Manager

The IN-VSM manager requires Java 2 Runtime Environment to be installed on the PC. The Java 2 Runtime is available on the installation CD with the file name j2re-1_4_1_01-windows-i586.exe. The latest version details of Java 2 Runtime environments is available at:

http://www.java.sun.com

Running **setup.exe** installs both server and client. Details about the configuration of the server are available in the documentation CD or at this Web site:

http://www.cnt.com

Configuration is also covered in the redbook Getting started with INRANGE FC/9000 FICON® Director, SG24-6858, which can be found at:

http://www.redbooks.ibm.com/redbooks/pdfs/sg246858.pdf

1. From the installation media, select **setup.exe**.
2. The welcome menu is displayed, click **Next**.
3. Read the license agreement and click **I agree** and click **Next**.
4. Check the **EM server** and **EM client** box and click **Next**.
5. Select the folder of installation by clicking **Browse**, and then click **Next** to start the installation.
6. If the directory is created new, we get the create directory confirmation screen. Click **Yes** to proceed.
7. We track the progress of the installation in the EM.
8. When the installation is complete, click **Finish**.

IN-VSM Enterprise Manager

The IN-VSM Enterprise Manager helps in centralized monitoring and control of multiple fabrics and all vital network functions from a single console. In-VSM Enterprise Manager help in configuration management, port zoning, performance and availability management, port statistics, and event logging.

- Menu Bar: Consists of five selections File, View Traps, Director and Help. Each of these menu selections represents a "function" of the Enterprise Manager software and displays a menu of related option when selected.
- Tool Bar: The tool bar is a shortcut for navigating through the menu. There are seven buttons in normal mode and eight if in maintenance mode.
- Navigation Tree: Is used to display a screen specific to the Fabric, switch, boards, or ports
- Details Panel: Information specific to the component you choose in the Navigation tree is displayed here, You may also effect changes upon certain components, view and acknowledge traps and perform other tasks related to switch and software.
- Message Panel: A scrolling display of the 25 most recent events of severity 3 or less
- Status Line: Includes status of the switch, error log count and login information.

Monitoring INRANGE Switches

In-VSM Enterprise Manager helps to monitor and manage the fabric with functions like auto-sense arbitrated loop, trap alarms, event logs, utilization monitor, port fail indicators, fabric security, save config and backup and restore of database

- Auto sense Arbitrated loop: There are three levels of enabling auto sensing arbitrated loop.
  - Director Level
  - Switch Level
  - Port Level
- Trap alarm: Include pot and fabric level.
- Event log: Shows information for everything controlled including fabric switch and ports at the bottom right of EM window.
- Utilization monitor: Monitors throughput in Mbps.
- Port Fail Indicators: Red indicator around the port indicates a failed port.
- Fabric security: Helps to enable or disable the ability to control the ISL of directors within a fabric.
- Save config: Gives the Administrator the ability to save the port configuration information
- Backup and restore Database: Backup and restore of the Database can be done using EM. Restoring a database should not be performed without consulting IBM CE.

Zoning using IN-VSM Enterprise Manager

This section explains how to implement zoning using the IN-VSM Enterprise Manager:

1. Login to the IN-VSM Manager using the IN-VSM Manager client.
2. Select the **Zoning** tab.
3. Select the Fabric in which zoning needs to be done and select **Zoom in** tab.
4. In the Domain View select the domain and select **Zoom in** tab.
5. Create a New Zone by entering the name of the new zone under the **All zones** tab.
6. Add the members to the new zone by clicking the **Zoom in** tab. Once the members are added to the zone, click the **Save** tab to save the changes.

7. Click **Zoom out** tab to return to zoning menu and create a zone set up entering the name of the new zone set.
8. Zones can be added to the new zone set by right-clicking on the new zone set, clicking **Add Zones**, selecting the required zones, and clicking **Add**. Click **Save** to save the changes.
9. The new zone set can be activated by right-clicking on the zone set and clicking **Activitate**.
10. Check for the active zone.

3.3.4 Cisco Switches

The Cisco MDS9000 Family of switches can be accessed and configured in many different ways, and support standard management protocols. There are different protocols that are supported in order to access, monitor, and configure the Cisco MDS9000 Family of switches. Some of them are:

- Telnet/SSH
- FTP/SFTP/TFTP
- SNMPV1,V2c,V3
- HTTP
- ANSI T11 FC-GS3

Cisco Fabric Manager is commonly used to discover and view fabric topology and to manage zones and zones sets. Cisco Device Manager is a utility used to monitor individual switches, this also gives a detailed information for verifying or troubleshooting device-specific configuration that is available in the Fabric Manager.

Note: Zones and zone sets can only be created and configured in the Fabric Manager.

Installing Cisco Fabric Manager and Device Manager

To install the software for the first time, we need to access the supervisor module with a Web browser. Fabric Manager requires Java Virtual Machine on the PC.

1. Open browser and enter the host name or IP address of the supervisor nodule in the address or location field of the browser.
2. The menu gives two options:
   – Install/Run Fabric Manager
   – Install/Run Device Manager
3. Click on the icon to start the installation.
4. Login screen indicates the end of installation. Entering the Server host name (Device name), user name and password enters the Fabric Manager or Device Manager.

Note: Cisco Fabric Manager does not allow multiple sessions to login simultaneously.

Cisco Fabric Manager

The Fabric Manager displays a view of the network fabric, including third party switches and end devices.

The main window has a menu bar, tool bar, message bar and status bar.

- VSAN/Switch pane: Displays a tree of configured VSANs and zones on the VSANs/Zones tab and a menu tree of available configuration tasks on the **Switch** tab.

- Information pane: Displays information about whatever option is selected in the menu tree.
- Map pane: Displays a map of the network fabric, including switches, hosts and storage. It also provides tabs for displaying log and event data.
- Menu bar: Provides options for managing and troubleshooting the current fabric and for controlling the display of information on the map pane. The menu bar provides the following menus:
  - File: Open a new fabric, rediscover the current fabric, locate switches, set preferences, print the map, and clear or export the map pane log.
  - Edit: Manage zones, zonesets, and various elements on the Fabric Manager map.
  - View: Change the appearance of the map (these options are duplicated on the Map pane tool bar).
  - Reports: Display summary reports.
  - Troubleshooting: Verify and troubleshoot connectivity and configuration.
  - Help: Display online help topics for specific dialog boxes in the information pane.
- Toolbar: Provides buttons for accessing the most commonly used menu bar options. The map pane tool bar provides buttons for managing the appearance of the map. The information pane tool bar provides buttons for editing and managing the Information pane.
- Status bar: Show the short term, transient messages (such as the number of rows displayed in the table), and long term discovery issues.

Device Manager

Device Manager provides a physical representation of your switch chassis, with the modules, ports, power supplies and fan assemblies. Most tasks performed with Device Manager can also be performed for multiple switches using the Fabric Manager. However, Device Manager provides more detailed information for verifying or troubleshooting device-specific configuration than what is available from the Fabric Manager.

- **Menu bar**: Provides access to options, organized into menus that correspond to the menu tree in Fabric Manager.
- **Tool bar**: The tool bar includes icons for opening devices, refresh display, command line interface, configure selected switch, SNMP Event log, Threshold manager, VSANs, SNMP Users and roles, save configuration, copy configuration file between server and switch, reset switch and online help.
- **Legend**: Indicates port Status
  - Green: The port is up.
  - Brown: The port is administratively down.
  - Red: The port is down or has failed.
  - Gray: The port is unreachable

Monitoring Cisco Switches

Fabric Manager and Device Manager can monitor events and reports errors. Fabric Manager provides the following features and responses to network events:

- SNMP events: These are pre configured notifications, including SNMPv2 traps and SNMPv3 informs.
- RMON alarms: These are configurable notifications that you can set based on thresholds for various network events.
- Call Home: This is a feature that lets you configure automatically generated e-mail messages or other responses to specific events.

- Syslog: This is a standard message log that records various network and system events.

Note: The Fabric Manager allows to manage events on multiple switches. The Device Manager allows to manage events on a single switch.

Zoning using Cisco Fabric Manager

Zones are configured within VSANs (Virtual SANs) but zones can be configured without configuring VSANs by configuring them within the default VSAN.

Note: VSANs allow you to separate devices that are physically connected to the same fabric, and thus provide higher security and scalability in the network fabric.VSAN could be described as multiple logical SANs over a common physical infrastructure.

The steps involved in creating a zone and zone set are explain below:

1. From the Fabric Manager select the **Zone** menu and click **Edit Full Database on Switch**.
2. The VSAN dialog box is display, select the VSAN and click **OK**.
3. To add a zone right-click and select **Insert**, add the name of the zone and click **OK**. Zone can be made read only by selecting the **read only** tab below the zone name.
4. Members can be added to the zone by right-clicking the new zone and clicking **Inset**.
5. In the Add Members menu select **Switch ports** tab, select the **Switch** to be zoned and enter the **port** details. The details of the ports can be collected from Device Manager, this is described in “Port Name Identification in Cisco Device Manager” on page 60.
6. To create a zone set right-click the zone set and click **Inset**.
7. Enter the name of the new zone and click **OK**.
8. To add zones to the new zone set, right-click the new zone set and click **Inset**.
9. Select the zones to be added and click **Add**.
10. To Activate the zone set right-click the zone set and click **Activate**. Active zone set is shown in italic type. This configuration is distributed to other switches in the network fabric.

Note: Changes made using Fabric Manager are applied to the running configuration of the switches you are managing and the changes may not be saved when the switch restarts. After you make a change to the configuration or perform an operation (such as activating zones), the system prompts you to save your changes before you exit.

Port Name Identification in Cisco Device Manager

1. In the device manager click the **Interface** menu and click **All FC ports**.
2. In the All FC ports menu select the other tab, this gives the complete list of ports on the switch.

Part 3

External disk

Part 3 covers the attachment of external ESS disk to iSeries.

4 Planning for external disk

When you configure Enterprise Storage Server (ESS) for iSeries it is a challenge to achieve the optimal balance between capacity, performance, and price.

It is important to remember that performances of disk system significantly influence performances of iSeries applications. We distinguish between two types of response time:

- Application response time - the response time of an application transaction. This is the response time which users experience and is important for an iSeries customer.
- Disk response time - the time in which an input/output (I/O) operation on disk is accomplished.

Although there are many factors in an iSeries that influence application response time (power of processor, size of memory, disk response time, etc.), disk response time is among the most important factors.

This can be observed on the graph in Figure 4-1. It is based on measurements that were performed in the Storage lab in Tucson with ESS model 800 and iSeries and shows how the response time of external disks influences the response time of iSeries applications. For a detailed description of the graph, refer to the presentation IBM TotalStorage Enterprise Storage Server (ESS) + iSeries performance on the following Web page:

http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geoID=All&prodID=Disk&docID=essiseriesperfpre

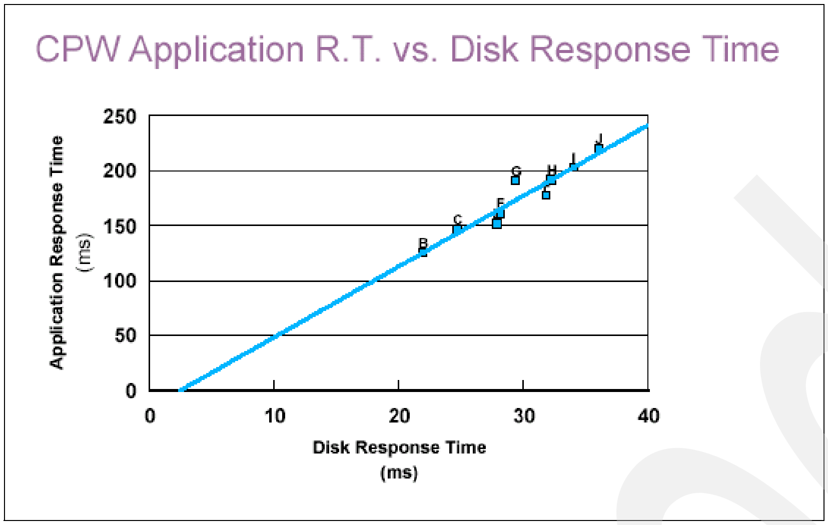

Figure 4-1 Application response time influenced by disk response time

To make a plan for ESS for iSeries as good as possible, it is necessary to consider all the components that are involved in I/O process: iSeries buses, Input Output Processors (IOPs), iSeries Fibre Channel Disk Controllers (FC adapters), SAN Fabric, ESS host adapters, clusters, device adapters, physical disks, logical disks, etc. Figure 4-2 basically shows the flow of I/O operations through these components.

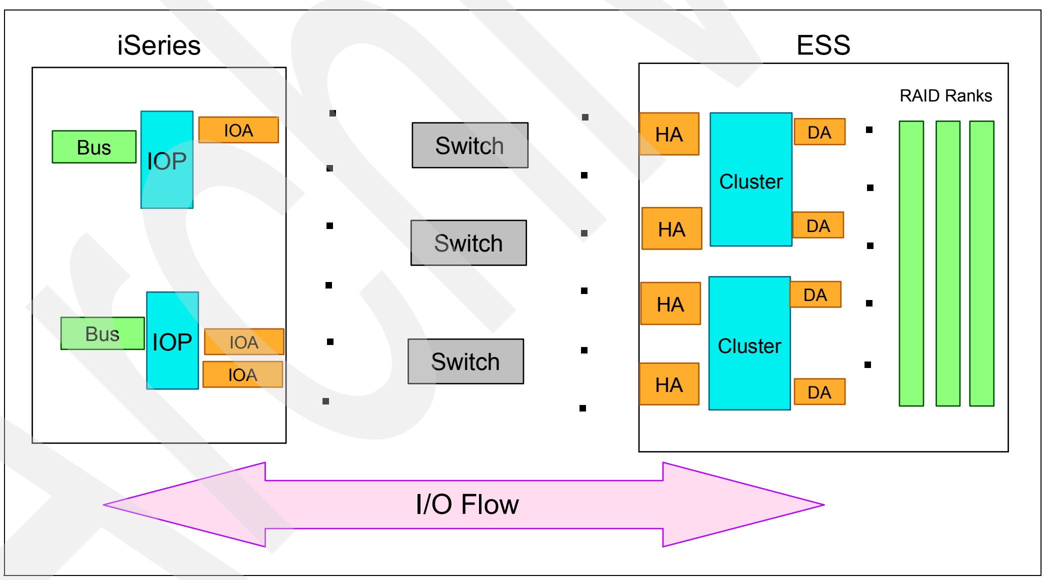

Figure 4-2 I/O flow

In this chapter we describe how these components relate to performances of ESS with iSeries and give basic guidelines for planning each of them.

Part of the chapter is dedicated to the tool Disk Magic with which we model ESS for iSeries based on the data from iSeries present workload. We can predict how iSeries disk response time will change if any of the ESS components change.

In this chapter you will find a checklist of activities, which we recommend that you perform during planning and implementation of ESS with iSeries.

4.1 ESS and iSeries hardware components

This section looks at the various components of the ESS and iSeries which are related to sizing and performance.

Figure 4-3shows the ESS components as well as a scheme of ESS. Along the top of Figure 4-3, you see up to 16 2Gb host adapters supporting one Fibre Channel port. Host adapters are placed in host adapter bays, 4 adapters in each bay. Each host adapter is connected to both clusters through the Common Parts Interconnect (CPI) buses, so that either cluster can handle I/O from any host adapter. You can see the two clusters that contain the RISC processors, the cache, and the Non-Volatile Storage (NVS) for the opposite cluster. Within each cluster, there are up to four device adapters. They are always installed in pairs, and the disk arrays are attached through an SSA loop to both device adapters in a pair. Each disk array (or rank) can be configured as a set of logical volumes which stripe across the disks in an array. They can be allocated (assigned) to Fibre Channel adapters in a host server and are seen by the host server as disk units.

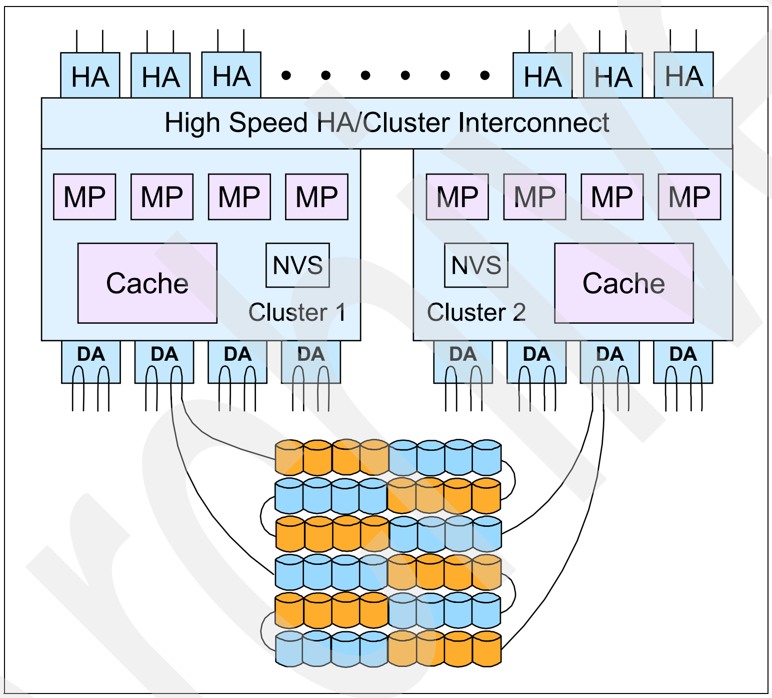

Figure 4-3 ESS hardware components

A part of iSeries architecture is hierarchy of microprocessors, which is shown in Figure 4-4. It uses input/output processors (IOPs) for handling I/O to storage devices and communications.

IOP is used to transfer data between main storage and a group of controllers and adapters (IOAs) which are attached to it. Workload for devices that are attached to these IOAs is off-loaded from the main processor to the IOP, so the main processor is available for additional workload.

IOA is attached to the IOP and transfers data between the IOP and attached devices. iSeries Fibre Channel Disk Controller is an IOA and is controlled by an IOP. Note: In this chapter we use the more popular term iSeries Fibre Channel adapter (FC adapter) instead of iSeries Fibre Channel Disk Controller.

IOPs and IOAs are accommodated on I/O cards that fit into slots on the system buses. The current iSeries product line uses Pheripheral Component Interconnect (PCI) or high-speed link (HSL) technology to attach IOPs.

Figure 4-4 also shows and example of iSeries Fibre Channel Disk Controller (FC adapter) with feature number 2787 attached to IOP with feature number 2844.

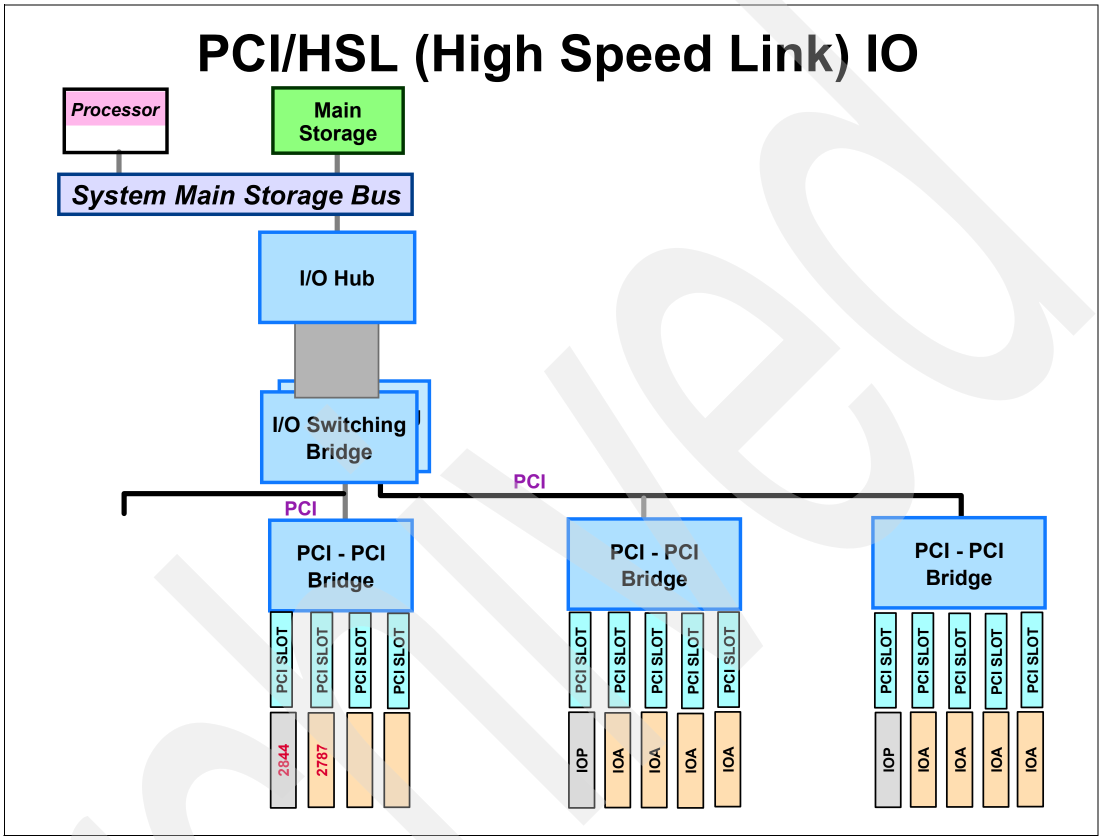

Figure 4-4 iSeries Microprocessors hierarchy

4.2 Consideration for planning ESS with iSeries

In this section we describe the factors that should be considered when planning ESS for iSeries. They are grouped in the following three areas which are the most important for performances of ESS connected to iSeries:

- Capacity (size) and speed of ESS disks, and the number of ESS ranks
- Number of iSeries Fibre Channel Adapters
- Capacity and allocation of logical volumes
- Considerations when consolidating storage from iSeries and other servers
- SAN fabric

Note: Guidelines and rules of thumb for sizing ESS with iSeries can be found in the section 4.3, "Checklist for planning ESS with iSeries" on page 72. Modeling ESS for iSeries with Disk Magic is described in 4.4, "Using Disk Magic for planning ESS with iSeries" on page 84.

4.2.1 iSeries workloads

When planning ESS for any server platform it is important to know the characteristics of the workload that will use ESS disks. This is also true for iSeries.

Many iSeries customers applications tend to follow the same patterns as iSeries benchmark workload Commercial Processing Workload (CPW). These applications have typically many jobs running brief transactions with database operations, similar to Online Transaction Processing (OLTP).

Other applications tend to follow the same patterns as iSeries benchmark workload Compute Intensive Workload (CIW): fewer jobs running transactions which spend a substantial amount of time in the application itself. An example of such workload is Domino® Mail and Calendar on iSeries**.**

In general, iSeries batch workloads can be I/O intensive or compute intensive. For I/O intensive batch applications the overall batch performance is very dependant on the speed of disk subsystem, for compute intensive batch workloads the run time is likely to depend on power of iSeries processor. Typically batch jobs would run during the night and for some enterprises it is very important that they finish in a certain amount of time which is called Batch window.

4.2.2 Number of ESS ranks, capacity and speed of ESS disks

The size and speed of ESS physical disks or Disk Drives Modules (DDMs) which are grouped in eight-packs (arrays or ranks*) can have significant effect on performance of ESS with iSeries.

Spreading the capacity across multiple ESS ranks ensures that as many disk arms as possible are available; this is especially important for mission critical applications. Therefore, when you plan the allocation of LUNs for iSeries, spread them across as many ranks as possible. It is not necessary that you fill up the rank with iSeries LUNs, important is that as many ranks as possible are included in iSeries disk space

On the other hand, you may plan to consolidate disk capacity from iSeries and other servers on the same ESS. For performance reason we recommend not to mix iSeries critical workload and workload from other systems within the same rank. So it wouldn't be wise to fill up the ranks with disk capacity for other workloads. In such case you may consider to increase disk capacity for iSeries rather to leave some disk space unallocated.

For example: say that you plan 700 GB disk capacity for iSeries, and you plan ESS with 36 GB DDMs. In order to use enough ranks for iSeries critical application you may want to spread the iSeries capacity over 4 ranks. This way you will allocate about 175 GB on each rank for iSeries, leaving about 45 GB unused disk space on each rank. If you plan to allocate disk capacity for other servers on the same ESS. If the workloads are not critical you would fill up the ranks with disk space for other servers. But with critical workloads it may not be a good idea to allocate the LUNs for other servers out of remaining 45 GB on the rank. In this case you may plan to increase the disk capacity for iSeries so that it fills all the space on 4 ranks.

4.2.3 iSeries Fibre Channel adapters

ESS can be attached to iSeries via two FC adapters with iSeries feature numbers 2766 and 2787. Each of them can be attached to one of the following IOPs: 2842, 2843, 2844.

The most accurate planning for the number of iSeries FC adapters can be done based on throughput in MB/second. You should take into account the following capabilities of FC adapters:

FC adapter 2766 is capable of transferring:

- 34 - 37 MB/s for transaction workload with high cache hit rations
- 57 MB/s for sequential read with 100% cache hits and large transfer size

FC adapter 2787 is capable of transferring the following I/O rate:

- 34 - 54 MB/s for transaction workload with high cache hit ratios
- 3.5 - 11 MB/s for transaction workload with low cache hit ratios
- 85 MB/s for sequential read with 100% cache hits and large transfer size

On the iSeries side:

- FC adapter 2787 controlled by IOP 2844 is capable of 2555 IO/s at 70% utilization
- FC adapter 2766 on IOP 2844 is capable of 1995 IO/s at 70% utilization
- FC adapter 2766 on IOP 2843 is capable of 1610 IO/s at 70% utilization

Take into consideration the following:

- I/O towers are used with iSeries to house components beyond the capacity of the system unit. As shown in Figure 4-5, I/O Tower, I/O tower houses, IOPs, IOAs, and internal disks; when iSeries uses external storage, the I/O tower houses IOPs and IOAs - FC adapters to connect to ESS. Multi-adapter bridges connect adapters to buses.

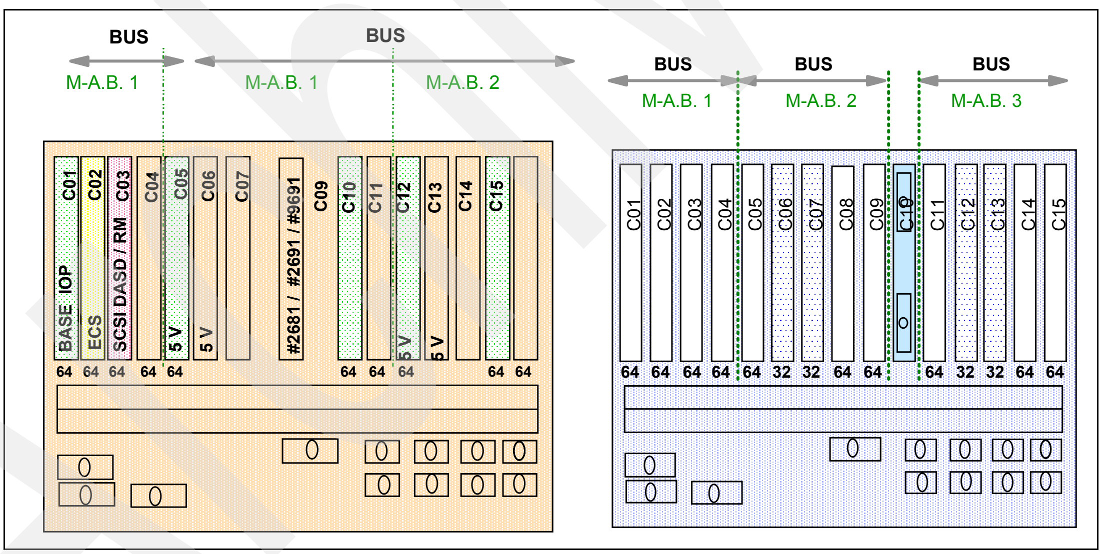

Figure 4-5 Multi-access bridge layout in I/O expansion units

- High Speed Link (HSL) bus technology was introduced to iSeries in year 2000. This flexible and powerful technology enables much better performance than previous iSeries buses. I/O towers and devices can be attached to iSeries on HSL loops, as shown in Figure 4-6 Multiple HSL Loops.

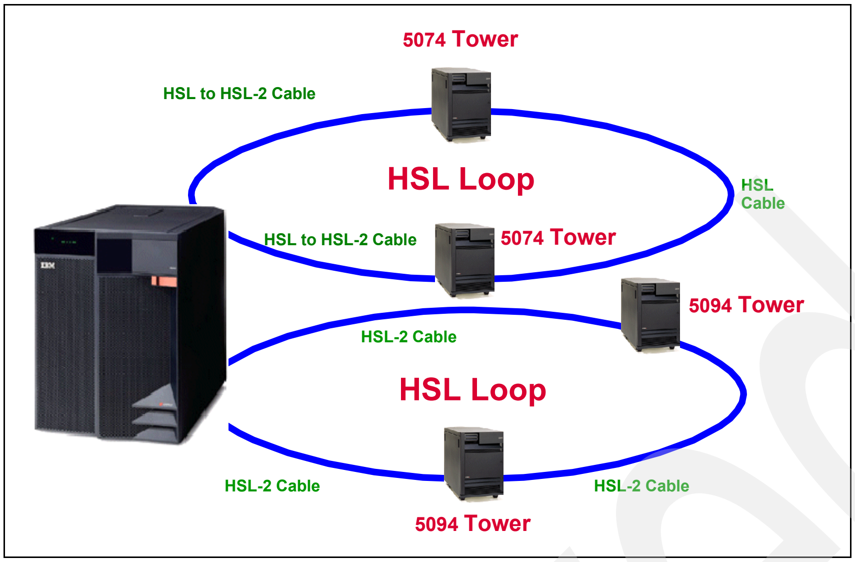

Figure 4-6 Multiple HSL loops

- Each FC adapter must be attached to an IOP which is dedicated for this FC adapter. iSeries configurator takes care of this.
- Consider one FC adapter-IOP pair per multi-adapter bridge.
- Consider no more than 6 FC adapter IOP pairs per one HSL loop.
- Spread FC adapter-IOP pairs as evenly as possible among I/O towers, and spread I/O towers as evenly as possible among HSL loops.
- When planning for throughput you should also remember that the performances of applications with large I/O activity are significantly better when FC adapters are attached to the IOP 2844, as shown in Figure 4-7.

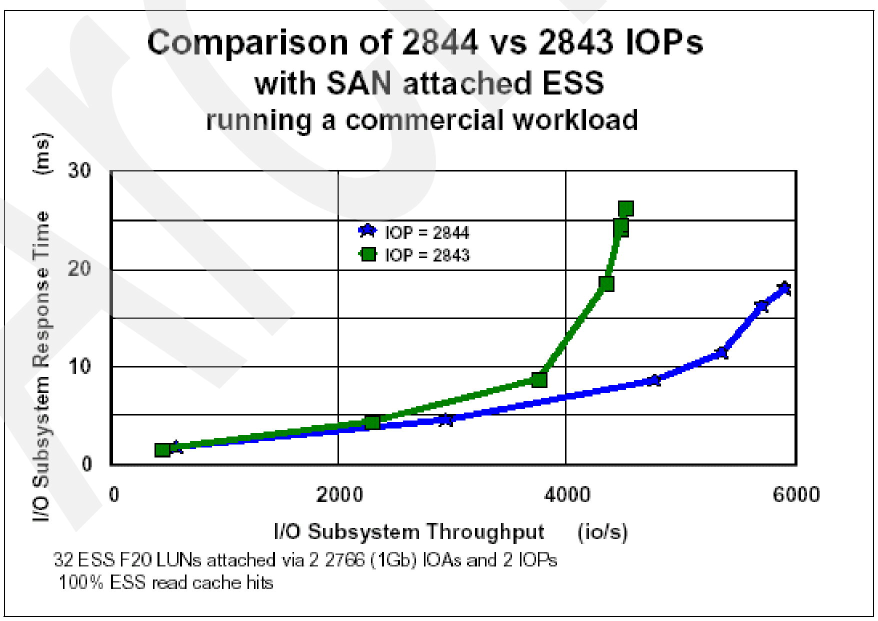

Figure 4-7 Influence of IOP throughput on I/O response time

4.2.4 Capacity and allocation of logical volumes

With iSeries internal disks one physical unit means also one logical unit. But with ESS the logical units are stripes of multiple DDMs from one array, as shown in Figure 4-8. We determine the size of a LUN when we define it in ESS. After the LUNs are defined in ESS we assign the to iSeries FC adapters and iSeries sees each LUN as a physical disk unit.

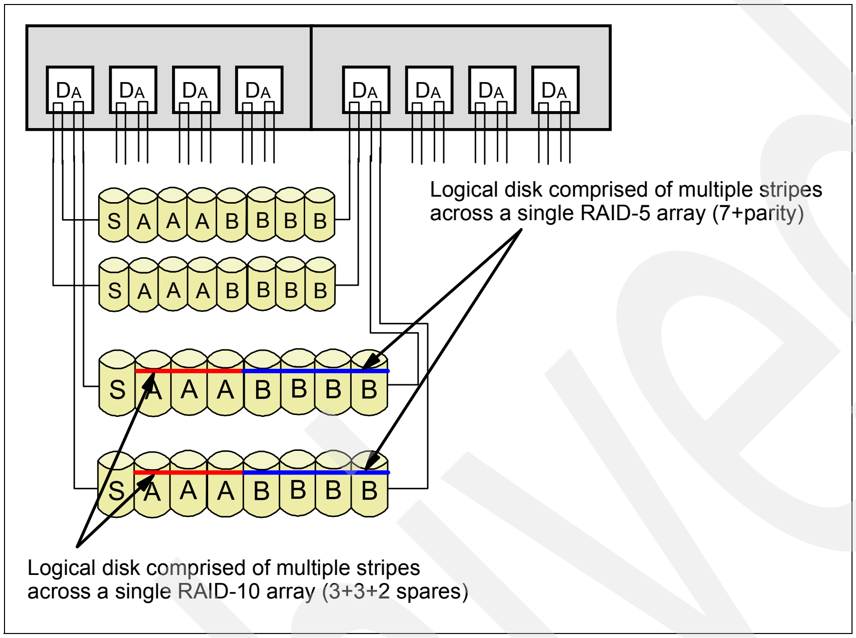

Figure 4-8 Logical units in ESS

From the ESS perspective, the size of a logical device has no effect on its performance. But, it is different from iSeries side: iSeries only allows one I/O operation per disk device at a time, which means that only one I/O operation per LUN can be done at a time. The smaller the LUNs are, more I/O operations at a time can be done, and consequently disk response time is better.

Also, observe the following difference between internal drives ad ESS LUNs: with internal drives the cache is on disk adapters. An I/O request is processed via cache. If it cannot be satisfied via cache, it is sent to drive and queued if the drive is busy. When ESS is attached to iSeries the I/O requests are queued before they reach the cache on ESS. So, it is important to keep I/O queues as short as possible, which is again accomplished by smaller sizes of LUNs.

When allocating the LUNs for iSeries consider the following guidelines for better performance:

- Balance activity to ESS clusters. When selecting RAID arrays for a critical application, spread them across separate clusters. Since each cluster has separate memory buses and cache memory, this maximizes the use of those resources.
- Balance activity to device adapters. When selecting RAID arrays within a cluster for a critical application, spread them across separate device adapters.
- Balance activity to adapter bays. When selecting host adapters to assign to iSeries, spread them across different adapter bays.

Following table shows size and number of LUNs which we can define from different ESS ranks. See Table 4-1.

Note: the residuals will be still adjusted based on the answer from Storage US ATS.

Table 4-1 LUN capacities

Array size	Effective capacity	2105-A01 2105-A81 8.59 GB	2105-A02 2105-A82 17.54 GB	2105-A05 2105-A85 35.16 GB	2105-A03 2105-A83 36.00 GB	2105-A04 2105-A84 70.56 GB
18.2 GB 6+P+S	105.20 GB	12 LUNs + 2.12 GB	5 LUNs + 17.50 GB	2 LUNs + 34.88 GB	2 LUNs + 33.20 GB	1 LUN + 34,64 GB
18.2 GB 7+P	122.74 GB	14 LUNs + 2.48 GB	6 LUNs + 17.50 GB	3 LUNs + 17.26 GB	3 LUNs + 14.74 GB	1 LUN + 52.18 GB
36.4 GB 6+P+S	210.45 GB	24 LUNs + 4.29 GB	11 LUNs + 17.51 GB	5 LUNs + 34.65 GB	5 LUNs + 30.45 GB	2 LUNs + 69.33 GB
36.4 GB 7+P	245.53 GB	28 LUNs + 5.01 GB	13 LUNs + 17.51 GB	6 LUNs + 34.57 GB	6 LUNs + 29.53 GB	3 LUNs + 33.85 GB
72.8 GB 6+P+S	420.92 GB	49 LUNs + 0.01 GB	23 LUNs + 17.50 GB	11 LUNs + 34.16 GB	11 LUNs + 24.92 GB	5 LUNs + 68.12 GB
72.8 GB 7+P	491.08 GB	57 LUNs + 1.45 GB	27 LUNs + 17.50 GB	13 LUNs + 34.00 GB	13 LUNs + 23.08 GB	6 LUNs + 67.72 GB
145.6 GB 6+P+S	841.84 GB	98 LUNs + 0.02 GB	47 LUNs + 17.46 GB	23 LUNs + 33.16 GB	23 LUNs + 13.84 GB	11 LUNs + 65.68 GB
145.6 GB 7+P	982.16 GB	114 LUNs + 2.9 GB	55 LUNs + 17.46 GB	27 LUNs + 32.84 GB	27 LUNs + 10.16 GB	13 LUNs + 64.88

4.2.5 Consolidating storage form iSeries and other servers

When using ESS, you can combine data and workloads from different kinds of independent servers into a single shared resource. This may be a mix of production and test servers in an open system environment or a mix of S/390 and open systems. In this kind of independent server environment, servers may rarely, if ever, contend for the same resource.

Sharing resources in ESS has advantages for storage administration and the servers that share it, but it has some implications for workload planning. The benefit of sharing is that a larger resource pool (for example, disk drives or cache) is available for critical applications. However, you should ensure that uncontrolled or unpredictable applications do not interfere with mission-critical work. This requires the same kind of workload planning you use when mixing various types of work on a server.

If your workload is truly mission critical, you may want to consider isolating it from other workloads. This is particularly true if other workloads are unimportant (from a performance perspective) or very unpredictable in their demands. You can achieve this by placing the iSeries LUNs on separate arrays. Note that S/390 and open system data are automatically placed on separate arrays. This reduces contention for disk use.

4.2.6 SAN fabric

In many cases you will consider to connect ESS to iSeries via SAN switches. For example:

- For high availability and disaster recovery you may connect an ESS locally and another ESS on long distance to iSeries, and use OS/400 mirroring between LUNs on local and remote ESS. In case of disaster a backup iSeries performs Load source recovery from remote ESS. To connect ESS on long distance you will use switches and Inter Switch Links (ISLs).
- If you plan to share one ESS host adapter between two iSeries FC adapters you will use a switch to connect both iSeries adapters to ESS.
- If you plan to share one ESS host adapter between iSeries and another server you will connect them both of them to ESS via a switch.

For accurate planning the connections, switches and ISLs consider the bandwidth of a switch, fibre cable and ISL, and the transfer (MB/s) of iSeries workload.

For the bandwidth of a particular switch and a Fibre cable refer to 4.3.4, “SAN fabric” on page 82.

Consider the following way to calculate the transfer of iSeries workload from iSeries performance report Resource Interval Report - Disk Utilization Summary: For each interval multiply the number of IO/s (Average I/O /second) with the transfer size (Average K per I/O). This way you will get the transfer KB/s for each interval. For this calculation you may import Performance report to Excel Spreadsheet and create a new column in which the transfer values are calculated. Choose the maximal transfer value and use it for planning SAN connections.

For example: You plan to connect three iSeries via a switch 2109-F16 to ESS. In order to make sure that one switch will have enough bandwidth for workloads for the three iSeries calculate the transfer form iSeries as described above. Calculations from performance reports show the following maximal transfer values:

- iSeries 1: 342 I/O per second multiplied by 100 KB per IO...34.2 MB/second
- iSeries 2: 740 I/O per second multiplied by 35 KB per I/O... 25.9 MB/second
- iSeries 3: 1850 I/O per second multiplied by 17 KB per I/O... 31.45 MB/second

The transfer from all three iSeries workloads will be 34.2 MB/s + 25.9 MB/s + 31.45 MB/s = 91.55 MB/s. Since the bandwidth between each two ports of the switch 2109-F16 is 200 MB/s you can be sure that the switch will not be a bottleneck for iSeries I/O operations.

4.3 Checklist for planning ESS with iSeries

In this section we present a checklist of actions that you should perform when planning ESS for iSeries. We took into account two initial situations:

- Performance Tools reports (PT reports) from current iSeries are available. For details about when and how to produce PT reports, refer to 4.4.1, “iSeries performance reports” on page 85.
- PT reports are not available, but you have information about the type of workload and possibly capacity.

If possible we certainly recommend that you get PT reports from current iSeries since they show exact statistics of I/O activity. But we are aware that in some cases you have to plan ESS for an iSeries for which you have very limited information. In such cases try to identify the type of present iSeries workload and place it in one of the following categories: Transaction application, Bath jobs, Compute intensive workload. It is possible that the workload contains a certain percentage of each type: for example 30% of Transaction application, 30% of Compute intensive and 40% of batch.

In this checklist we present some characteristics of workload for each of these three categories. We also give rules of thumb for initial sizing of ESS based on workload statistics, or characteristics.

For planning ESS follow the steps that are shown in the flowchart in Figure 4-9:

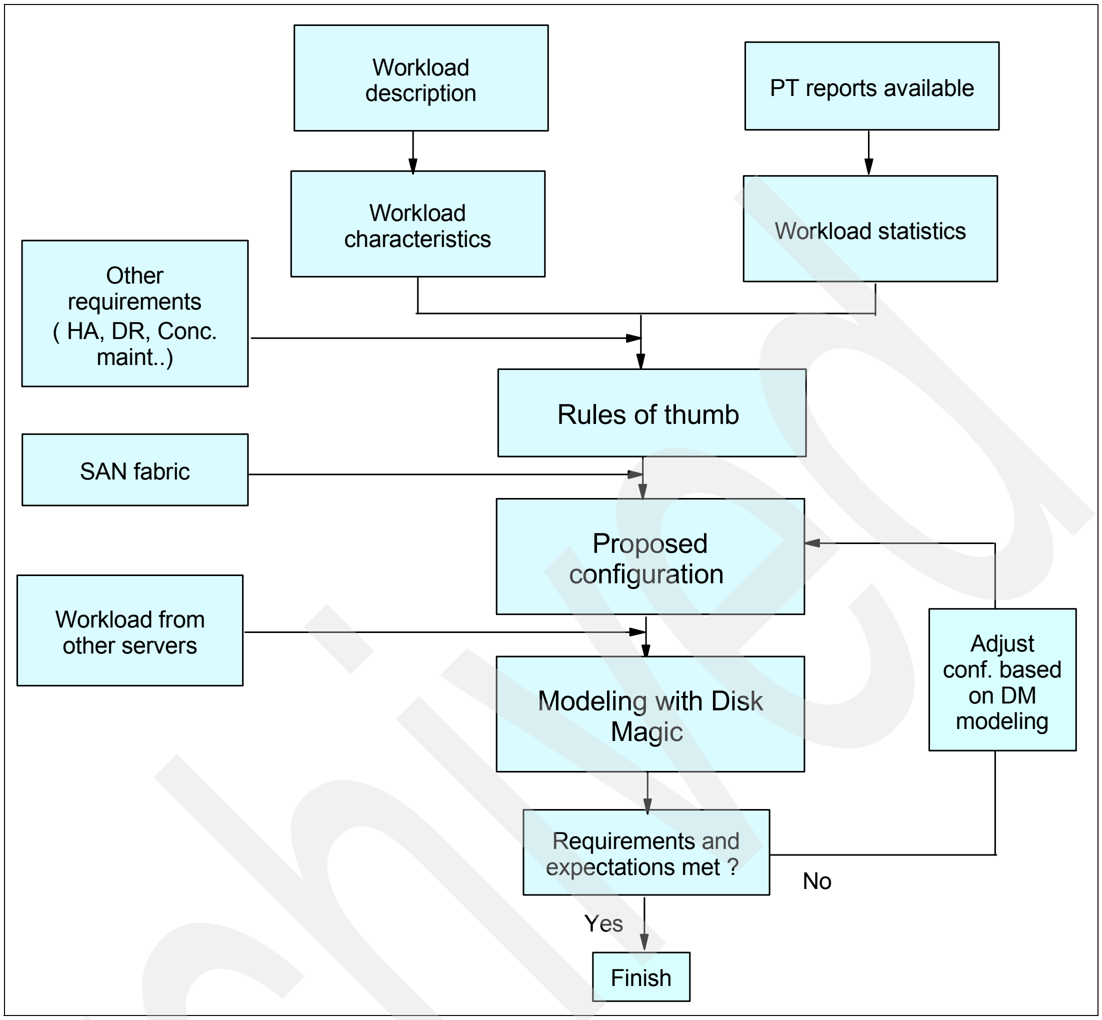

Figure 4-9 Flowchart for planning ESS with iSeries

In the following sections we describe each of the steps and give guidelines how to perform it.

Connectivity

Is iSeries model and OS/400 level appropriate to support FC connectivity? Fibre Channel (FC) connectivity is supported on iSeries models 270 and 8xx, and on @server i5 models. It is supported by iSeries software levels V5R1 and higher.

Is Multipathing considered by the customer? Are enough iSeries FC adapters ordered to implement Multipathing? When planning for Multipath it is important to know that we still support one FC adapter per IOP and 32 LUNs per adapter. With Multipath we really support 32 'LUN paths' per FC adapter, i.e. 2 FC adapters are needed to connect 32 LUNs with 2-way multipath.

Are SAN switches considered? Are the performance recommendations for SAN connectivity met? As a 'rule of thumb' performance guideline it is recommended that max 2 iSeries FC adapters are connected to one ESS FC port via SAN.

Are appropriate cable connectors considered? ISeries FC adapters use Lucent Connectors (LC). When connecting to devices which use Subscriber connectors (SC) appropriate cables and / or interposers have to be considered.

Performance expectations		
Has the customer's critical period been identified ?		
Were performance data taken ? Were they taken in the critical period ?		
If performance data are not available, are characteristics of the workload available ?		
Was Disk Magic used with iSeries performance reports to model ESS response time ?		
Are enough ESS ranks considered ?		
Is appropriate size of DDMs considered ?		
Is appropritae size of LUNs considered ?		
Are enough iSeries FC adapters considered ?		
Does the customer consider sharing of ESS among iSeries and other servers ? Does the customer consider sharing of ESS among multiple iSeries ?		
Is Metro mirror, Global mirror or Global copy used ?		
Are the recommendations of placing IOP/IOA pairs on HSL loops met ?		

Performance expectations

Has the customer's critical period been identified? Before producing the performance reports from collected data the performance requirements and expectations should be considered. In some cases batch window is the most critical period for the customer, for example, if every night a batch process must be run and must complete within defined duration. In other cases transaction workload during the day is most important. In some cases a variable workload is experienced: during certain time periods (5 - 15 minutes) the I/O rates are extremely high when compared to the 5 - 30 minutes immediately preceding or following.

Was performance data taken? Was it they taken during a critical period?

To collect Performance data use the following iSeries commands:

1. GO PERFORM.
1. Select 2. Collect Performance Data.
2. To start collecting data, select 1. Start Collecting Data.
3. On Start Collecting Data menu you can change the values like collection interval, or you can use the defaults.
4. To stop collecting data, select 2. Stop Collecting Data.

To produce performance reports for a critical period, use the following iSeries commands:

1. GO PERFORM.
2. Select 3. Print Performance Report.

If necessary type in the name of the library which contains collected data.

On Print Performance Report - Sample data panel select which type of report you need, for example Resource Report. On select Sections for Report panel select desired section for example Disk Utilization.

When modeling with Disk Magic the following performance reports are required:

- System Report - Disk Utilization
- Resource Interval Report - Disk utilization detail
- Component Report - Disk Activity

To use Disk Magic with iSeries performance reports, each report can be a separate.txt file or all of them concatenated into one.txt file. Use iSeries Navigator to copy performance reports from iSeries spool files to.txt files.

If performance data are not available, are characteristics of the workload available?

We strongly recommend that the customer takes iSeries Performance data for sizing ESS. However there are cases when performance data are not available. In such cases it is helpful to have basic characteristics of the customer's workloads in order to size ESS as good as possible. Some typical iSeries workloads are:

- Transaction workload working with iSeries database
- Batch jobs
- Compute intensive workload like WebSphere or Domino mail
- Following are some characteristics of these workloads:
- Transaction workloadCompute intensive workload
- Read / write ratio 5 / 21 / 2
- Read cache hits 7% 4%
- Write cache efficiency 64% - 79% 61% - 69%
- Transfer size 7 KB/IO14 KB/IO
- IO/second 1.4 CPW
- Access density (IO/second/GB) 1.4 CPW * (70% of disk capacity)

Note: CPW (Commercial Processing Workload) is a value which gives a relative performance rating for each @server i5, iSeries or AS/400 model.

Note: If Performance Reports are available calculate Access Density the following way: divide average IO/second by capacity of occupied disk space. The needed values can be obtained from performance reports.

Are enough ESS ranks considered? Following table gives basic guideline for the number of ESS ranks needed:

- Read / write ratio 100 Consider max 400 - 500 IO/second per rank
- Read / write ration 2.3 Consider max 275 - 325 IO/second per rank
- Read / write ratio 1Consider max 200 - 300 IO/second per rank

Is appropriate size of DDMs considered? Following table gives basic guideline for the size of DDMs:

- Access Density greater than 2 IO/second/GBConsider 18 GB 15 K RPM
- Access Density between 1 and 2 IO/second/GBConsider 18 GB 10 K RPM or 36 GB 15 K RPM
- Access Density lower than 1 IO/second/GB or cache hit ratio greater than 95% Consider 36 GB 10 K RPM or 73 GB 15 K RPM

Consider performance implications of large DDMs. 73 GB 10 K RPM or 145 GB may not give enough arms for acceptable performances.

Is appropriate size of LUNs considered? For performance reason consider as small LUNs as possible, the basic guideline is: at least 2 LUNs per DDM.

Are enough iSeries FC adapters considered? Following table gives basic guideline for the number of FC adapters, based on capacity:

- FC adapter 2787 on IOP 2844(1022 / Access density) GB per FC adapter
- FC adapter 2766 on IOP 2844 (798 / Access density) GB per FC adapter
- FC adapter 2766 on IOP 2943 (644 / Access density) GB per FC adapter

Note: If Access Density is not known you may consider as a ‘rule of thumb’: Access Density = 2

Does the customer consider sharing of ESS among iSeries and other servers? Does the customer consider sharing of ESS among multiple iSeries? Following table shows basic guidelines for sharing ESS:

- iSeries workload Other workload or further iSeries workloadAllocation of LUNs
- Very importantSignificantDivide workloads between ESS clusters
- Medium importantUnimportantDedicate ranks to one or another platform
- Not importantNot significantConsider sharing ranks between iSeries and other servers

If you want to reallocate storage from iSeries to another platform or vice versa the whole array must be reformatted which can be disruptive to the remaining data on the rank, therefore we don't recommend mixing platforms on the same rank.

Is Metro Mirror, Global Mirror or Global Copy used? Consider modelling Metro Mirror with Disk Magic. If extenders are used for connection on long distance consider the latency as specified by the vendor. If dark fibre is used consider sizing bandwidth the following way: multiply writes/second by transfer size, this way you get the estimated I/O rate per second. Make sure that bandwidth is capable of the estimated I/O rate.

Are the recommendations of placing IOP/IOA pairs on HSL loops met? For performance reasons it is recommended to place 1 IOP/IOA pair per Multi Adapter Bridge, and maximum 6 IOP/IOA pairs per HSL loop.

Disaster recovery solutions		
Are the cusotmer's goals well understood and clarified ?		
Does the customer consider ESS remote copy of IASP ? Is Toolkit by Rochester considered ?		
Is iSeries model appropriate to support Toolkit for copy services of IASP ?		
Is recovery procedure for planned and unplanned outages well understood by the customer ?		
Is expected recovery time for planned and unplanned outages well understood by the customer ?		
Does the customer consider reverse of ESS remote copy ?		
Is enough ESS capacity planned on production and DR site ?		
Are enough iSeries FC adapters considered on production and DR site ?		
If ESS copy services of IASP are considered, are the skills available to setup iSeries applications in an IASP ?		

Disaster recovery solutions

Does the customer consider ESS remote copy of IASP? Is the Toolkit ESS Copy Services by Rochester considered? An iSeries independent disk pool, or independent auxiliary storage pool (IASP), is a collection of disk units that can be brought online or taken offline independently of the rest of the disk storage on a system. ESS copy functions in combination with iSeries IASPs enable customers to use Metro Mirror, Global Mirror and Global Copy with minimal recovery time. This is only supported by using the iSeries Copy Services for ESS Toolkit offered by Rochester Continuous Availability Team. The Toolkit provides necessary code and services for running ESS copy functions of an IASP.

As an exception to remote copy of IASP it is possible to use remote copy of entire iSeries disk space. This solution needs longer recovery time and doesn't prevent from a high availability exposure which is prevented by using Copy services of IASP. However, we strongly recommend you should use the Toolkit of IASP.

Is iSeries model appropriate to support Toolkit for copy services of IASP? Toolkit for copy services of IASP is presently not yet supported on @server i5. The support is expected beginning of 2005. For the customers who are looking for this DR solution it is advisable to obtain iSeries model 8xx and upgrade it to @server i5 when the Toolkit will be supported on these model range, or implement i5 with IASP and wait until the Toolkit is available with i5.

Is expected recovery time for planned and unplanned outages well understood by the customer? When using ESS copy services of IASP the experienced recovery time for planned outages is about 5 - 10 minutes plus the time it takes to end the applications; for unplanned outages you may experience recovery times of about 2 hours.

Does the customer consider reverse of ESS remote copy? If Global Mirror or Global Copy are used it may be necessary to purchase additional capacity in order to be able to reverse the remote copy after the failed site is fixed and running again.

For example: you may want to install 73 GB 15 K RPM drives instead of 36 GB 15 K RPM drives. This would provide double capacity with the same number of arms required for performance.

Is enough ESS capacity planned on production and DR site? In case Global mirror or Global copy are used the capacity for additional Flashcopy has to be considered, on primary site as well as on remote site.

Are enough iSeries FC adapters considered on production and DR site? If DR solution contains multiple IASPs you should consider that each IASP requires at least one iSeries FC adapter, being on production site or remote site.

Solution for minimizing backup window		
Are the cusotmer's goals well understood and clarified ?		
Does the customer consider ESS Flashcopy of IASP ? Is Toolkit by Rochester considered ?		
Is iSeries model appropriate to support Toolkit for copy services of IASP ?		
Is the procedure for taking backups from Flashocpy well understood by the customer ?		
Is the estimated downtime of applications well understood by the customer ?		
Is usage of BRMS in connection with this solution well understood by the customer ?		
In case the procedure needs to be fully automated is the Windows server for Flashcopy CLI available ?		
If Flashcopy of of IASP is considered, are the skills available to setup iSeries applications in an IASP ?		

Solution for minimizing backup window

Does the customer consider ESS Flashcopy of IASP? Is Toolkit by Rochester considered? An iSeries independent disk pool, or independent auxiliary storage pool (IASP), is a collection of disk units that can be brought online or taken offline independent of the rest of the disk storage on a system. ESS copy functions in combination with iSeries IASPs enable customers to use Flashcopy for taking backups with minimal downtime of production applications. This is only supported by using the iSeries Copy Services for ESS Toolkit offered by Rochester Continuous Availability Team. The Toolkit provides necessary code and services for running ESS copy functions of an IASP.

Is iSeries model appropriate to support Toolkit for copy services of IASP? Toolkit for copy services of IASP is presently not yet supported on @server i5. The support is expected beginning of 2005. For the customers who are looking for this solution it is advisable to obtain iSeries model 8xx and upgrade it to @server i5 when the Toolkit will be supported on these model range, or implement i5 with IASP and wait until the Toolkit is available with i5.

Is the estimated downtime of applications well understood by the customer? If Toolkit for Copy services of IASP is used for the estimated downtime of IASP and therefore of the applications which run in it, is about 10 - 15 minutes plus the time it takes to end the applications.

Is usage of BRMS in connection with this solution well understood by the customer? iSeries software Backup, Recovery and Media Services (BRMS) provides users with planning, controlling and automating of backup, recovery and media management. When BRMS is used to take backups from Flashcopy of IASP, the BRMS Networking feature should be considered in order to share backup information between the backup server / partition and the production server / partition.

In case the procedure needs to be fully automated is the Windows server for Flashcopy CLI available? With Flashcopy solutions for iSeries, CLI for Flashcopy can be done via a Windows server: iSeries sends to Windows server the command which triggers the CLI for Flashcopy.

High availability		
Is Recovery Time Objects (RTO) versus Recovery Point Objects (RPO) understood by customer ?		
Is Multipathing considered by the customer ?		

High availability

Is Multipathing considered by the customer? When properly implemented Multipathing provides high availability in case an iSeries adapter, an ESS host adapter or connection between the two fails. It is built-in OS/400 or i5/OS, there is no need for a driver. It requires at least iSeries software level V5R3.

The customers who implemented iSeries system mirroring in connection with ESS before Multipath was available, may consider the following:

- Local High Availability in case of failure of iSeries or ESS adapters or connection between them, is provided by Multipath.
- Functionality provided by possibility to recover from remote mirrored half of volumes, can be implemented by using ESS Metro Mirror of an IASP.

Skills and documentation		
Are iSeries and Storage specialists included in the project ?		
If Business partner is running the project, are iSeries and Storage people from BP involved ?		
If ESS Copy services for IASP are used, is ATS involved ?		
Is the redbook iSeries in Storage area Networks SG24-6220-00 used ?		
Is the ESS on iSeries Proposal Planning Guide used ?		
Are appropriate IBM Web links used ?		

Skills and documentation

If ESS Copy services for IASP are used, is ATS involved? ATS requests can be created as Siebel Service Requests in EMEA and assigned to the queue ‘EMEA ATS STORAGE’.

Information about EMEA Storage ATS can be found on the following Web link:

http://ebcweb.mainz.de.ibm.com/ATSSE/atsse.nsf/AllNetdocs/About?OpenDocument

Is the ESS on iSeries Proposal Planning Guide used? iSeries Proposal Planning Guide can be found on the following Web page:

http://w3-1.ibm.com/support/ESSonassure/assur30i.nsf/WebIndex/SA553

Are appropriate IBM Web links used? It is recommended to use the following Web links:

iSeries Technology Center (iTC):

http://www-1.ibm.com/servers/eserver/iseries/service/itc/technicalmarketing.htm

iSeries Storage and SAN solutions:

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

IBM Global Services - Rochester iSeries Services Group (expert, fee based consulting):

http://www-1.ibm.com/servers/eserver/iseries/service/igs/

ISeries Workload Estimator:

http://www-912.ibm.com/supporthome.nsf/document/16533356

iSeries Storage and SAN solutions:

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

iSeries SAN Interoperability Support:

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

4.3.1 Obtain workload characteristics from available information

We basically divide iSeries workload to the following three types:

- Transaction applications working with iSeries database
- Batch jobs
- Compute intensive workload, like WebSphere® or Domino mail

In case you have only limited information about current iSeries, like rough description of workload and needed capacity, figure out to which listed type of workload does it belong. Use the following table to obtain basic characteristics for each of these workload types:

Table 4-2

Characteristics	**Workload**		
	Transaction appl.	Batch jobs	Compute intensive
Read/Write ratio	2.5		0.5
Read cache hits	7%		4%
Write cache efficiency	64%-79%		61%-69%
Transfer size in KB/IO	7KB		14KB
I/O per second	1.4		
Access density IO/second/GB	1.4 CPW (70% capacity)		

Note: As is indicated in the table, calculate IO/second by multiplying processor CPW of iSeries model with 1.4. For example: iSeries model 825 (6 way) has processor CPW 6600. Characteristics of Transaction application workload on this model is: 1,4 * 6600 = 9240 IO/second.

In our experience, it is most likely that the iSeries Access densities (IO/second/GB) will be between 1 and 2.

4.3.2 Obtain workload statistics from PT reports

If PT reports of present iSeries are available, use the instructions in the following table to obtain workload statistics:

Table 4-3 Workload statistics

Statistics	Where to get it
IO/second	Obtain Average IO/second from Resource Interval Report - Disk Utilization Summary
Read/Write ratio	Obtain Average read/second Average write/second from Resource Interval Report - Disk Utilization Summary. Divide read/se by write/second
Cache hits	Obtain Cache hit statistics from Component report - Disk Activity
Transfer size in KB/IO	Obtain Average K per I/O form Resource Interval Report - Disk Utilization Summary
Access density IO/second/GB	From System report - Disk utilization obtain Total size of disk units and Percent full (Percentage of disk space capacity in use) Calculate “Occupied disk space” by the following formula: Total size of disk units * Percent full / 100 Divide Average IO/second by Occupied disk space

4.3.3 Other requirements

If there are requirements for High availability, concurrent maintenance, Disaster recovery, etc., you should take them into account when you will use rules of thumb for planning ESS. Some of such requirements are listed below:

- If there is need for Fault tolerance in case one iSeries adapter fails and concurrent maintenance of an ESS host bay, consider OS/400 mirroring within ESS, each iSeries FC adapter connected to one ESS Host adapter. For more information about mirroring see Chapter 7, “Mirroring to an ESS” on page 165.
- If there is requirement for concurrent maintenance if an ESS Host bay, consider partial failover scenario with a switch. In this scenario there is one link from iSeries FC adapter to a switch and two links from the switch to two ESS host adapters. The LUNs are assigned to iSeries
- If disaster recovery solution is required consider PPRC between to ESSs
- If there is a requirement to minimize backup window consider Flashcopy
- If fault tolerance and disaster recovery is required consider mirroring between internal disks and ESS, or mirroring between local and remote ESS

Table 4-4 on page 82 shows these requirements and what has to be taken into account for each of them:

Table 4-4 Other requirements

Scenario to meet a requirement	What has to be taken into account
Mirroring within ESS	Plan for double disk capacity, double iSeries FC adapters, double ESS Host adapters
Partial failover	Plan two ESS host adapters for one iSeries FC adapter
PPRC	Plan ESCON links on local and this may limit the number of available ESS host adapters
Flashcopy	Plan for double disk capacity
Mirroring between internal disks and ESS	Plan for the same number and size of LUNs as are internal disks For bus level mirroring plan for FC adapters on other iSeries buses than the buses with internal disks
Mirroring between local and remote ESS	Plan the same disk capacity for both ESSs

4.3.4 SAN fabric

If you plan to connect multiple iSeries FC adapters to one ESS host adapter through switches and you don-t have PT reports consider the following rule of thumb:

Table 4-5

Workload	Number of iSeries FC adapters per ESS HBA
Transaction appl.	
Batch jobs	
Compute intensive	

If PT reports are available refer to 4.2, “Consideration for planning ESS with iSeries” on page 66 for more accurate planning.

4.3.5 Rules of thumb

After you obtained and calculated features of the present iSeries workload use the sizing rules of thumb, depending of the workload features. Following table shows these rules of thumb for the most important sizing factors: Number of ranks, capacity and speed of DDMs, number of FC adapters, capacity of LUNs.

Table 4-6 Rules of thumb

Rules of thumb	
Number of ESS ranks	
Read/Write ratio 100	Consider maximum 400 - 500 IO/second per rank
Read/Write ratio 2.3	Consider maximum 275 - 325 IO/second per rank
Read/Write ratio 1	Consider maximum 200 - 300 IO/second per rank
If PT reports are available consider to use iSeries ESS Sizing Spreadsheet which is available through ATS	

Rules of thumb	
Capacity and speed of DDMs	
Access density greater than 2 IO/second/GB	Consider 18 GB 15,000 RPM
Access density between 1 and 2 IO/second/GB	Consider 18 GB 10,000 RPM or 36 GB 15,000 RPM
Access density lower than 1 IO/second/GB, or Cache hit ratio greater than 95%	Consider 36 GB 10,000 RPM or 73 GB 15,000 RPM
Carefully plan for performances	73 GB 10,000 RPM or 145 GB 15,000 RPM
Number of iSeries FC adapters	
FC adapter 2787 on IOP 2844	1022 / Access density GB per FC adapter
FC adapter 2766 on IOP 2844	798 / Access density
FC adapter 2766 on IOP 2843	644 / Access density
If you plan direct connection of iSeries to ESS consider one ESS host adapter per one iSeries FC adapter. If you plan to connect it via switches refer to 4.3.4, "SAN fabric" on page 82	
Capacity of LUNs	
FC adapter 2787 on IOP 2844	Consider maximum 31 / Access density GB
FC adapter 2766 on IOP 2844	25 / Access density
FC adapter 2766 on IOP 2843	20 / Access density
If PT reports are available consider to use iSeries ESS Sizing Spreadsheet which is available through ATS	

Note: Recommended capacity per FC adapter is calculated the following way: we divide max IO/second per adapter at 70% utilization, by Access density (AD). This way we get the capacity which one FC adapter is capable of handling at 70% utilization. Since it is recommended that LUN utilization is 40%, we apply 40% to the calculated capacity per FC adapter.

For example: FC adapter 2787 on IOP 2488 is capable of 2555 IO/second at 70% utilization. We divide it by AD and take 40% of the result: 0,4 * (2555 / AD) = 1022 / AD

See 4.2.3, "iSeries Fibre Channel adapters" on page 67 for maximum IO/second that an FC adapter can handle at 70% utilization.

Note: We calculate maximal recommended capacity of LUN by dividing capacity per FC adapter with 32.

For example: max capacity for 2787-2844 is (1022 / AD). Maximal LUN size for this adapter-IOP is (1022 / AD) / 32 = 31 / AD

4.3.6 Proposed configuration

The listed rules of thumb help you to prepare initial ESS configuration. You may also prepare the configuration of additional FC adapters to iSeries and the configuration of switches. You may want to know the price of proposed configuration by using the configurator e-config.

4.3.7 Workload from other servers

If you plan to consolidate storage from iSeries and other servers, like Unix or Windows NT servers, to the same ESS use Disk Magic for modeling both iSeries and other servers workloads. In this case you should obtain information obit other servers workload in order to insert it in Disk Magic together with iSeries workload.

Use the following rules of thumb for allocating storage for iSeries and other servers:

Table 4-7 Allocating storage

iSeries workload	Other workload	Allocation of LUNs
Very Important	Significant	Divide workloads between ESS clusters
Medium important	Unimportant	Dedicate ranks to one or another platform
Not important	Not significant	Consider sharing ranks between iSeries and other servers

4.3.8 Modeling with Disk Magic

Insert proposed configuration into Disk Magic the following way:

- If PT reports are not available insert the data manually.
- If PT reports are available use them as input to Disk Magic and adjust the values so that they will reflect proposed configuration. For example: If you calculated that 6 FC adapters are needed, insert the PT reports to Disk Magic and afterwards change the number of FC adapters to 6.

Run Disk Magic for proposed configuration and observe the disk response time.

It is most likely that the proposed configuration will not meet requirements and expectations for disk response time, price, etc. In this case model changes in Disk Magic and observe response time. Adjust the proposed configuration the modeled changes. You may also use e-config in parallel with Disk Magic to observe the price of adjusted configuration.

Repeat modeling with Disk Magic until the proposed configuration meets all expectations as much as possible.

For instruction about how to use Disk Magic, refer to 4.4, “Using Disk Magic for planning ESS with iSeries” on page 84.

4.4 Using Disk Magic for planning ESS with iSeries

Disk Magic is a tool intended for sizing and modeling disk systems with for various servers.It can be used for modeling IBM and other disk systems attached to iSeries, zSeries, pSeries and other servers. It was developed by the company IntelliMagic and can be downloaded from the following Web pages:

- For IBM Professionals, go to the Web page:

  http://w3-1.ibm.com/sales/systems/

  Select Tools, and select Disk Magic.
- For Business Partners:

  http://www.ibm.com/partnerworld/sales/systems/

  Search for ‘Disk Magic’.

For modeling disk system for any server, Disk Magic requires information about current I/O performances of the server.

In this section we describe three examples of modeling ESS for iSeries with Disk Magic, and we also give instructions how to get the needed performance data from iSeries.

4.4.1 iSeries performance reports

iSeries performance reports can be produced on any iSeries server which has installed the licensed product IBM Performance Tools for iSeries. To be able to get iSeries performance reports you have to collect performance data on the iSeries. We recommend that performance data is collected during a period of 24 hours. From collected data you will be able to produce performance reports for different critical time periods.

To collect performance data use the following iSeries commands:

```
GO PERFORM
select 2. Collect performance data
```

To start collecting data:

```
select 1. Start collecting data
```

On Start Collecting Data menu you can change the values like collection interval, or you use default ones.

To stop collecting data:

```
select 2. Stop collecting data
```

On the Start Collecting Data menu you can change the values like collection interval, or you can use the defaults.

Before producing the performance reports form collected data you should consider the performance requirements and expectations. In some cases batch window is the most critical period. For example, if every night a batch process must be run and that process must complete within defined duration. In other cases transaction workload during the day is most important. In some cases the bursty workload is experienced: A bursty I/O workload is one where during certain time periods (5-15 minutes) the I/O rates are extremely high when compared the 5-30 minutes immediately preceding or following.

Based on performance requirements and expectations select the most important time period. Produce performance reports for that period using the following iSeries commands:

```
GO PERFORM
Select 3. Print performance report
```

If necessary type in the name of the library which contains collected data

On Print Performance Report - Sample Data panel, select which type of report you need, for example, Resource Report. If necessary select a member (members represent different collection periods).

On Select Sections for Report panel, select desired section, for example Disk Utilization.
On Select Categories for Report panel, select Time interval.
On Select Time Intervals panel, select intervals for the critical period.
On specify Report Options pane, specify report title.

The report will be produced through a batch job, and will be available as a spool file.

For modeling with Disk Magic the following performance reports are required:

- System Report - Disk utilization
- Resource Interval Report - Disk utilization detail
- Component Report - Disk activity

The following report is optional for Disk Magic: system Report - Storage pool utilization.

For use with Disk Magic, iSeries performance reports each must be in one.txt file or all of them concatenated in one.txt file. Use iSeries Navigator to copy performance reports from iSeries spool files to.txt files.

4.4.2 Examples of using Disk Magic for iSeries

In this section we describe how to use Disk Magic in three examples:

- Modeling ESS for a single iSeries
- Modeling ESS for two iSeries
- Modeling ESS for one iSeries and two Windows systems

Modeling ESS for one iSeries

In this example we use Disk Magic for sizing ESS for an iSeries which has presently only internal disks. Besides, we will use Disk Magic to predict how disk response time changes at the following conditions:

- When we use different number of iSeries FC adapters
- When we use different capacity of ESS DDMs
- When we use different capacity of LUNs

Currently internal disks are mirrored, but on ESS we plan RAID protection so we will not use OS/400 mirroring. The needed Performance reports are concatenated in one.txt file.

To start Disk Magic, from the screen Welcome to Disk Magic, select iSeries PT Report file - with extension TXT, as shown in Figure 4-10.

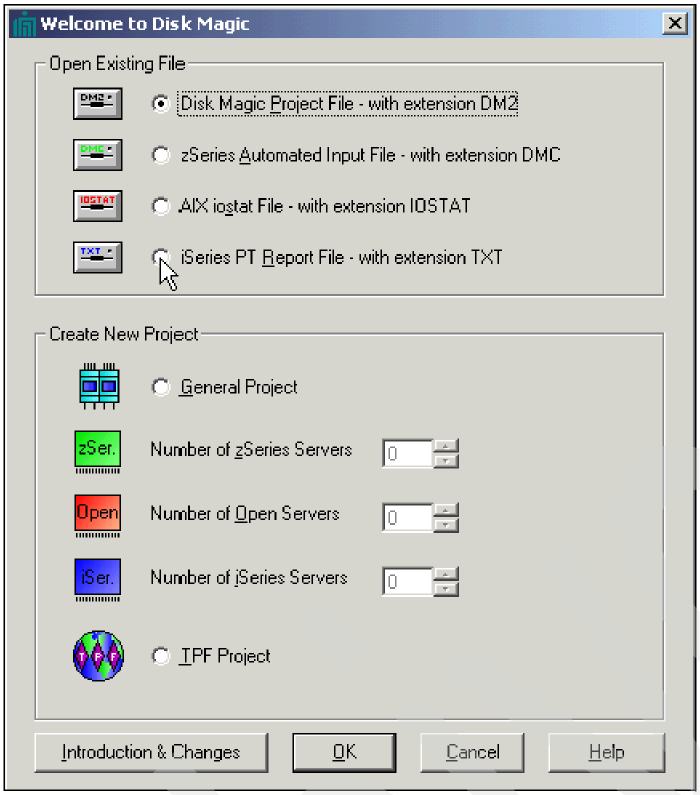

Figure 4-10 DM for single iSeries - importing files

Select the folder and .txt file which you want to import, as shown in Figure 4-11.

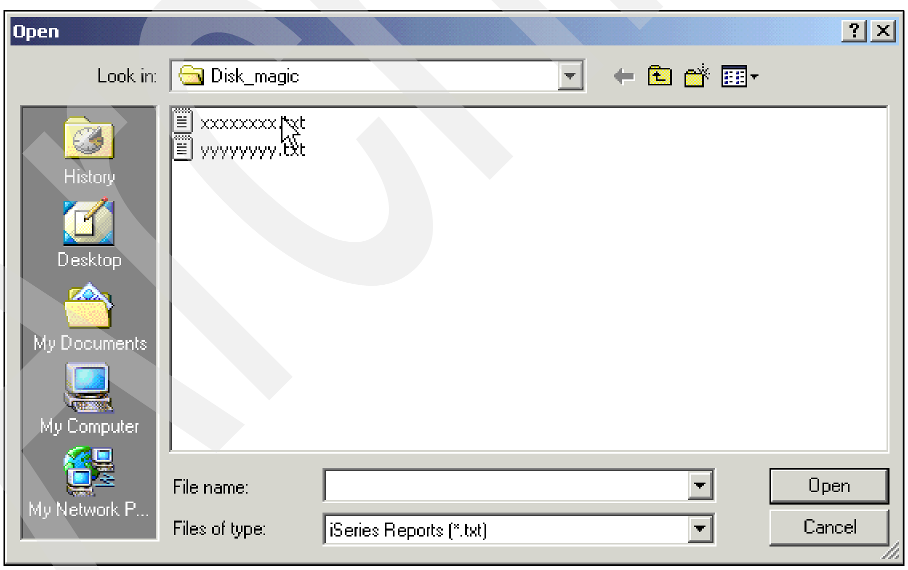

Figure 4-11 DM for single iSeries - selecting files

Because iSeries internal disk are mirrored Disk Magic offers a choice to continue mirroring within ESS, or to use only ESS RAID protection. IN this example we select only ESS RAID protection, as seen in Figure 4-12.

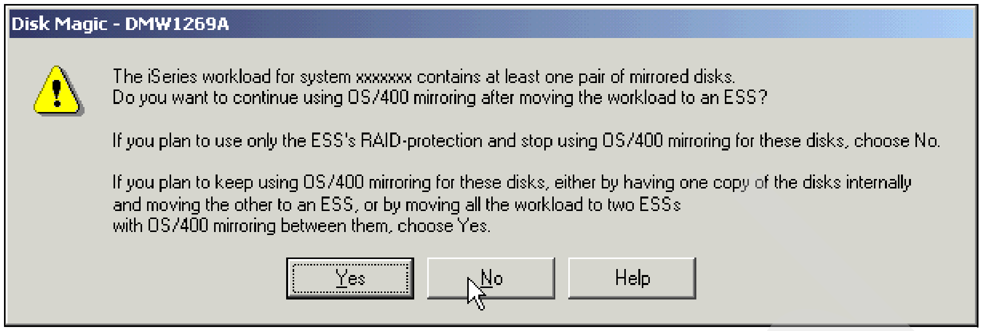

Figure 4-12 DM for single iSeries -selecting RAID protection

After the input file is successfully processed, Disk Magic shows the report of present I/O activity on internal disks.

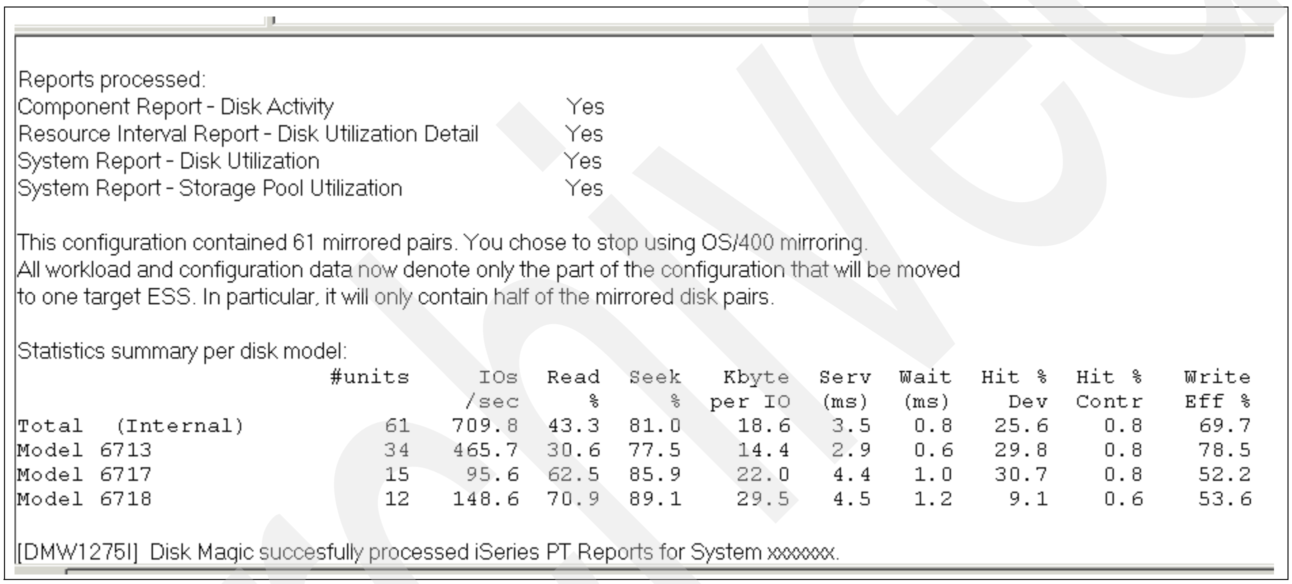

Reports processed:

Component Report - Disk Activity	Yes
Resource Interval Report - Disk Utilization Detail	Yes
System Report - Disk Utilization	Yes
System Report - Storage Pool Utilization	Yes

This configuration contained 61 mirrored pairs. You chose to stop using OS/400 mirroring.
All workload and configuration data now denote only the part of the configuration that will be moved to one target ESS. In particular, it will only contain half of the mirrored disk pairs.

Statistics summary per disk model:

	#units	IOs /sec	Read %	Seek %	Kbyte per IO	Serv (ms)	Wait (ms)	Hit % Dev	Hit % Contr	Write Eff %
Total (Internal)	61	709.8	43.3	81.0	18.6	3.5	0.8	25.6	0.8	69.7
Model 6713	34	465.7	30.6	77.5	14.4	2.9	0.6	29.8	0.8	78.5
Model 6717	15	95.6	62.5	85.9	22.0	4.4	1.0	30.7	0.8	52.2
Model 6718	12	148.6	70.9	89.1	29.5	4.5	1.2	9.1	0.6	53.6

[DMW1275I] Disk Magic succesfully processed iSeries PT Reports for System xxxxxxx.

Figure 4-13 DM for Single iSeries - present I/O activity

iSeries disk subsystem (currently internal disks) is represented in Disk Magic with the green icon. For modeling we double-click the Disk subsystem icon. Diks Magic shows the Disk subsystem window, as seen in Figure 4-14.

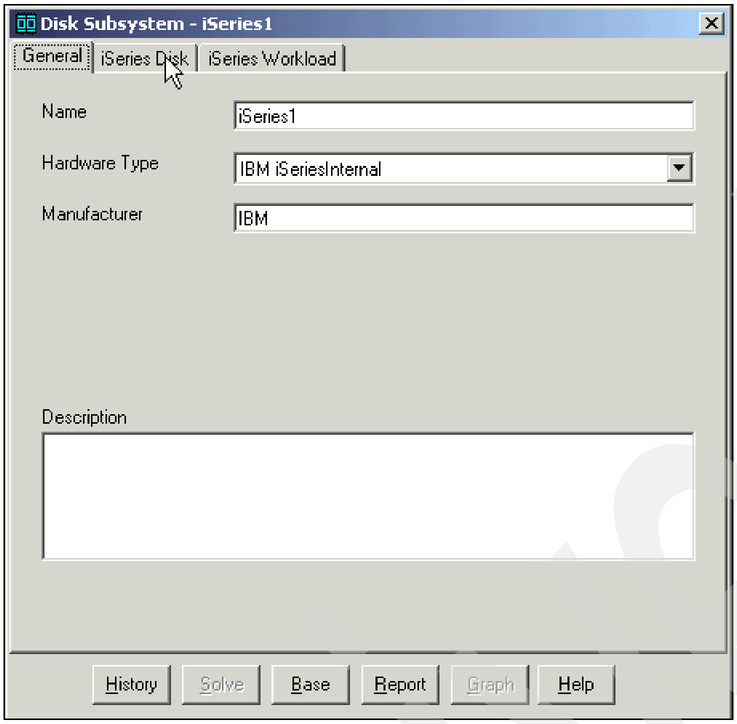

Figure 4-14 DM for Single iSeries - Disk Subsystem window

By clicking the tab iSeries Disk on Disk subsystem window we can observe present disk capacity and number of disk devices, as shown in Figure 4-15.

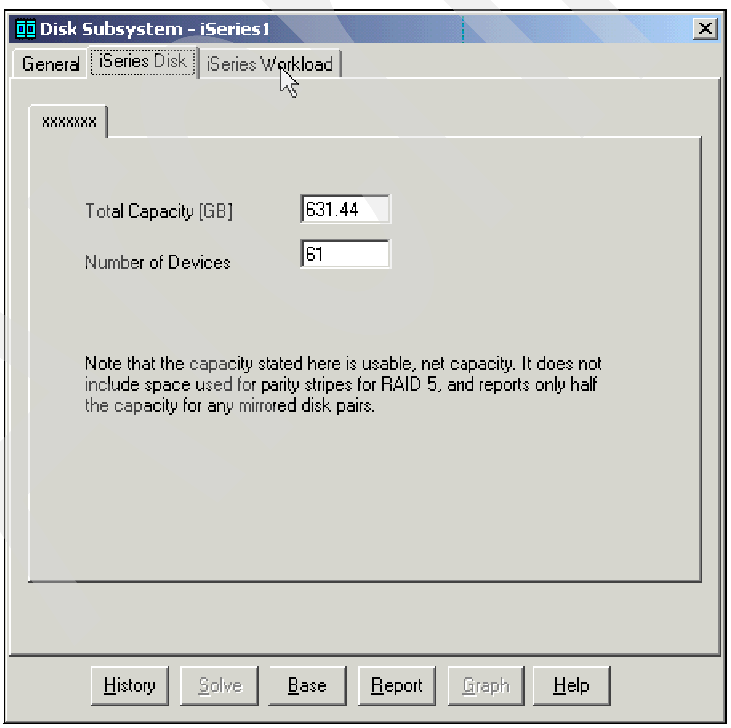

Figure 4-15 DM for Single iSeries - Present disk capacity

By clicking the tab iSeries workload on Disk Subsystem window we can observe the characteristics of the present iSeries workload, like I/O per second, read/write ratio and transfer size. This is shown in Figure 4-16.

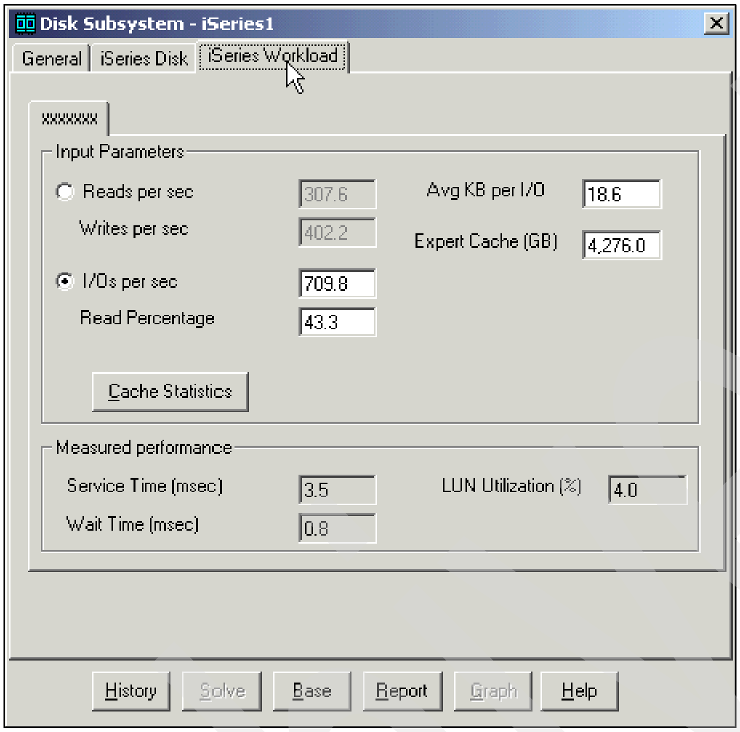

Figure 4-16 DM for Single iSeries - Present I/O activity

You can see cache statistics from current iSeries workload by clicking the button Cache Statistics.

Save iSeries present Disk subsystem as the base for further modeling. This is accomplished by clicking button Base on Disk Subsystem window, as shown in Figure 4-17.

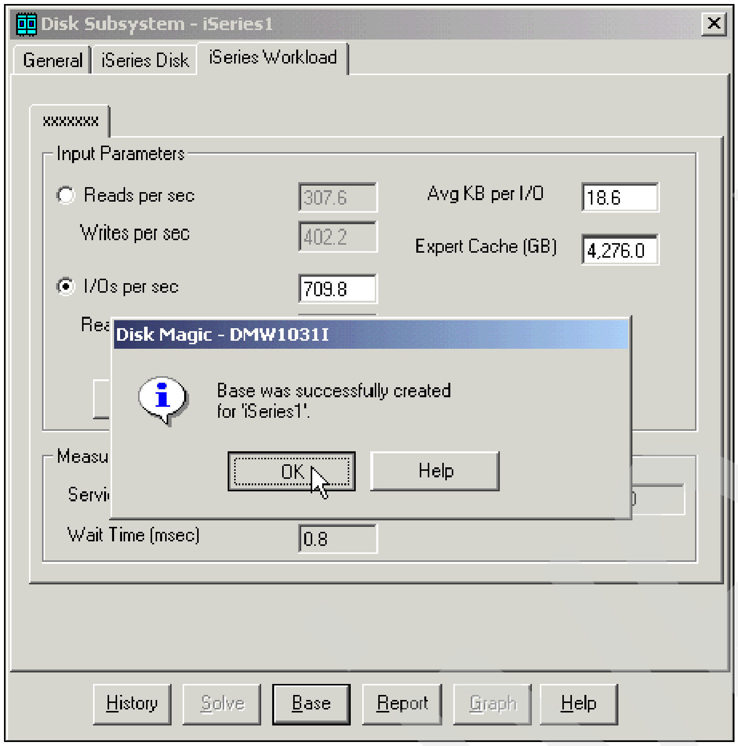

Figure 4-17 DM for Single iSeries - Create base for modeling

After the base is created, start the modeling disk subsystem for ESS. On the General tab of the Disk subsystem window, select Hardware type ESS 800, as shown in Figure 4-18.

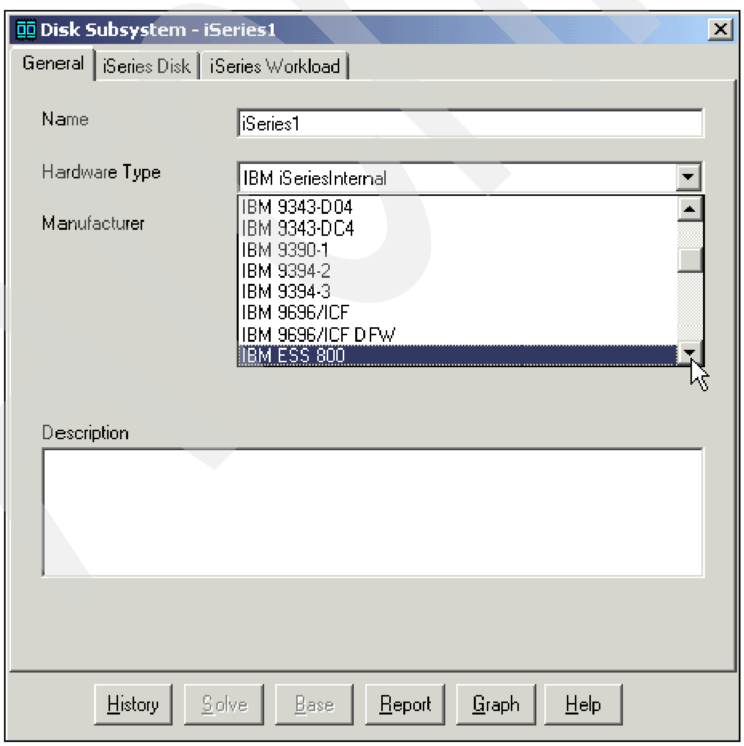

Figure 4-18 DM for Single iSeries - Selecting ESS as disk subsystem

Click Hardware Details. Disk Magic shows a window in which you can adjust the number of host adapters and cache size, as shown in Figure 4-19.

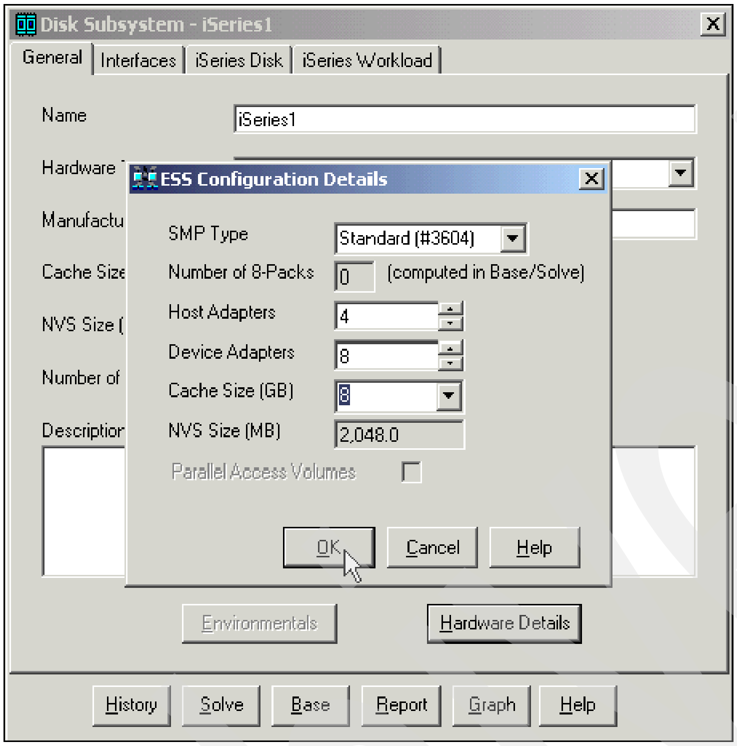

Figure 4-19 DM for Single iSeries - Adjusting ESS

By clicking on iSeries Disk tab you see the initial disk capacity, DDM capacity, type of protection and number of LUNs and number of ranks, as sized by Disk Magic. This is shown in Figure 4-20.

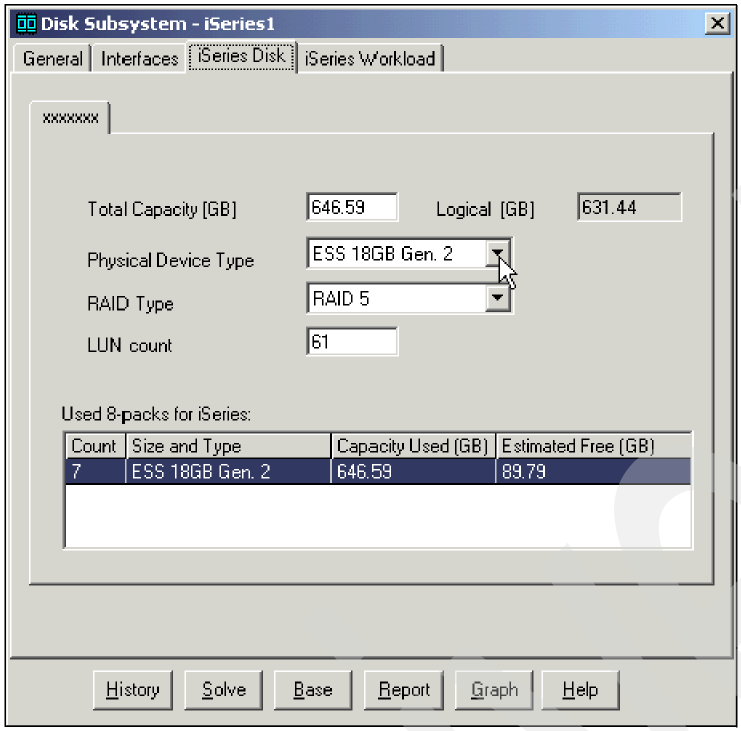

Figure 4-20 DM for Single iSeries - Initial values

You may want to change the type of DDMs. In this example we changed them to 36 GB / 15000 RPM ones. You may also consider to change the number of FC adapters. To achieve this, click the Interfaces table and change the number of interfaces from servers and from disk subsystem. In this example we changed them to 2, as shown in Figure 4-21.

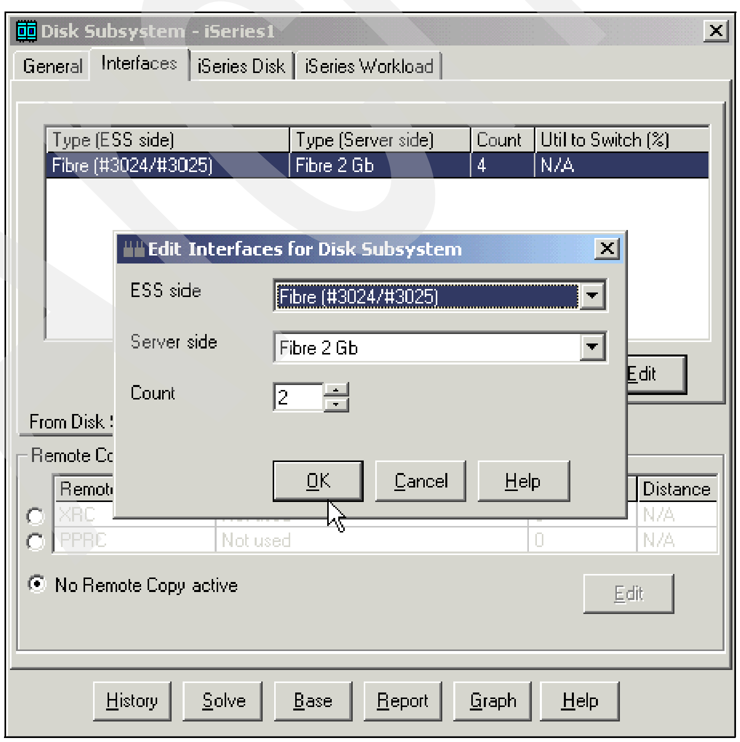

Figure 4-21 DM for Single iSeries - Changing the number of FC adapters

After you have changed all the values that you consider necessary, solve the new configuration by clicking Solve. This is shown in Figure 4-22.

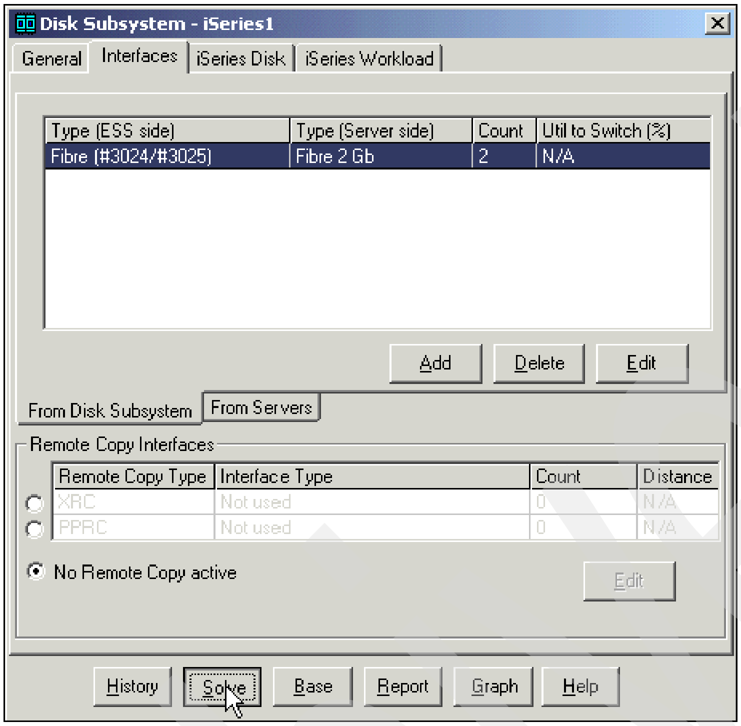

Figure 4-22 DM for Single iSeries - Solving ESS configuration

To create a report of I/O activity, click Report. Disk Magic shows the report of several values, such as disk response time, host interface utilization, iSeries LUN utilization, etc. This report is shown in Figure 4-23.

```
Advanced ESS Outputs ():
Avg. SMP Utilization:            2.2
Highest HDD Utilization:        11.7
Highest Device Adapter Util:     2.8
Upper Internal Bus Utilization:  3.1
Avg. Host Interface Utilization: 14.4
Avg. iSeries LUN Utilization:    2.7


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit% Hit%  Eff % Cnt Util
Average         710    18.6    2.3    43   59  100     70  61    3
xxxxxxx         710    18.6    2.3    43   59  100     70  61    3

The following parameters are estimated during the modeling process
       Host interfaces            14% busy
       Disk interfaces             3% busy
```

Figure 4-23 DM for Single iSeries - Report for initial ESS configuration

If you want to show the disk response time in graphical representation, click Graph, select the data you want to see in the graph and click Plot. For this example we selected to plot Response Time components in ms, as shown in Figure 4-24.

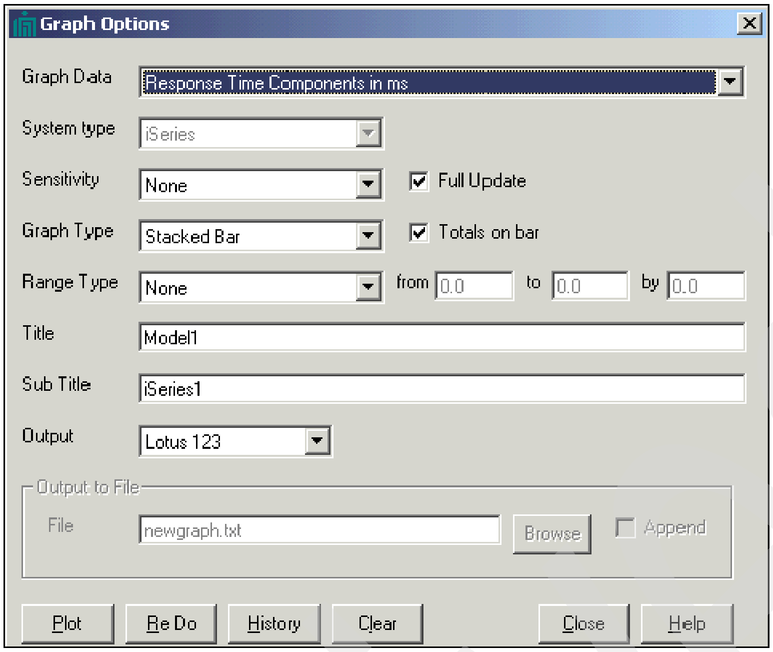

Figure 4-24 DM for Single iSeries - Graphical representation

Disk Magic creates a spreadsheet with the bar which represents selected data. In our example the bar graph contains Disk response time components: Service time and wait time. This is shown in Figure 4-24.

Figure 4-25 DM for Single iSeries - Bar graph of disk response time

For modeling with Disk Magic, change ESS configuration and observe how the change influences I/O activity. In this example we first changed the number of FC adapters from 2 to 3, as shown in Figure 4-26.

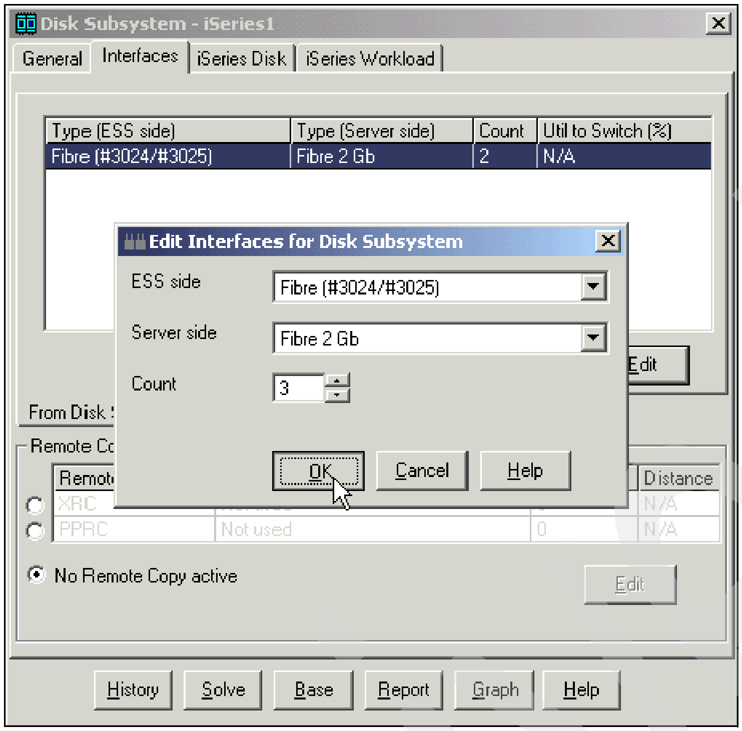

Figure 4-26 DM for Single iSeries - 3 FC adapters

Solve the changed configuration and produce the report. In our example we can see from report that disk response time didn’t change, but we see that host interface utilization is lower than with previous reports (2 FC adapters). The new report is shown in Figure 4-27.

```
Cache Size / Backstore Sensitivity          6.0

Advanced ESS Outputs ():
Avg. SMP Utilization:                2.2
Highest HDD Utilization:            11.7
Highest Device Adapter Util:         2.8
Upper Internal Bus Utilization:      3.1
Avg. Host Interface Utilization:     9.6
Avg. iSeries LUN Utilization:     2.7


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit%  Hit% Eff % Cnt Util%
Average         710    18.6    2.3   43   59   100     70  61    3
xxxxxxx         710    18.6    2.3   43   59   100     70  61    3

The following parameters are estimated during the modeling process
       Host interfaces             10% busy
       Disk interfaces              3% busy
```

Figure 4-27 DM for Single iSeries - 3 FC adapters - Report

To model another change in ESS configuration first restore initial ESS configuration. To achieve this click History, select the initial ESS configuration and click Restore. This is shown in Figure 4-28.

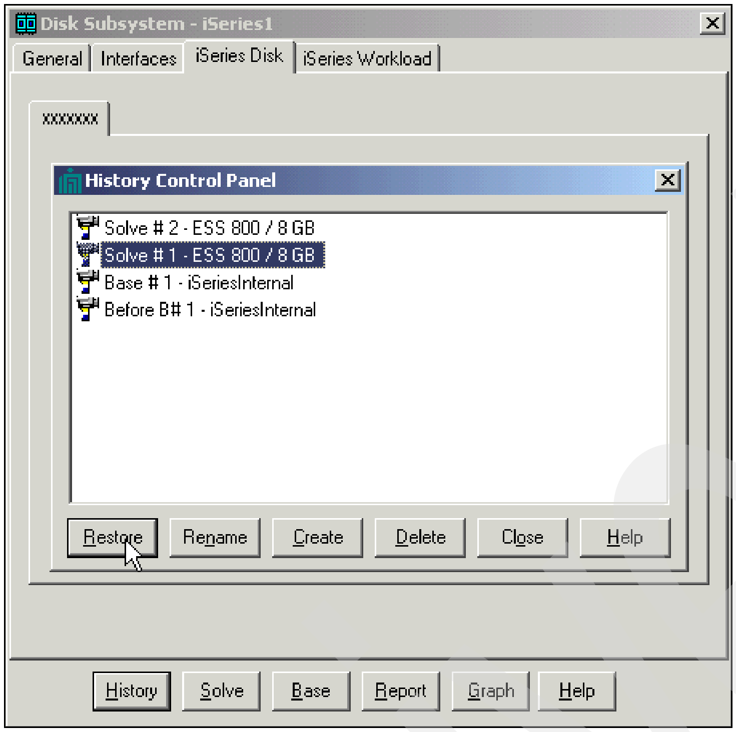

Figure 4-28   DM for Single iSeries - Restore initial ESS configuration

Now you can model another change in configuration. For this example we changed the capacity of DDMs from 36 GB / 15000 RPM to 18 GB/ 15000 RPM, as shown in Figure 4-29.

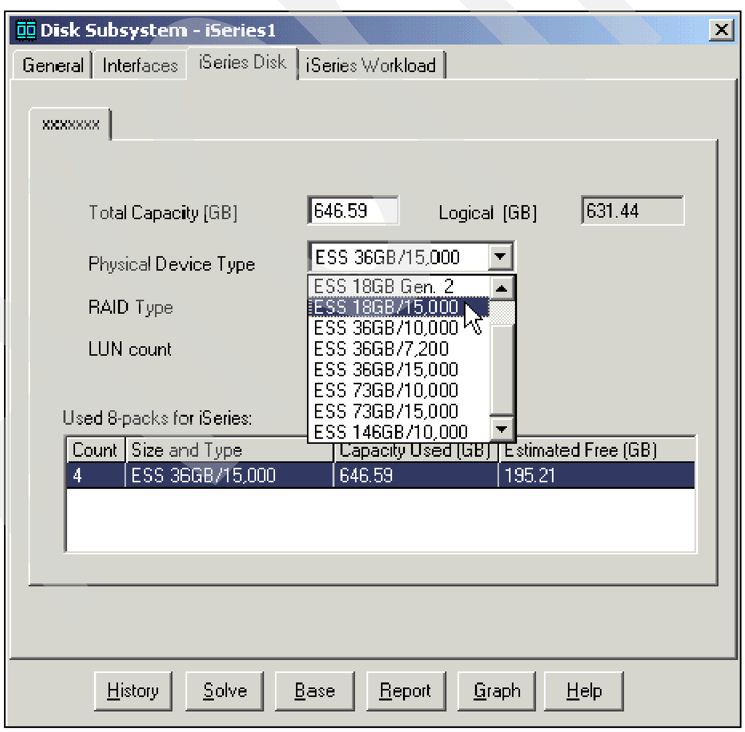

Figure 4-29   DM for Single iSeries - 18 GB DDM

Solve the changed configuration and produce the report. In our example the smaller capacity of DDMs influences lower disk response time and lower LUN utilization. This can be observed in the report shown in Figure 4-30. You can also produce a graphical report which contains a bar graph of the new disk response time.

```
Upper Internal Bus Utilization:    3.5
Avg. Host Interface Utilization:  14.4
Avg. iSeries LUN Utilization:     2.4


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit%  Hit% Eff % Cnt Util%
Average         710    18.6    2.1    43   59  100     70  61    2
xxxxxxx         710    18.6    2.1    43   59  100     70  61    2

The following parameters are estimated during the modeling process
       Host interfaces           14% busy
       Disk interfaces            2% busy




Base model for this configuration
Project                    1
```

Figure 4-30   DM for Single iSeries - 18 GB DDMs - Report

To model another change in configuration first restore the initial ESS configuration and make a desired change. For this example we changed the number and size of LUNs by specifying 32 LUNs on iSeries disk panel of Disk subsystem window. This is shown in Figure 4-31. We also changed the number of FC adapters to 1.

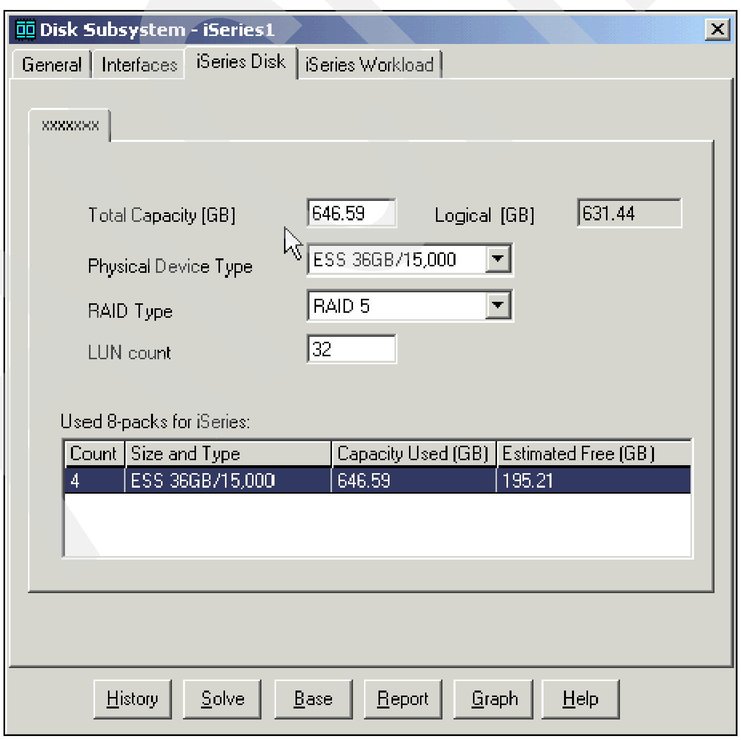

Figure 4-31   DM for Single iSeries - 32 LUNs and 1 FC adapter

After performing changes solve the configuration and produce the report. In our example the larger capacity of LUNs and less FC adapters influences disk response time, LUN utilization and host interface utilization, as shown in Figure 4-32.

```
Cache Size / Backstore Sensitivity          6.0

Advanced ESS Outputs ():
Avg. SMP Utilization:              2.1
Highest HDD Utilization:          11.1
Highest Device Adapter Util:       2.7
Upper Internal Bus Utilization:    3.1
Avg. Host Interface Utilization:  28.8
Avg. iSeries LUN Utilization:     5.4


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit%  Hit% Eff % Cnt Util%
Average         710    18.6    2.5    43   59  100     70  32    5
xxxxxxx         710    18.6    2.5    43   59  100     70  32    5

The following parameters are estimated during the modeling process
       Host interfaces           29% busy
       Disk interfaces            3% busy
```

Figure 4-32 DM for Single iSeries - 32 LUNs and 1 FC adapter - Report

You may also produce a bar graph of the new disk response time. The bar graphs made by Disk Magic for initial configuration and 2 modeled changes (18 GB DDMs and 32 LUNs) are shown in Figure 4-33.

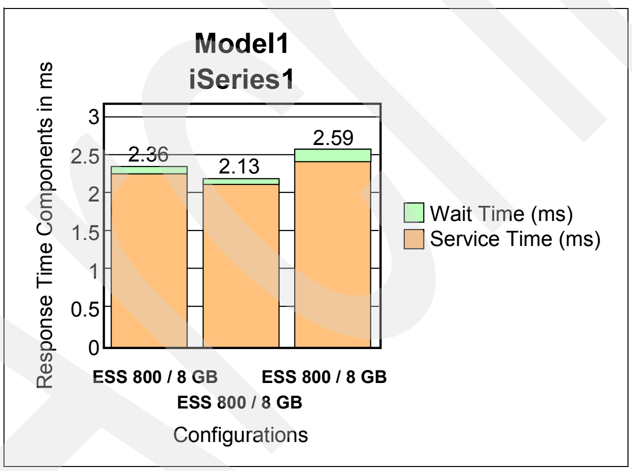

Figure 4-33 DM for Single iSeries - bar graphs of modeled changes

Modeling ESS for two iSeries

In this example we use Disk Magic for modeling ESS which will be shared by two iSeries that presently use internal disks. We will use Disk Magic to predict how disk response time changes with the following conditions:

- Both iSeries use 36 GB 15,000 RPM DDMs, and ESS ranks are shared between the two systems. Each iSeries connects through 2 FC adapters.
- One iSeries uses 18 GB 15,000 RPM DDMS and 2 FC adapters, the other uses 73 GB 15,000 RPM ddms and 1 FC adater.

Currently internal disks are mirrored, but on ESS we plan RAID protection so we will not use OS/400 mirroring.

The necessary Performance reports from each iSeries are concatenated in one.txt file.

Start Disk Magic, from the screen Welcome to Disk Magic, select iSeries PT Report file with extension TXT, as shown in Figure 4-10.

After the first PT report was successfully processed, insert the second PT report with the following method. Right-click the Disk subsystem icon and add PT reports from other iSeries system, as shown in Figure 4-34.

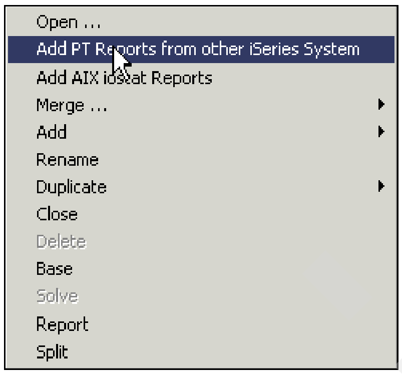

Figure 4-34 DM for two iSeries - Insert the second system

After both PT reports are successfully processed double-click the Disk Subsystem icon to bring up the Disk subsystem window, as shown in Figure 4-14.

In the Disk subsystem window observe that the iSeries disk table now contains two tables, each showing data for one iSeries, as shown in Figure 4-35. Also, the iSeries workload table contains two tables, each for one iSeries.

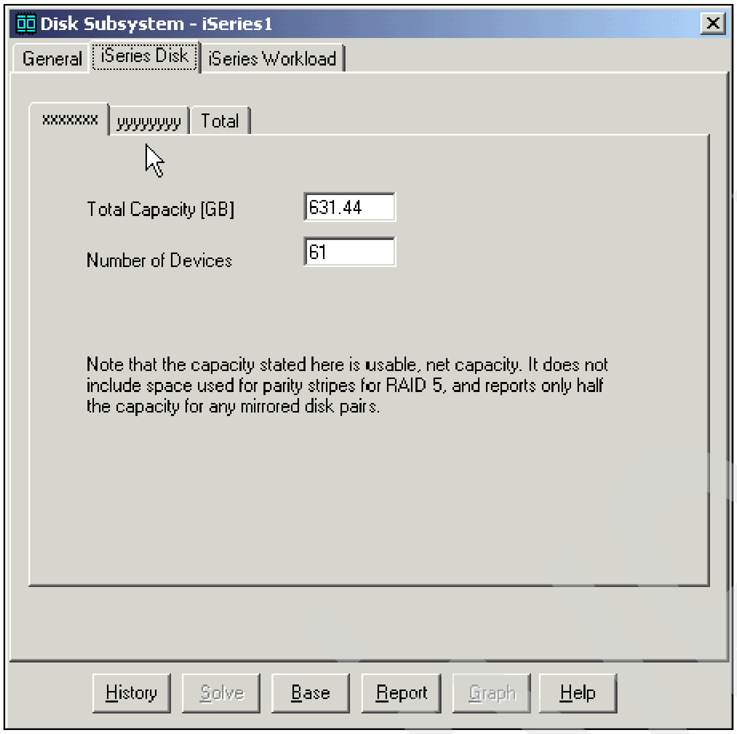

Figure 4-35 DM for two iSeries - Disk subsystem window

Save present disk subsystem as a base, by clicking Base on the Disk subsystem window.

On the General tab of the Disk subsystem window, select ESS as the Hardware type, as shown in Figure 4-36.

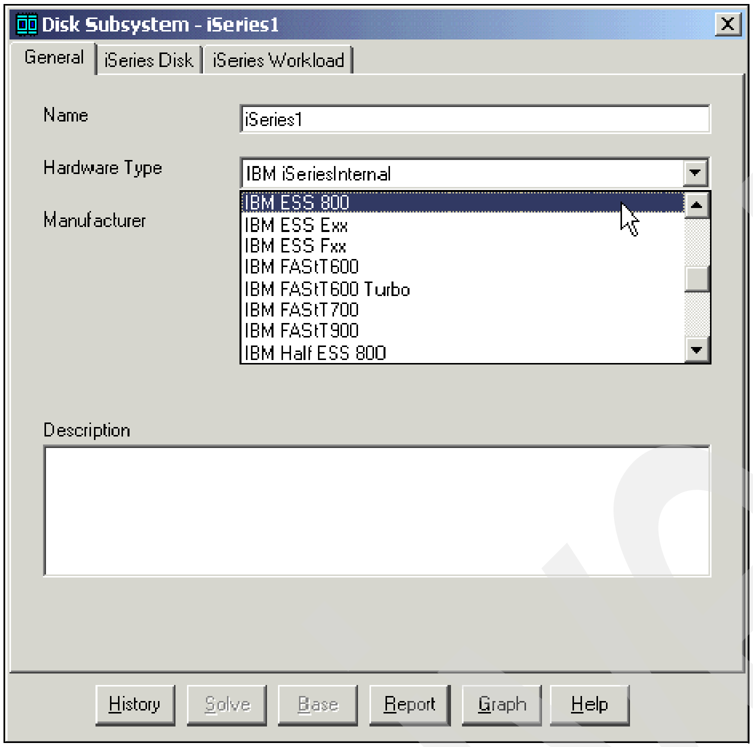

Figure 4-36 DM for two iSeries - Select ESS

If necessary, change the cache size and number of ESS adapters by clicking Hardware details and changing the relevant values.

If necessary change the number of FC adapters for each iSeries. To accomplish this, click the Interfaces tab, select the interfaces for a particular iSeries, click Edit and change the number of interfaces, as shown in Figure 4-37. For this example we specified 2 FC adapters for each iSeries.

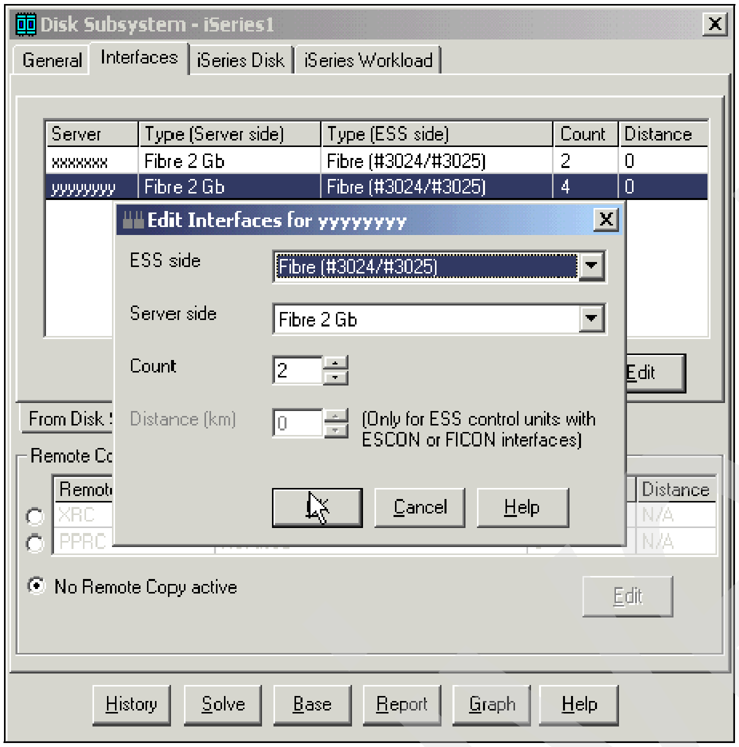

Figure 4-37 DM for two iSeries - Change the number of FC adapters

If necessary change the capacity and speed of DDMs for each iSeries. To accomplish this, click the iSeries Disk tab, click a tab for a particular iSeries and change the capacity of DDMs, as shown in Figure 4-38. For this example we used 36 GB 15,000 RPM DDMs for each iSeries.

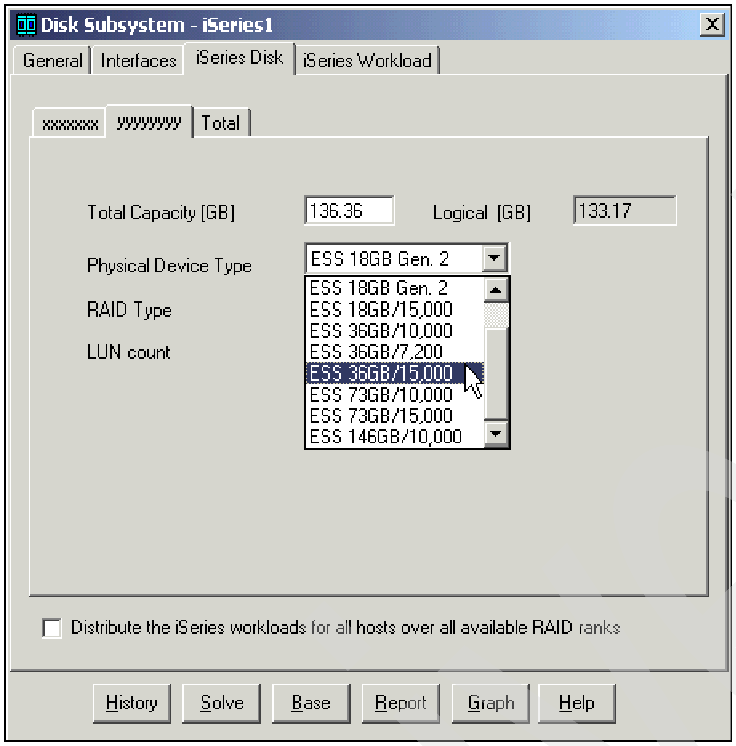

Figure 4-38 DM for two iSeries - capacity and speed of DDMs

If you plan to share the ESS ranks between the two iSeries, select the check box Distribute the iSeries workloads for all hosts over all available RAID ranks, on the iSeries disk tab. This is shown in Figure 4-39.

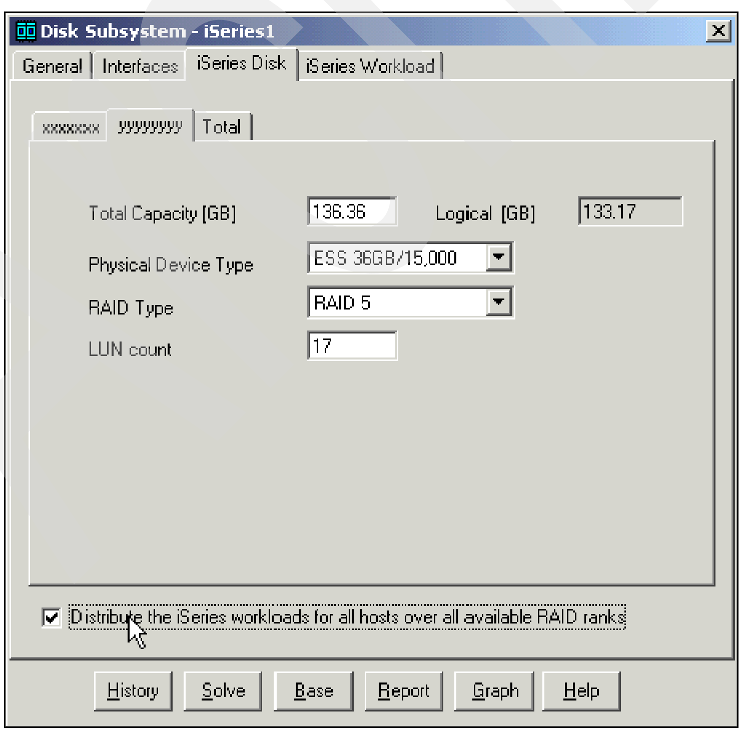

Figure 4-39 DM for two iSeries - Distribute the workload over ranks

After you have changed the ESS as desired, solve the Disk subsystem by clicking Solve, as shown in Figure 4-40.

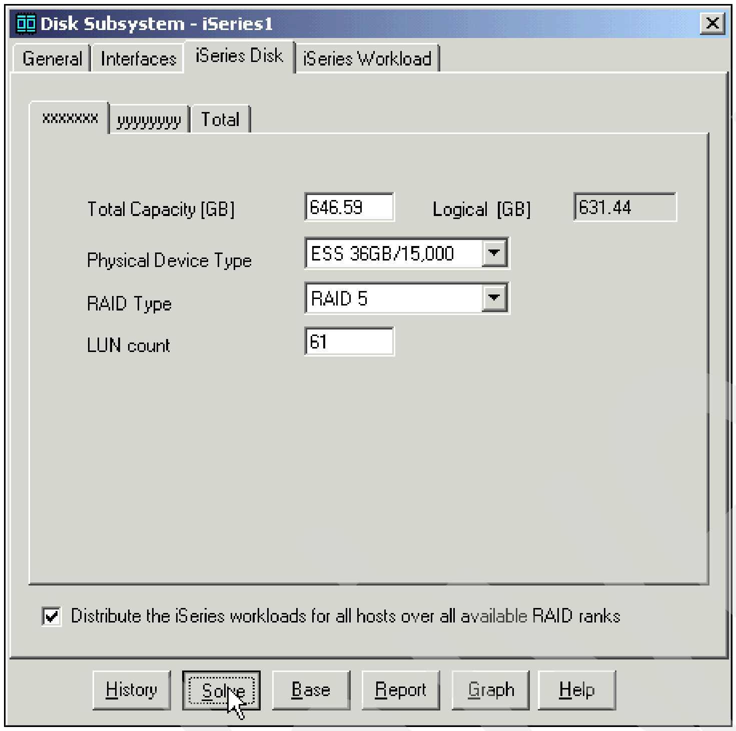

Figure 4-40 DM for two iSeries - Solve disk subsystem

If the number of LUNs is too large, so that the capacity of LUNs will be less than 8 GB, Disk Magic presents a warning and a possibility to lower the number of LUNs, as shown in Figure 4-41.

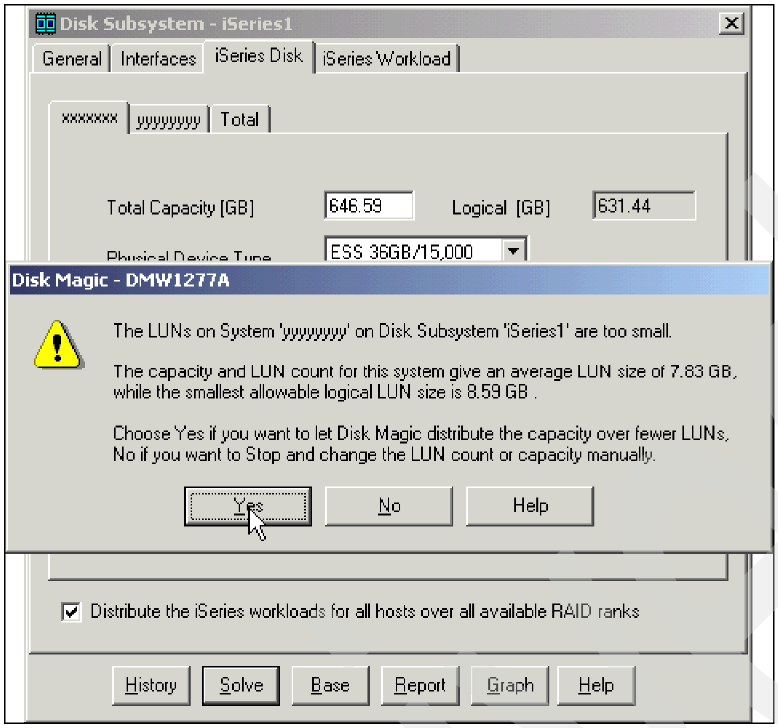

Figure 4-41 DM for two iSeries - Too many LUNs

Produce a report of modeled ESS with two iSeries by clicking Report on the Disk subsystem window. The report from our example is shown in Figure 4-42. Observe the predicted disk response time for each iSeries.

```
Total Capacity:    783.0 GB
  approximately         4 8-packs of disk type ESS 36GB/15,000

Remote copy type:  No Remote Copy Active or Remote Copy Not Support

Cache Size / Backstore Sensitivity         6.0

Advanced ESS Outputs ():
Avg. SMP Utilization:             3.9
Highest HDD Utilization:         15.1
Highest Device Adapter Util:      3.6
Upper Internal Bus Utilization:   5.0
Avg. Host Interface Utilization: 10.5
Avg. iSeries LUN Utilization:    6.1


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit% Hit% Eff %  Cnt Util%
Average        1063    16.6    2.5    50   58  100     67  77    3
xxxxxxx         710    18.6    2.4    43   57  100     70  61    3
yyyyyyyy        353    12.4    2.7    63   61  100     61  16    6

The following parameters are estimated during the modeling process
```

Figure 4-42 DM for two iSeries - Report of first modeling

If the current ESS configuration doesn't meet your expectations make further changes. In our example we changed the Capacity and speed of DDMs so that iSeries xxxxxxx uses 18 GB

15,000 RPM DDMs and iSeries yyyyyy y uses 73 GB 15,000 RPM DDMs. Changing the DDMs for iSeries xxxxxxx is shown in Figure 4-43.

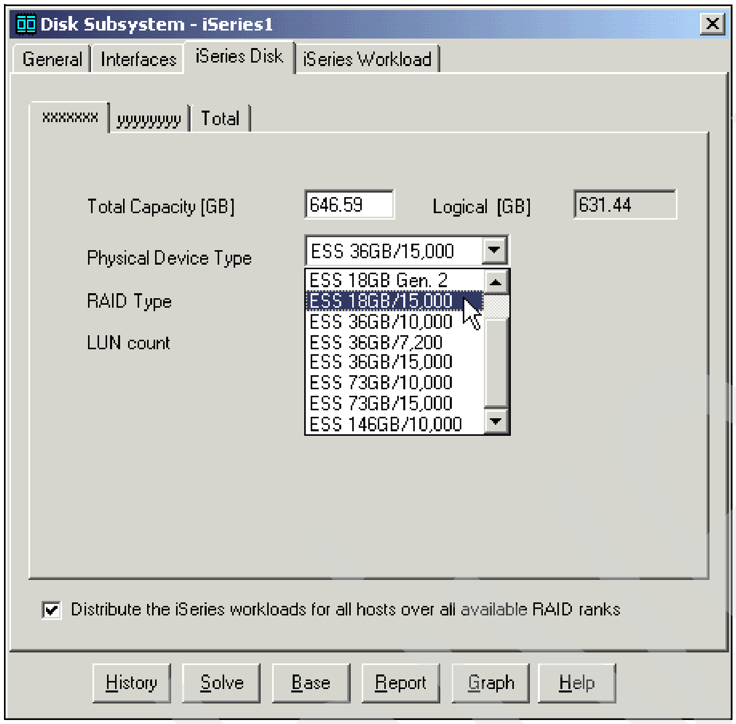

Figure 4-43 DM for two iSeries - Changing DDMs

In this example we plan to dedicate ranks of different DDMs to each iSeries, so we uncheck the checkbox Distribute the iSeries workload for all hosts over all available RAID ranks, as shown in Figure 4-44.

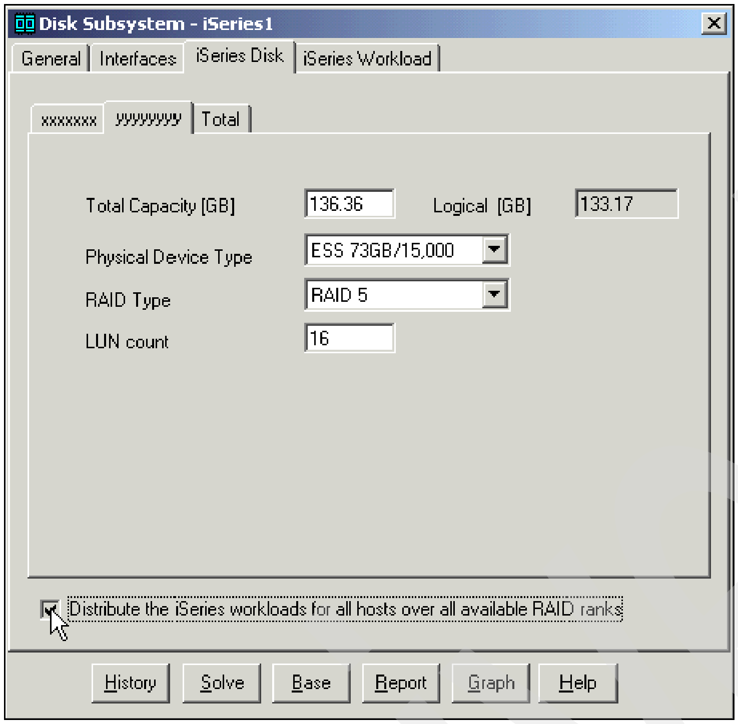

Figure 4-44 DM for two iSeries - Dedicate ESS ranks to each iSeries

As seen in Figure 4-45 there will be 7 ranks for iSeries xxxxxxx and 1 rank for iSeries yyyyyyyy.

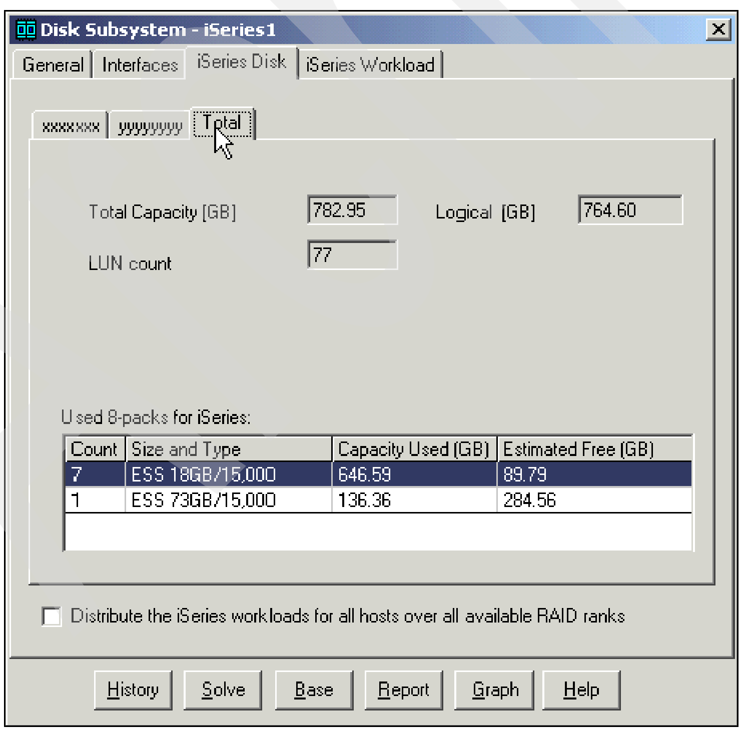

Figure 4-45 DM for two iSeries - ESS ranks for each iSeries

We also changed the number of FC adapters for iSeries yyyyyyy to 1. This is accomplished by editing the number of interfaces for iSeries yyyyyyy on the Interface tab of the Disk subsystem window, as shown in Figure 4-46. On the same tab, also change accordingly the number of interfaces from the disk subsystem.

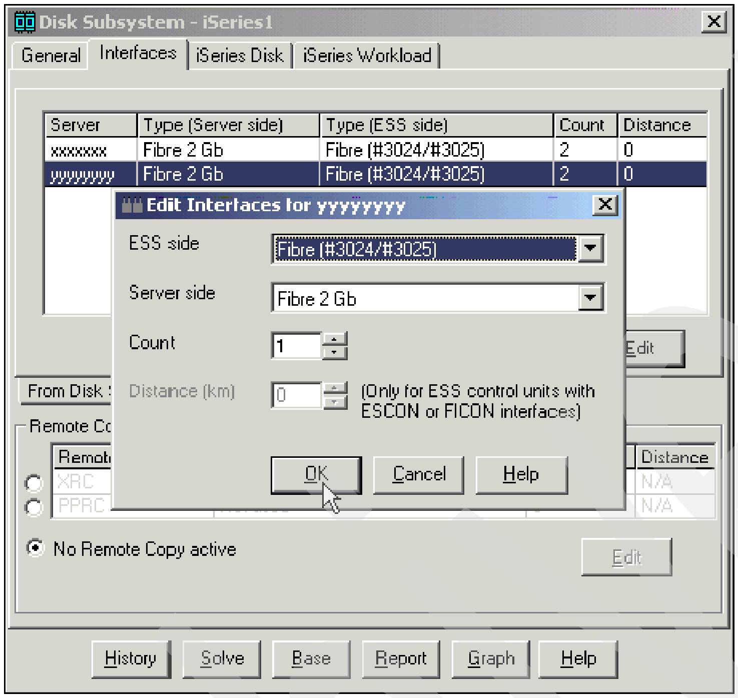

Figure 4-46 DM for two iSeries - 1 FC adapter for iSeries yyyyyyyy

After solving the changed model, obtain a report of predicted disk response time and other related values. This is accomplished by clicking Report. Disk Magic will present the report, as shown in Figure 4-47.

```
Total Capacity:    783.0 GB
  approximately        7 8-packs of disk type ESS 18GB/15,000
  approximately        1 8-pack  of disk type ESS 73GB/15,000

Remote copy type:  No Remote Copy Active or Remote Copy Not Supporte

Cache Size / Backstore Sensitivity         6.0

Advanced ESS Outputs ():
Avg. SMP Utilization:            4.0
Highest HDD Utilization:        19.5
Highest Device Adapter Util:     5.2
Upper Internal Bus Utilization:  5.0
Avg. Host Interface Utilization: 14.0
Avg. iSeries LUN Utilization:    5.9


iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit%  Hit% Eff % Cnt Util%
Average        1063   16.6    2.3    50   58   100     67  77    3
xxxxxxx         710   18.6    2.1    43   57   100     70  61    2
yyyyyyyy        353   12.4    2.7    63   61   100     61  16    6
```

Figure 4-47 DM for two iSeries - Different DDMs and 1 FC adapter for iSeries yyyyyyy - Report

Observe that in our example disk response time for iSeries xxxxxxx significantly decreased, while response time for iSeries yyyyyyyy is the same as before the changes were made. The reason for this is probably very low I/O activity of iSeries yyyyyyyy.

Modeling ESS for iSeries and other servers

In this example we use Disk Magic to model ESS for iSeries and two servers with Windows NT. PT report for iSeries is available, but for Windows servers we have only the following information:

- Server 1: 300 GB capacity, 350 IO/second
- Server 2: 700 GB capacity, 450 IO/second

Insert into Disk Magic the capacity and I/O rate of both Windows servers. This is accomplished in the following way:

On Welcome to Disk Magic window, select General Project and specify the number of open systems, as shown in Figure 4-48.

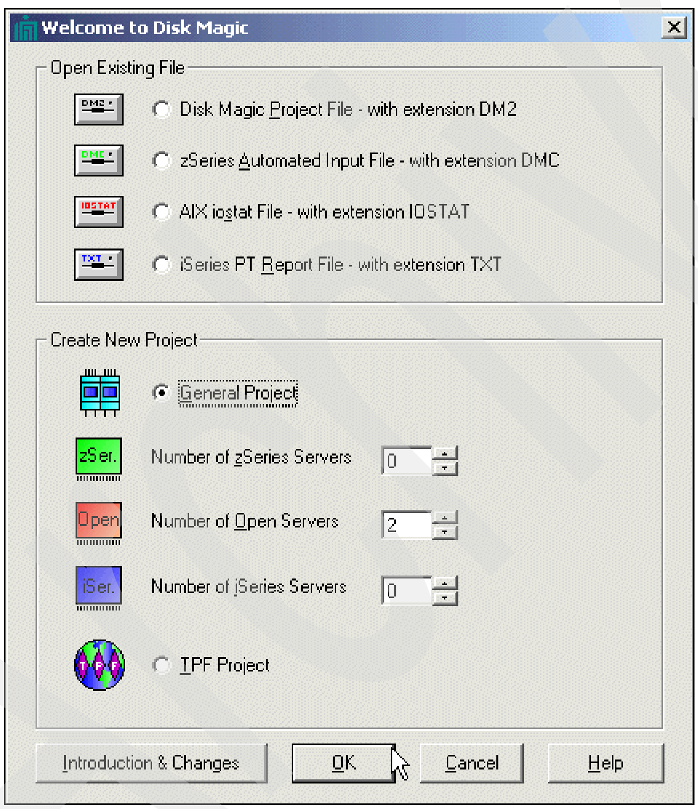

Figure 4-48 DM for iSeries and Windows - insert data from Windows

When you click OK on Welcome to Disk Magic window, the Disk Subsystem window opens. Click the Open Disk tab and insert capacity for both Windows servers, as shown in Figure 4-49.

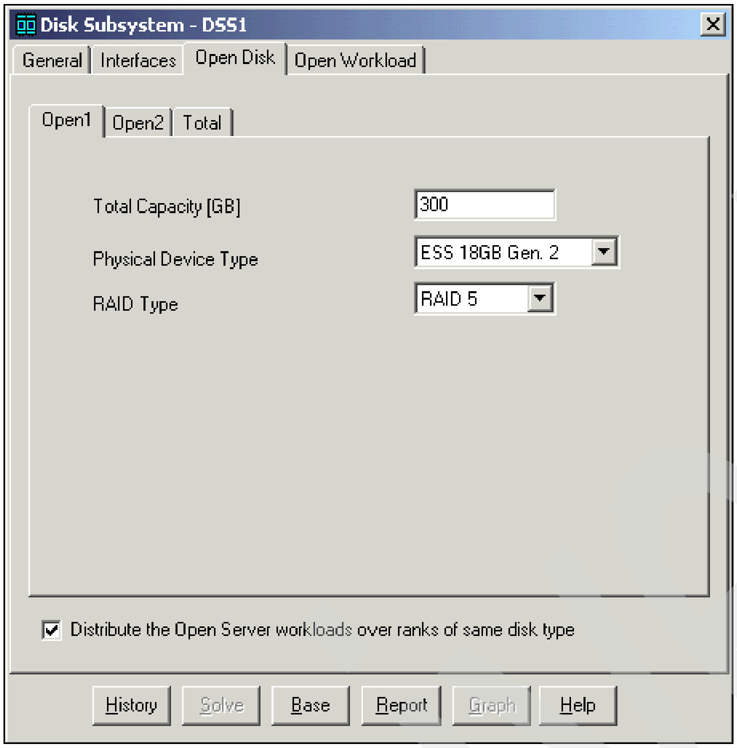

Figure 4-49 DM for iSeries and Windows - insert Windows capacity

Click the Open Workload tab of the Disk Subsystem window and insert the I/O rate for both Windows systems, as shown in Figure 4-50. Leave other values as defaults. At this point you don't need to change the features of ESS as desired, because you will change them later.

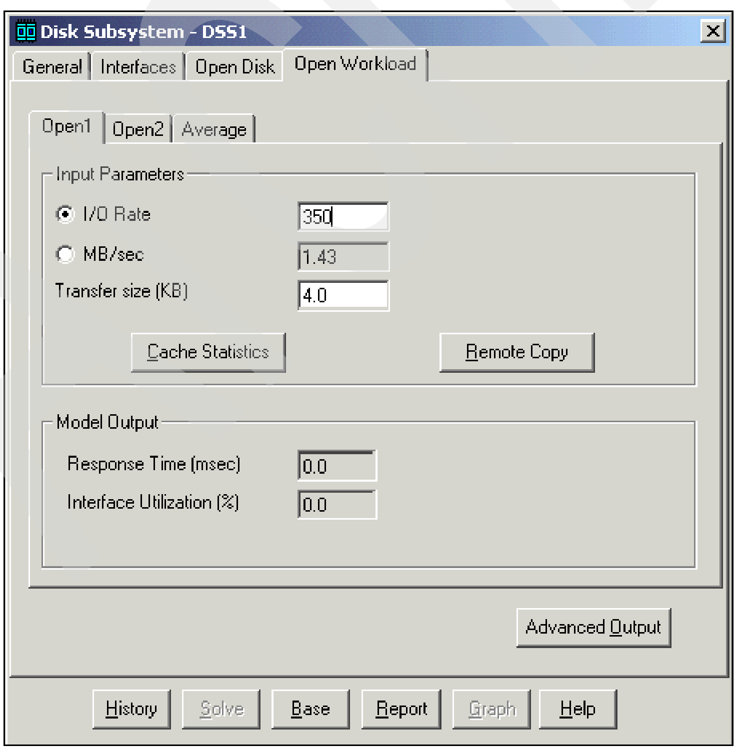

Figure 4-50 DM for iSeries and Windows - insert Windows I/O rate

To add iSeries workload, right-click the Disk Subsystem icon and select Add PT Reports from other iSeries System, as shown in Figure 4-51.

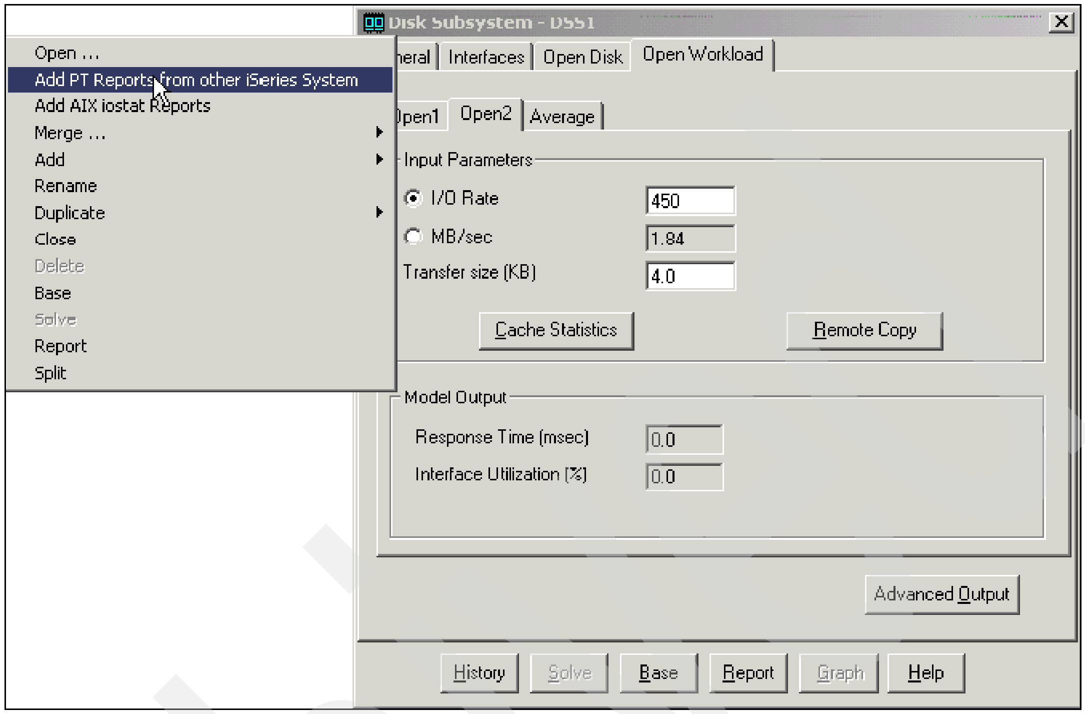

Figure 4-51 DM for iSeries and Windows - add iSeries workload

You are prompted for iSeries PT report, choose the one you want to process, as shown in Figure 4-52.

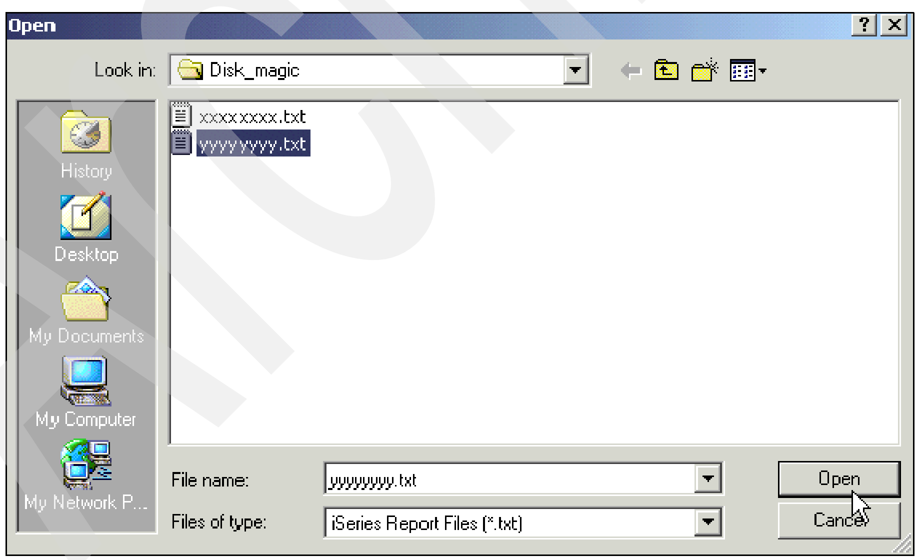

Figure 4-52 DM for iSeries and Windows - choose iSeries PT report

After the iSeries PT report is successfully processed it will result in a separate Disk subsystem.

To merge both iSeries and Windows workloads, click the Model1 icon, select Add Disk Subsystem from pull-down menu and select Merge Target DSS from pull-down menu. This way you create an intermediate disk subsystem which will be used for merging both iSeries and Windows workloads. This is shown in Figure 4-53. This will bring up the Merge Target window.

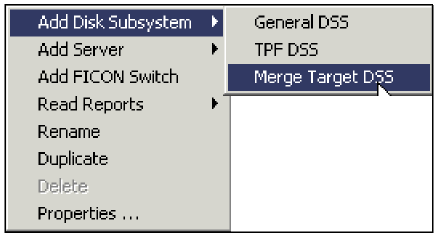

Figure 4-53 DM for iSeries and Windows - Add Merge Target

Right-click the Disk Subsystem icon for Windows, select Merge from pull-down menu and select Add to Merge Source Collection from the pull-down menu. This is shown in Figure 4-54.

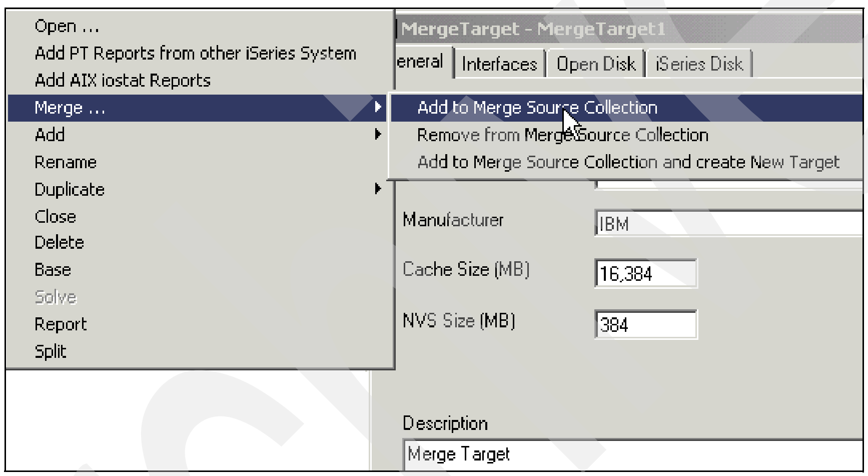

Figure 4-54 DM for iSeries and Windows - Add Windows workload to Source merge

Add the iSeries workload to the Source Merge collection the same way as the Windows workload.

After both workloads are added to Source Merge Collection, click Merge on the Merge Target window, as shown in Figure 4-55.

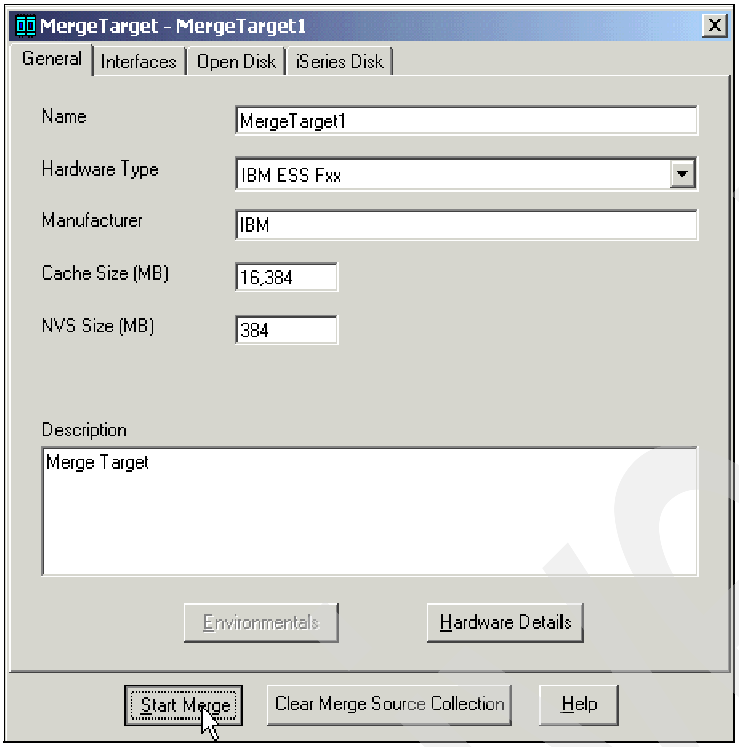

Figure 4-55 DM for iSeries and Windows - merge workloads

After workloads are successfully merged, Disk Magic will create a new Disk Subsystem - MergeResult1 which contains iSeries and both Windows servers. Disk Magic will open the Disk Subsystem window for MergeResult1.

In MergeResult1 Disk Subsystem, select the desired model of ESS, as shown in Figure 4-56.

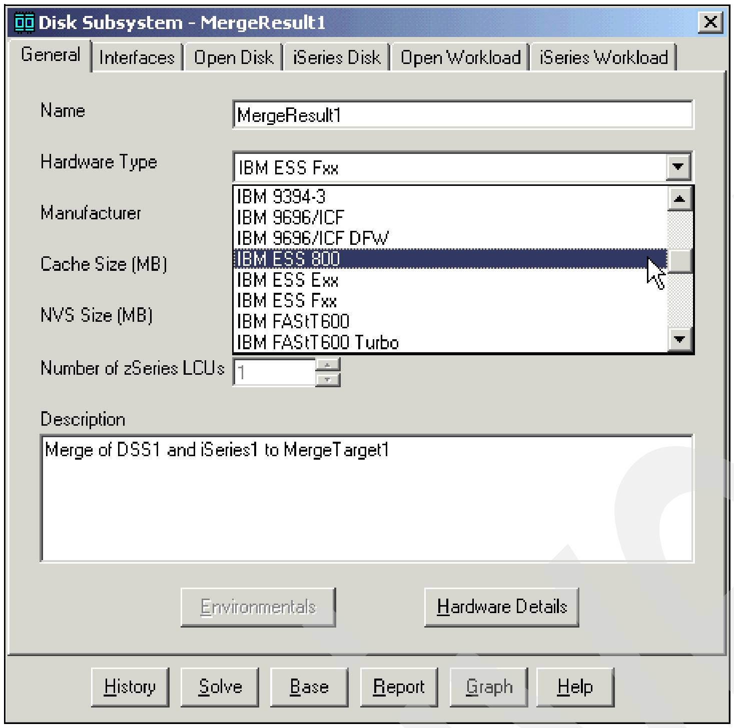

Figure 4-56 DM for iSeries and Windows - Select ESS model

If necessary change the number of adapters and cache size for ESS. To accomplish this click Hardware Details and change the relevant values as desired.

If necessary change the number of FC adapters for each system. To accomplish this click the Interfaces tab and edit the number of interfaces. Also specify the appropriate number of interfaces from Disk Subsystem.

For this example we specified 1 FC adapter for the first Windows server, 2 FC adapters for the second Windows servers and 2 FC adapters for iSeries. This is shown in Figure 4-57.

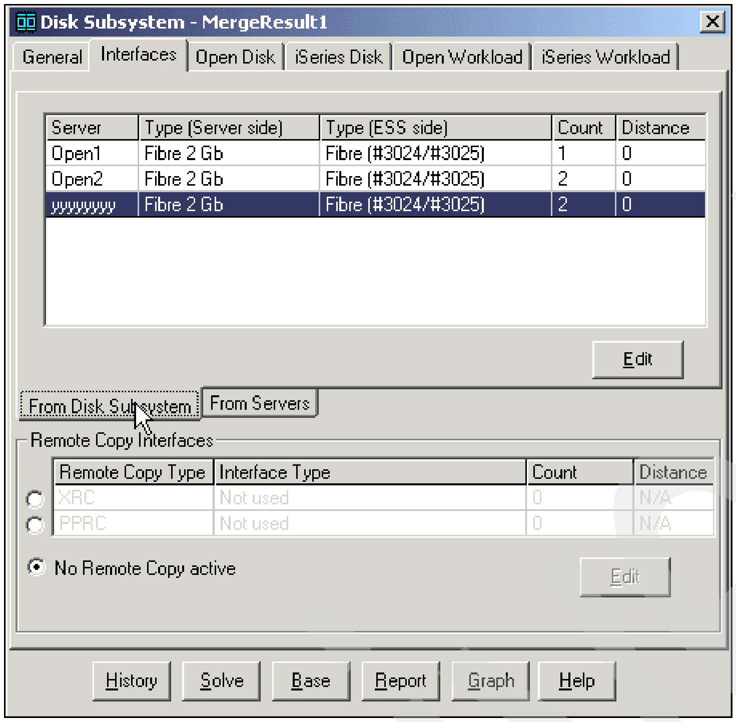

Figure 4-57   DM for iSeries and Windows - FC adapters

Change capacity and speed of DDMs for Windows servers as appropriate. To accomplish this click the Open Disk tab and change capacity and speed of DDMs as desired.

Change capacity and speed of DDMs for iSeries as appropriate. To accomplish this click the iSeries Disk tab and change capacity and speed of DDMs as desired.

For this example we use 73 GB 15,000 RPM DDMs for Windows servers and 36 GB 15,000 RPM DDM for iSeries. Figure 4-58 shows the changing capacity and speed of DDMs for Windows server 1. Observe that the checkbox Distribute the Open Server workloads over ranks of same disk type is checked by default.

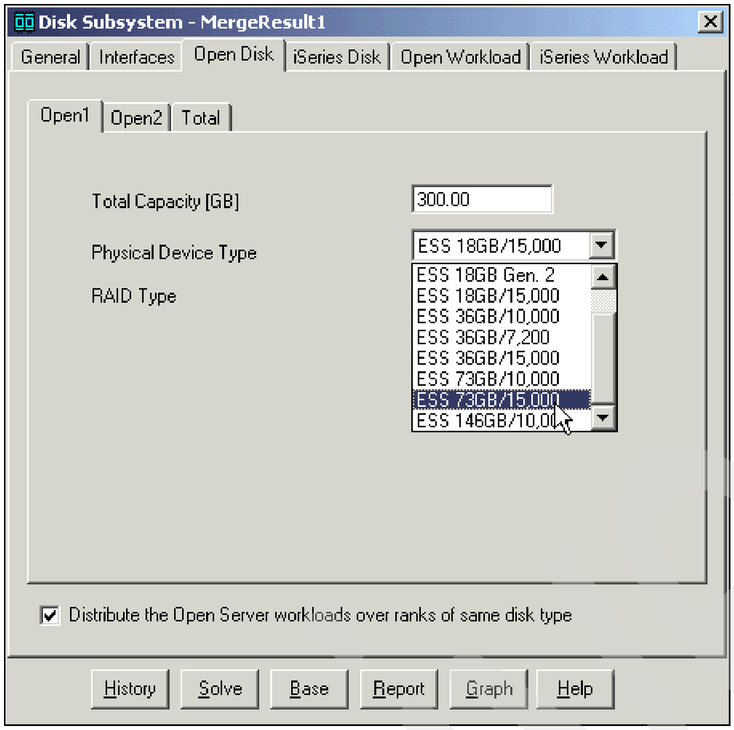

Figure 4-58 DM for iSeries and Windows - capacity and speed of DDMs for Windows server 1

After you have made desired changes, click Solve on the Disk Subsystem MergeResult1 window to solve the current configuration. Click Report on the same window to produce the report of present configuration. Disk response times for both Windows servers and iSeries are shown in part of the report in Figure 4-59.

```
Open Server     I/O  Transfer  Resp  Read Read Write Write C
                Rate Size (KB) Time  Perc Hit%  Hit% Eff%  U
Average          800      4.0   3.7    70   44   100    0
Open1            350      4.0   3.7    70   44   100    0
Open2            450      4.0   3.7    70   44   100    0

iSeries Server I/O  Transfer  Resp  Read Read Write Write L
                Rate Size (KB) Time  Perc Hit%  Hit% Eff % C
Average          353     12.4   2.4    63   70   100      61
yyyyyyyy         353     12.4   2.4    63   70   100      61
```

Figure 4-59 DM for iSeries and Windows - disk response time - Report

If necessary make other changes to the proposed configuration. For this example we changed the capacity of DDMs for iSeries to 73 GB. Accordingly we changed the capacity and the number of LUNs for iSeries. This is shown in Figure 4-60.

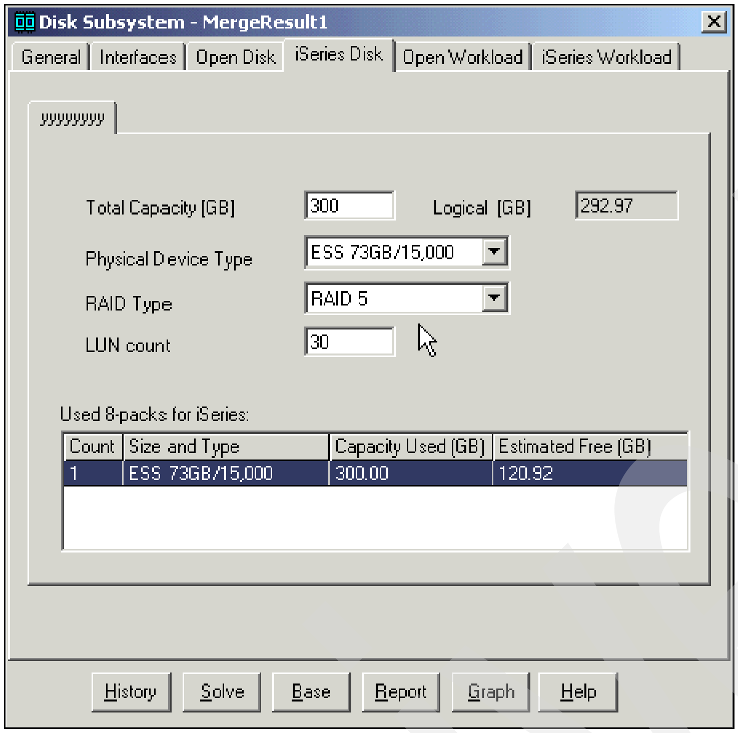

Figure 4-60 DM for iSeries and Windows - 73 GB DDMs for iSeries

After making desired changes solve the configuration and produce the report. Part of the report from our example is shown in Figure 4-61. Observe the same disk response time for iSeries as was on 36 GB DDMs and higher disk response time for Windows servers.

```
Open Server    I/O  Transfer  Resp  Read Read Write Write Chan
               Rate Size (KB) Time  Perc Hit%  Hit% Eff%  Util%
Average         800      4.0   3.9    70   42   100    0      2
Open1           350      4.0   3.9    70   42   100    0      2
Open2           450      4.0   3.9    70   42   100    0      1

iSeries Server I/O  Transfer  Resp  Read Read Write Write LUN LUN
               Rate Size (KB) Time  Perc Hit%  Hit% Eff % Cnt Uti
Average         353     12.4   2.4    63   66   100     61  30
yyyyyyyy        353     12.4   2.4    63   66   100     61  30
```

Figure 4-61 DM for iSeries and Windows - 73 GB DDMs for iSeries - Report

5

Adding ESS storage to the iSeries server

This chapter explains how to configure the iSeries for adding Enterprise Storage Server (ESS) storage.

The ESS can attach to the iSeries by either Small Computer Systems Interface (SCSI) for V4R5 or earlier via a migration tower, or fibre, for V5R1 only. The steps to add the external storage to the iSeries are the same, regardless of the adapter. With a SCSI adapter, the logical unit numbers (LUNs) report to the iSeries as device type 9337 (Table 5-1). For a Fibre Channel adapter, the LUNs report as 2105 device types (Table 5-2).

Table 5-1 SCSI attached drives

Size (GB)	Type	Protected	Unprotected
4.190	9337	48C	48A
8.589	9337	59C	59A
17.548	9337	5AC	5AA
35.165	9337	5CC	5CA
36.003	9337	5BC	5BA

Table 5-2 Fibre attached drives

Size (GB)	Type	Protected	Unprotected
8.589	2105	A01	A81
17.548	2105	A02	A82
35.165	2105	A05	A85
36.003	2105	A03	A83
70.564	2105	A04	A84

The steps defined here assume that the connection to the ESS is performed by the Customer Engineer (CE) or service support representative (SSR). The CE or SSR should be trained on the ESS. The CE or SSR performs the following tasks:

1. Installs the ESS.
2. Installs the Disk Drive Modules (DDMs) into the ESS.
3. Configures Copy Services on the ESS if applicable.

4. Uses the worksheets to configure:
   - Hosts to the ESS
   - Host adapters on the ESS
   - Disk groups on the ESS
   - LUNs or volumes on the ESS

You must have completed the planning worksheets from Chapter 4, “Planning for external disk” on page 63, before adding the disk drives to the iSeries. You will add disk units to an ASP and then start device parity protection on those drives.

If you plan to implement remote load source mirroring, see Chapter 11, “Load Source Mirroring in the ESS” on page 303, before you proceed with these steps. If you plan to use an advanced feature such as PPRC, read Chapter 12, “Peer-to-Peer Remote Copy” on page 327, before you proceed with these steps.

1. IPL the iSeries in manual mode to Dedicated Service Tools (DST).
2. Select option 3 (Use Dedicated Service Tools) as shown in Figure 5-1.

```
                         IPL or Install the System

Select one of the following:

     1. Perform an IPL
     2. Install the operating systems
     3. Use Dedicated Service Tools (DST)
     4. Perform automatic installation of the operating systems
     5. Save Licensed Internal Code



Selection  3


F3=Exit        F10=Command entry        F12=Cancel
```

Figure 5-1 Selecting to use DST from the IPL or Install the System display

3. Sign on to DST with the Service tools QSECOFR user ID and password as shown in Figure 5-2.

```
                     Dedicated Service Tools (DST) Sign On
> > > > > > > > > > > > > > > > > > > > > > > > > > >System: AST32
Type Choices, press enter

             DST User . . . . . . . . . . . .
             DST Password . . . . . . . . . .





F3=Exit      F12=Cancel

Selection
```

Figure 5-2 DST sign-on display

4. Select option 4 (Work with disk units) from the DST menu.

```
                Use Dedicated Service Tools (DST)

 Select one of the following:

>1. Perform an IPL
>2. Install the operating system
>3. Work with Licensed Internal Code
>4. Work with disk units
>5. Work with DST environment
>6. Select DST console mode
>7. Start a service tool
>8. Perform automatic installation of the operating system
>9. Work with save storage and restore storage
10. Work with remote service support
11. Work with system partitions


   Selection
      4

   F3=Exit      F12=Cancel
```

Figure 5-3 DST main menu

5. Select option 2 (Work with disk configuration) as shown in Figure 5-4.

```
                              Work with Disk Units

 Select one of the following:


      1. Work with disk configuration
      2. Work with disk unit recovery






 Selection 2


 F3=Exit         F12=Cancel
```

Figure 5-4 Work with Disk Units display

At this point, you have the option to add disks to an ASP. We recommend that you start parity protection before you add disks to ASPs.

6. On the Work with Disk Configuration display, select option 5 (Work with device parity protection).

```
                        Work with Disk Configuration

 Select one of the following:

>>>>> 1. Display disk configuration
      2. Work with ASP threshold
      3. Work with ASP configuration
      4. Work with mirrored protection
      5. Work with device parity protection
      6. Work with disk compression




 Selection
      5

 F3=Exit         F12=Cancel
```

Figure 5-5 Work with Disk Configuration display

7. On the Work with Device Parity Protection display, select option 2 (Start device parity protection) as shown in Figure 5-6.

```
                  Work with Device Parity Protection

  Select one of the following:

        1. Display device parity status
        2. Start device parity protection
        3. Include unit in device parity protection
        4. Exclude unit from device parity protection







  Selection
     2

  F3=Exit      F12=Cancel
```

Figure 5-6 Work with Device Parity Protection display

8. On the Start Device Parity Protection display (Figure 5-7), where you can you can select disks to start parity protection. Select the units for RAID-5 protection. Press Enter to begin the process.

> **Note:** With Fibre Channel attachment, the display is identical, but the device type is 2105.

```
                          Start Device Parity Protection

Select the subsystems to start device parity protection.

Type choice, press Enter.
  1=Start device parity protection

         Parity  Serial                    Resource
Option     Set   Number      Type   Model  Name
             3   10-8275001  9337   5AC    DC01
             4   10-8275001  9337   5AC    DC01
             5   10-8275001  6751   001    DC01



F3=Exit         F12=Cancel
```

Figure 5-7 Start Device Parity Protection display

9. The system confirms the starting of device parity protection on the selected drives (Figure 5-8). Press Enter to confirm the start of parity protection.

```
                    Confirm Starting Device Parity Protection

 During the preparation for starting device parity protection, data
 will be moved from parts of some disk units. This may take several
 minutes for each subsystem selected.

 Press Enter to continue.
 Press F12=Cancel to return and change your choice.

         Parity           Serial                     Resource
 Option     Set  ASP  Unit  Number      Type  Model  Name
    1        3              10-8275001  9337    5AC    DC01
    1        3    *     *   83-3005466  9337    5AC    DD003
    1        3    *     *   33-00188    9337    5AC    DD017
    1        3    *     *   83-13256    9337    5AC    DD018
    1        3    *     *   83-3007795  9337    5AC    DD001



 F12=Cancel
```

Figure 5-8 Confirm Starting Device Parity Protection display

10. The system presents a display that shows the status of the operation. You can wait until the next display returns or press F16 to return to DST to wait. When the starting device parity protection operation finishes, you return to the Work with Device Parity Protection display.

11. Press F12 to return to the Work with Disk Configuration display.

12. Select option 3 (Work with ASP configuration) as shown in Figure 5-9.

```
                          Work with Disk Configuration

 Select one of the following:

      1. Display disk configuration
      2. Work with ASP threshold
      3. Work with ASP configuration
      4. Work with mirrored protection
      5. Work with device parity protection
      6. Work with disk compression




 Selection
      3

 F3=Exit         F12=Cancel
```

Figure 5-9 Work with Disk Configuration display

13. On the Work with ASP Configuration display, select option 3 (Add units to ASP) as shown in Figure 5-10.

```
                          Work with ASP Configuration

 Select one of the following:

      1. Display disk configuration capacity
      2. Delete user ASP
      3. Add units to ASP
      4. Delete ASP data
      5. Change ASP threshhold
      6. Move units from one ASP to another
      7. Remove units from configuration
      8. Add units to ASPs and balance data




 Selection
        3

 F3=Exit         F12=Cancel
```

Figure 5-10 Work with ASP Configuration display

14.On the next display (Figure 5-11), specify the ASP number next to the desired units and press Enter.

```
                        Specify ASPs to Add Units to

 Specify the ASP to add each unit to.

 Specify  Serial                           Resource
   ASP    Number       Type  Model  Capacity  Name
          68-0142152  6607   050      4194  DD011
          68-0118035  6607   050      4194  DD013
          68-26409    6607   050      4194  DD002
          68-09544    6717   050      8589  DD016
          68-40090    6607   050      4194  DD005
          68-57776    6607   050      4194  DD019
          68-09179    6717   050      8589  DD007
          68-56547    6607   050      4194  DD004
          68-57706    6607   050      4194  DD009
          68-53774    6607   050      4194  DD008
          83-3005466  6714   050     17548  DD003
          68-0118156  6607   050      4194  DD010
          83-3007795  6714   050     17548  DD001
                                                                       More...

 F3=Exit     F5=Refresh     F11=Display disk configuration capacity
 F12=Cancel
```

Figure 5-11 Adding units to ASP

15. The Confirm Add of Units display (Figure 5-12) appears for review. If everything is correct, press Enter to continue.

```
                              Confirm Add Units

 Add will take several minutes for each unit.  The system will
 have the displayed protection after the unit(s) are added.

 Press Enter to confirm your choice for Add units.
 Press F9=Capacity Information to display the resulting capacity.
 Press F12=Cancel to return and change your choice.

             Serial                    Resource
 ASP  Unit   Number      Type Model    Name         Protection
   1                                                Mirrored
 >>>>>> >1   68-09011    6717  050     D023        >Bus
          1  75-1044000  9337  59A     DD037        Bus
          2  68-0142152  6607  050     DD011        Controller
          2  68-0118035  6607  050     DD013        Controller
          3  75-2044000  9337  59C     DD038        Device Parity
          4  75-3044000  9337  5AC     DD039        Device Parity
          5  75-4044000  9337  5AC     DD040        Device Parity
          6  75-5044000  9337  5AC     DD041        Device Parity
                                                                  More...

 F9=Resulting Capacity           F12=Cancel
```

Figure 5-12 Confirm Add Units display

Depending on the number of units you are adding, this step could take some time. When it completes, display your disk configuration to verify the capacity and data protection. If everything is correct, finish running an IPL on the system.

6

Creating iSeries storage in ESS

The IBM StorWatch Enterprise Storage Server Specialist (ESS Specialist) is the Web user interface that is included with your IBM Enterprise Storage Server (ESS). You can use ESS Specialist to view the current status of the ESS and to modify the ESS configuration.

This section explains how to:

- Use task maps (quick guides to configuration tasks)
- Use the ESS Specialist interface
- Access the ESS Specialist interface
- Get started using ESS Specialist
- Check the ESS status
- Use the Storage Allocation panel
- Configure open system (iSeries) storage
- Sign off ESS Specialist and close all browser windows

For more information, refer to *IBM TotalStorage Enterprise Storage Server Web Interface User's Guide*, SC26-7346.

6.1 Using task maps

Table 6-1 lists various configuration tasks for Small Computer Systems Interface (SCSI) or Fibre Channel-attached hosts and the sections in this chapter that explain them. If you have previously used ESS Specialist, consult the sections in this table to go directly to task you need.

Table 6-1 ESS Specialist task map for SCSI or Fibre Channel-attached hosts

Configuration task	Related section
Add, modify, or delete host systems	Section 6.5.1, "Defining host systems attached to the storage server" on page 145
Configure an open system port	Section 6.6, "Configuring host adapter ports" on page 151
Configure disk group	Section 6.7.2, "Configuring fixed block disk groups" on page 155
Add volumes	Section 6.7.4, "Adding volumes" on page 157
Assign or modify volumes	Section 6.8.1, "Modifying volume assignments" on page 161

6.1.1 Using the ESS Specialist user interface

The ESS Specialist user interface consists of four major components:

- **Information frame**: The long bar at the top of the Web page that identifies the interface. It contains the global help button and the IBM logo on the right side of the bar.

  Click the question mark (**?**) at any time to bring up a new browser window that displays the Help system. See 6.1.3, "Using the help system" on page 130, for a description of the Help system.
- **Navigation frame**: The column of buttons on the extreme left of the interface. These buttons indicate the categories of the ESS status and configuration that you can view or modify with ESS Specialist. You can click any of these buttons at any time. Then, the work frame switches to display the information associated with the particular button that you clicked.
- **Work frame**: The main area in the center of the browser window. It contains the actual information being displayed by the interface, and it allows you to perform configuration tasks. Every panel displayed in the work frame contains a page-help icon in the upper right-hand corner (a page icon with a question mark (?)).

  When you click this icon, it links to the Help system and displays the information specific to the current work frame panel. You may have to switch from the main browser window to the help browser window after you click the Help icon to actually see the displayed help information.

  In addition to the Help icon, certain panels displayed in the work frame also contain buttons that allow you to manipulate the interface. In general, these additional buttons can be one of three types:

  – Buttons near the top of the work frame manipulate the view of the information currently being displayed. These types of buttons generally do not link to other panels, but display an alternate view of the same information in the current panel. This may take the form of a tabular view versus a graphical view (for example, on the Storage Allocation panel). Or it can be as simple as modifying the sorting order in one of the tables (for example, on the Modify Fixed Blocked Volume Assignments panel).

– Buttons near the bottom of the work frame generally link to some other panel in the same configuration category as the current panel. For example, the Problem Notification panel links to the Modify Emails, Modify Pagers, and Modify SNMP panels.

– The *Perform* Configuration Update and *Cancel Configuration Update* buttons, which also appear at the bottom of the configuration panels, allow you to apply the indicated configuration changes to the storage server or to cancel out of the current configuration panel.

  The majority of the configuration panels eventually return to their individual parent panels after you click the Perform or Cancel buttons.

- **Message frame**: The message frame is the long bar at the extreme bottom of the interface. ESS Specialist uses this frame to display messages to the user when necessary. The two most commonly displayed messages are:

  – `– Message 2108,Retrieving the machine configuration data`

    The storage server displays this message when it obtains the latest configuration data. This message disappears once the latest data is received from the machine.

  – `– Message 1533,Downlevel Data...please wait a few minutes and try again`

    This message indicates that the absolute latest configuration data is not currently available and that the information displayed on the panel does *not* represent the latest configuration.

    This message does not disappear until you manually reload the panel. You should wait a minute or two, and then click the last-clicked button again or click **Data Refresh** to reload the configuration data.

6.1.2 Using ESS Specialist tables

In several of the ESS Specialist windows, information is presented in the form of tables, with a fixed number of columns and a variable number of rows. (You can print the ESS Specialist tables from a Web browser.) This section describes some of the common behavior of tables in the ESS Specialist user interface.

Sorting on table columns

Tables that are sortable have a set of drop-down list boxes with sort choices displayed above each of the table columns. To sort these tables, follow these steps:

1. In the drop-down list boxes above each column in the table, there are several choices:

   – No sort...
   – Sort First...
   – Sort Second...
   – Sort Third...

   Select one of the choices for those columns that you want to sort by. When you select a sort choice for one of the columns, it is no longer shown (and therefore cannot be selected) for the other columns.

2. Click **Perform Sort** near the top of the window to rearrange the table according to the sort options that you selected.

   The sort order is always in ascending order. The rows in the table are ordered so that:

   – The row with the smallest data value in the column selected with the Sort First... choice is the first row in the table

   – The row with the largest data value in the column selected with the Sort First... choice is the last row in the table.

If one or more table rows have the same data value in that column, the order of those rows is determined by their respective data values in the column selected with the Sort Second... choice. If one or more table rows have the same data value in that column, then the order is determined by the column selected with the Sort Third... choice. The order of any table rows that contain the same data values in all three Sort... columns is undetermined.

Selecting multiple table rows

To select multiple contiguous rows, follow these steps:

1. Click the first row in the sequence.
2. Scroll down to the last row in the sequence.
3. Hold down the Shift key.
4. Click the last row.
5. Release the Shift key. All rows between the first and last rows inclusive are selected.

To select multiple non-contiguous rows, follow these steps:

1. Click the first row in the sequence.
2. Hold down the Control (Ctrl) key.
3. Click each desired row.
4. Release the Control key. Each of the individually clicked rows are selected.

6.1.3 Using the help system

The ESS Specialist provides three types of help:

- **Content help**: Includes a list of topics. When you click the topic of interest, the ESS Specialist displays information about the topic.
- **Glossary help**: Includes a list of the terms that are used in the Web interface. Click a term to see its definition.
- **Task help**: Includes a list of tasks that you perform to configure the ESS. Click the task you want to perform to see instructions for it. Task help also includes scenarios that show examples of how to:
  - Prepare for configuration
  - Check the status of the ESS
  - Add a SCSI or Fibre Channel-attached host system
  - Add storage
  - Create a shared volume
  - Re-assign storage from one host system to another
  - Detach a volume from a host system
  - Configure using a non-RAID format
  - Change storage format
  - Define S/390® storage
  - Use Parallel Access Volumes (PAV) for S/390 hosts
  - Reformat a RAID-5 array
  - Use standard and custom volumes

6.2 Using Web browsers

IBM includes a Web browser as part of the IBM Enterprise Storage Server Network (ESSNet). You can access ESS Specialist through the ESSNet workstation.

You can access ESS Specialist from your own workstation if you connect the ESSNet into your intranet using one of the ports on the ESSNet hub. Access ESS Specialist using a browser that is fully enabled for Java 1.1, such as Netscape Navigator or Microsoft® Internet Explorer.

Note: If you experience problems using a browser at a certain level, upgrade to the highest-available level.

If you do not choose to attach ESSNet into your intranet, then you can access ESS Specialist from the browser on the ESSNet workstation (an IBM personal computer).

You can use the dotted decimal Internet Protocol (IP) address or hostname alias that you assigned to each ESS cluster controller as the Web address to access the user interface from the browser. Each cluster provides the same information, so you need to connect to only one at any given time. IBM provides two cluster controllers for fault tolerance.

6.2.1 Accessing the user interface

You can connect to the ESS Specialist interface from the ESSNet workstation, or from any workstation on the local intranet, if the ESSNet is connected to the local intranet.

Start the Netscape Navigator or Microsoft Internet Explorer browser on the workstation. Next type the URL of one of the cluster controllers. Then the ESS Specialist Welcome page (Figure 6-1) is displayed.

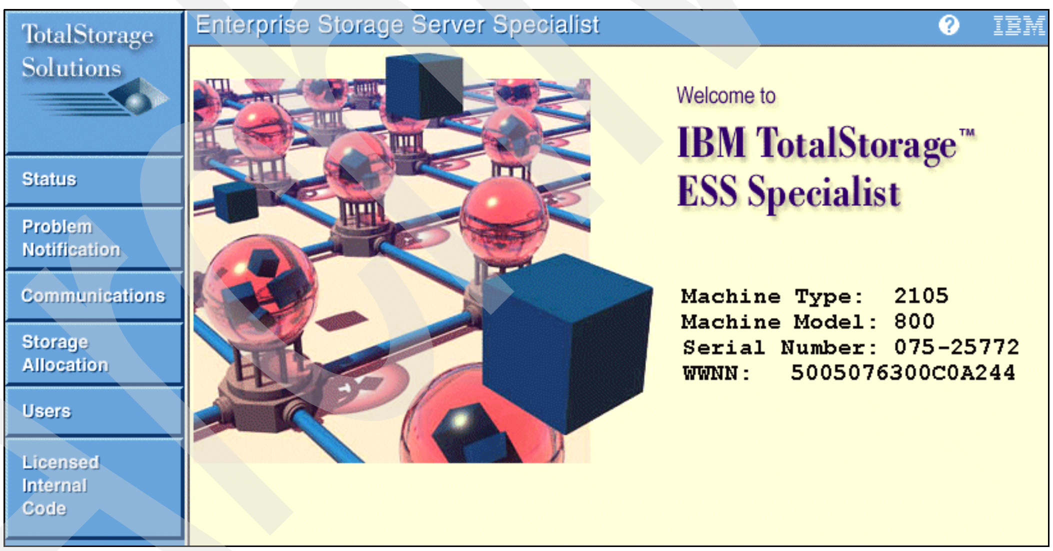

Figure 6-1 ESS Specialist Welcome panel

The Welcome panel displays the main menu options and provides the following information about the ESS machine you accessed:

- Machine Type, for example, `2105`
- Machine Model, for example `E20, F20 or 800`
- Serial Number, for example `MMM–14419`
- World-wide node name (WWNN) of the ESS, for example `5005076300C00A40`

The navigation frame on the left side provides buttons for accessing the tasks you can perform with subsequent panels of the ESS Specialist.

6.2.2 Access security

All data sent between the ESS and the ESS Specialist is encrypted to avoid unauthorized modification of configuration commands during transit. As a result, the Web browser warns the user, by displaying a sequence of access panels, that an encrypted site is being accessed.

Click any of the option buttons in the navigation frame to initiate the access security panels. You see the New Site Certificate panels if you use the Netscape Navigator browser. You see Security Alert panels if you use the Microsoft Internet Explorer browser. The New Site Certificate panels are shown in the following panels.

Note: The access security panels also change if you change the Java permissions on Internet Explorer.

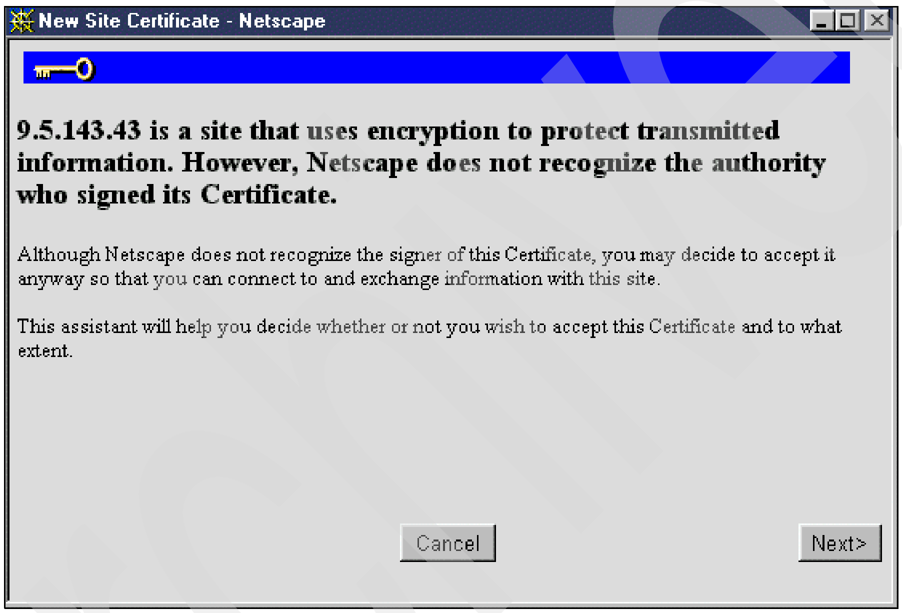

Figure 6-2 Security Alert panel (1 of 2)

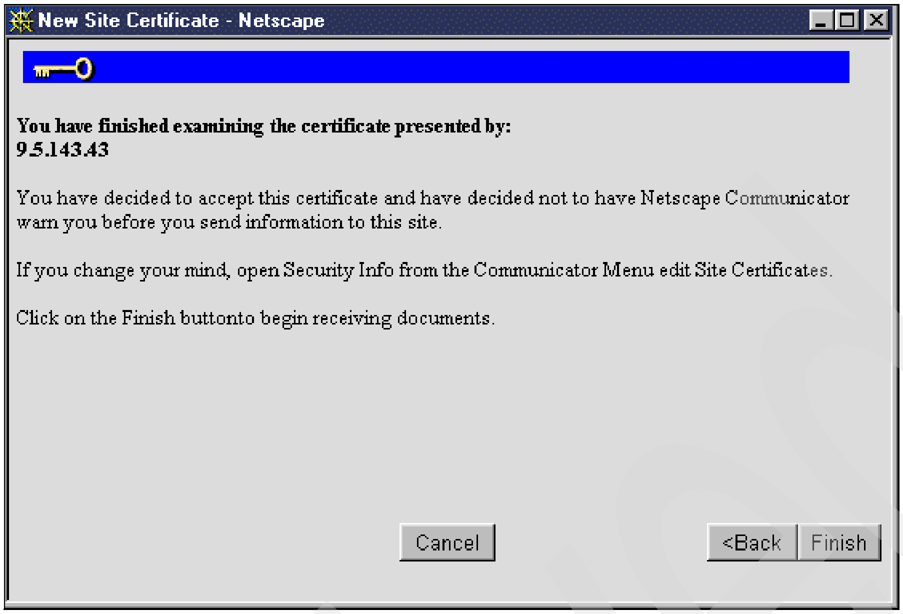

Figure 6-3 Security Alert panel (2 of 2)

With the site-access security panels, you tell your Web browser that the Web site represented by the ESS Web address should be treated as a trusted site. You also specify that encrypted communications should be allowed between your Web browser and this site.

6.2.3 User authentication

In addition to security by encryption, access to the ESS Specialist interface is protected by user names, passwords, and authorization levels. Access to each Enterprise Storage Server, through the ESS Specialist interface, requires a valid user name and password for that server. After clicking any of the buttons in the navigation frame and going through the encrypted-site sequence of windows, the Web browser prompts you for a user name and password.

If this is the first time you are accessing this IBM storage server, you can use the following default authorization:

1. Enter `storwatch` in the User Name field.
2. Enter `specialist` in the Password field.
3. Click **OK**.

Figure 6-4 Sign-on window

Note that both the user name and password are case sensitive. This default user name and password pair are automatically deleted from the ESS after the first administrative user ID is configured. Your first action as the storage server administrator is to set up a user ID with administration authority.

If the last user name and password with administration privileges are removed, the default user name and password are reinstated by the program.

If this is not your first time accessing this particular storage server, enter a valid user name and password combination. Your browser remembers your user name and password, so it does not prompt you again while you access the same storage server, until you close your browser application.

As an option, the ESS Specialist can be configured to require particular user name connections to originate from a specified IP address or range of IP addresses. See the description of the User Administration panel in 6.2.5, "Modifying user's access authority" on page 136, for more details about this feature.

6.2.4 Getting started with the ESS Specialist

Once you access the ESS Specialist Web interface, the navigation frame is displayed on the left side of the page (see Figure 6-1 on page 131).

For the initial configuration, click the buttons and perform the tasks in the following order:

1. **Users**: From the User Administration panel, click **Modify Users** to grant access to ESS Specialist for other users in your organization. See 6.2.5, "Modifying user's access authority" on page 136, for details.

   Note: Only a user with administrations level authority can modify this panel.

2. **Status**: When you click Status, the Status-Graphical View panel opens. See 6.3, "Checking ESS status" on page 138 for details.

   Note: Resolve any configuration problems *before* you select Storage Allocation.

3. **Communications**: From the Communications panel, select **Modify Remote Support** or **Reset PE Password** (at the bottom of the panel). Use these panels to enable callhome and remote service access and to set or change the passwords that authorize access by service personnel.

   The Communications panel also displays TCP/IP information. TCP/IP information is initially configured by IBM during installation. This information must be correct or ESS Specialist does not function correctly.

   Note: Only users with administration, configuration, or operation levels of authority can modify these panels.

4. **Problem notification**: From the Problem Notification panel, select **Modify e-mail**, **Modify Pager**, or **Modify SNMP**. Use these panels to enable the various types of problem notification and to provide the destination of these notifications.

   Note: Only users with administration or configuration levels of authority can modify these panels.

5. Storage allocation: From the Storage Allocation panel, you can view all the storage that is currently defined in your ESS. You can also access additional panels that allow you to define and configure other available storage.

   You use the worksheets in *IBM TotalStorage Enterprise Storage Server Configuration Planner*, SC26-7353, to plan your initial storage allocation for the ESS.

   Use the **Open System Storage** button to view and allocate fixed block (FB) storage for open system hosts. This is storage that is accessible by the hosts systems through the SCSI or Fibre Channel interfaces that are installed in the ESS.

   To allocate storage to disk groups and then to volumes, we recommend that you follow this sequence:

   Note: Only users with administration or configuration level authority can modify the storage allocation panels.

   a. Use the graphical view of the Storage Allocation panel to view all host and device adapters that IBM has installed on the ESS. Identify your attached host systems and any storage allocation already configured for you by IBM on the ESS. Your ESS may already be logically configured with one of these options (see 6.4, “Allocating storage” on page 141 for details):

      - OS/390® logical configuration formatted for 3390 Model 3 devices
      - OS/400 logical configuration, formatted for a maximum number of 9337 5BC/2105 Model logical unit numbers (LUNs) per array
      - Windows NT logical configuration, formatted with the LUN size equal to the array size
      - UNIX logical configuration, formatted for a maximum number of 8 GB LUNs per array
      - UNIX logical configuration, formatted for a maximum number of 16 GB LUNs per array

      You can use ESS Specialist to customize these basic logical configurations.

   b. Use the Open Systems Storage panel (Figure 6-11 on page 144) to display your iSeries or Open System configuration. Use the following panels to configure fixed-block storage:

      - *Modify Host Systems*: Use this panel to modify the list of host systems that are defined in the ESS Specialist. Host system nicknames are used in the ESS Specialist to aid in the configuration of the ESS.

        Note: The ESS adds an anonymous host whenever you configure the access mode to “access any” for the storage server. The anonymous host represents all hosts on the Storage Area Network that are not explicitly defined to the ESS, but that have access to one or more volumes configured on the ESS.

      - *Configure Host Adapter Ports*: Use this panel to configure the ESS Fibre Channel ports and to identify which ESS SCSI ports are connected to which host systems.

        Note: ESCON ports and their attached S/390 hosts are identified to the IBM storage server and ESS Specialist when you make the physical connection between the hosts and the storage server. This is a feature of the ESCON adapter cards.

- *Fixed Block Storage*: Use this panel to configure fixed-block disk groups as RAID or non-RAID disk groups. All disk groups that IBM has physically installed are displayed in a list on this panel. After you select and configure the disk groups, ESS Specialist displays the group on the Storage Allocation panel.
- *Add Volumes*: Use this panel to allocate fixed-block volumes in the available disk groups and assign the added volumes to a selected host or port.
- *Modify Volume Assignments*: Use this panel to:
  – Configure shared volume assignments with another host or port
  – Remove volume assignments from another host or port
  – Assign volumes to another host or port

6.2.5 Modifying user's access authority

The User Administration panel (Figure 6-5) displays a list of users who are authorized to use ESS Specialist. The list includes the user's level of access authority. With administration-level authority, you can grant access to other users in your organization.

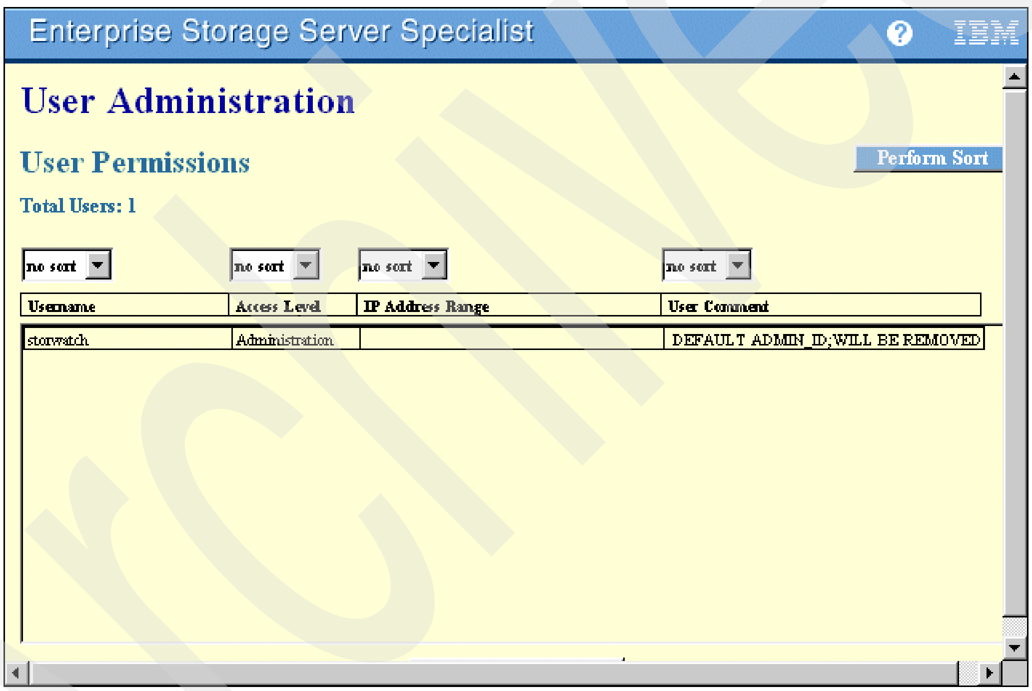

Figure 6-5 User Administration panel

Use the buttons at the top of each column to set sort priorities. For example, when your list of users is extensive, you can set the sort button for the Access Level column to sort first, and then click **Perform Sort**. You then see the type of access that you (or another administrator) assigned to the users.

To display the Modify Users panel, click **Modify Users** at the bottom of the User Administration panel (Figure 6-5).

From the Modify Users panel (Figure 6-6 on page 138), you can assign or change access authority, and add or remove users. To assign user access-level authority for each user, perform the following steps:

1. On the Modify Users panel, in the User Account box, complete the following information:
   a. Type the user name in the Name field.

b. Type the user password in the Password and Password Verification fields.

c. Select the user access authority in the Access Level field. The access authorities restrict a user ID to one of the following access levels:

- *View*: View-level authority only allows the user to view the ESS status and configuration.
- *Operation*: Operation-level authority allows the user to perform view and operation functions. This includes allowing the user to change the remote service and PE passwords.
- *Configuration*: Configuration-level authority allows the user to perform view, operation, and configuration functions. This includes allowing the user to define and modify the storage on the ESS.
- *Administration*: Administration-level authority allows the user to perform view, operation, configuration, and administration functions. This includes allowing the user to grant access to other ESS Specialist users.

d. Type the user's IP address range in the IP Address Range field. This is an optional entry that allows you to define the IP address or IP address range, from which the user can initiate client communications with the ESS Specialist.

> **Note:** You can use an asterisk (*) as a wild-card character to specify an IP address range. You can type comments about the user in the User Comment field. This is an optional field for entry of information about a user, such as their job title or how long they should have access authority.

2. Click **Add** to add the user to the User List.

To change a user's access level, perform the following steps:

1. Select the user name in the User List.
2. Click **Remove** to move the entry to the User Account window.
3. Select the new access level in the Access Level field.
4. Click **Add** to add the user, with the modified access level, to the User List.

To remove a user, perform the following steps:

1. Select the user name that you want to remove in the User List.
2. Click **Remove**.

> **Note:** Be sure that someone in your organization always has administration-level access. Without this level of access, the default user name and password are in effect, which could be a security issue.

Click **Perform** Configuration Update when you have finished adding access, modifying access, or deleting users. Click **Cancel Configuration Update** if you do not want to make the changes to the User List.

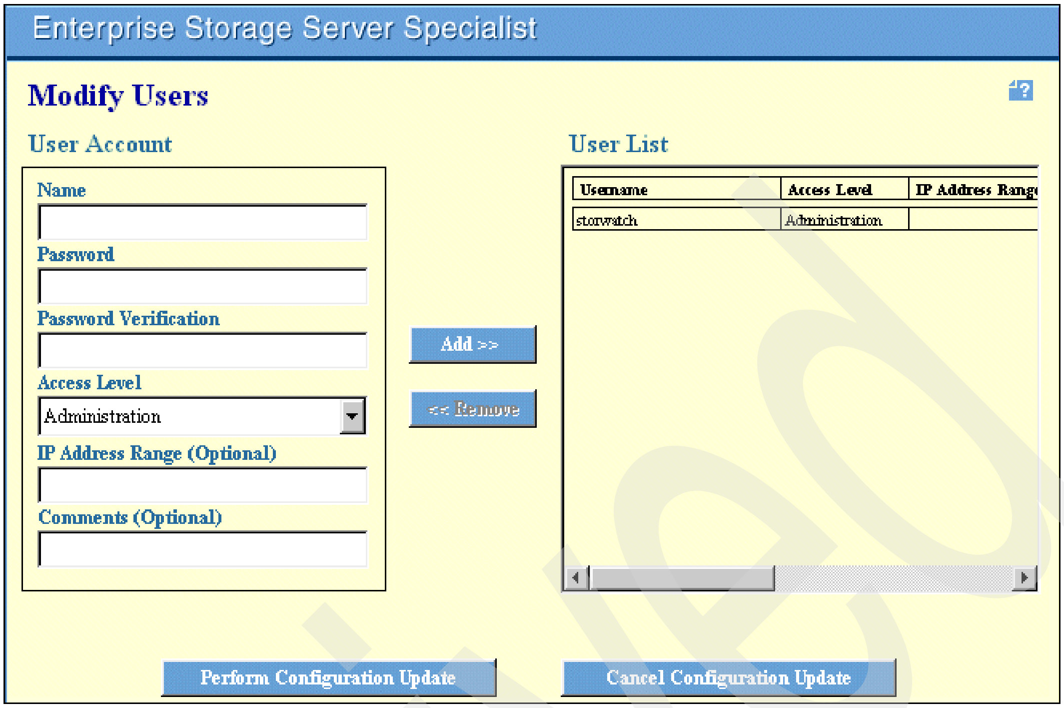

Figure 6-6 Modify Users panel

> **Note:** The ESS Specialist prompts for the user name and password every time it detects an insufficient authority level. For example, you may modify or delete the username and password that you used to log into the current ESS Specialist session, or you may add the first administration-level username and password (invalidating the default username/password). In either case, the ESS Specialist prompts you for the correct username and password from the list you configured.

6.3 Checking ESS status

Before you use ESS Specialist to initialize or modify the ESS configuration, review the ESS Specialist Status panel for problem records that require service attention. Ensure that the Status panel shows that all machine resources are in the normal state and that there are no problem records listed:

1. From the navigation frame, click **Status** to access the Status panel.
2. Click **Refresh Data** at the top of the Status panel to display the current status.

ESS Specialist displays a graphical view of the ESS.

6.3.1 Status graphical view

The graphical view of the Status panel (Figure 6-7) shows a simplified logical representation of the internal components of the server. The display highlights any components that are experiencing difficulty. Do not proceed with configuration changes until the problems are resolved by IBM or your service provider.

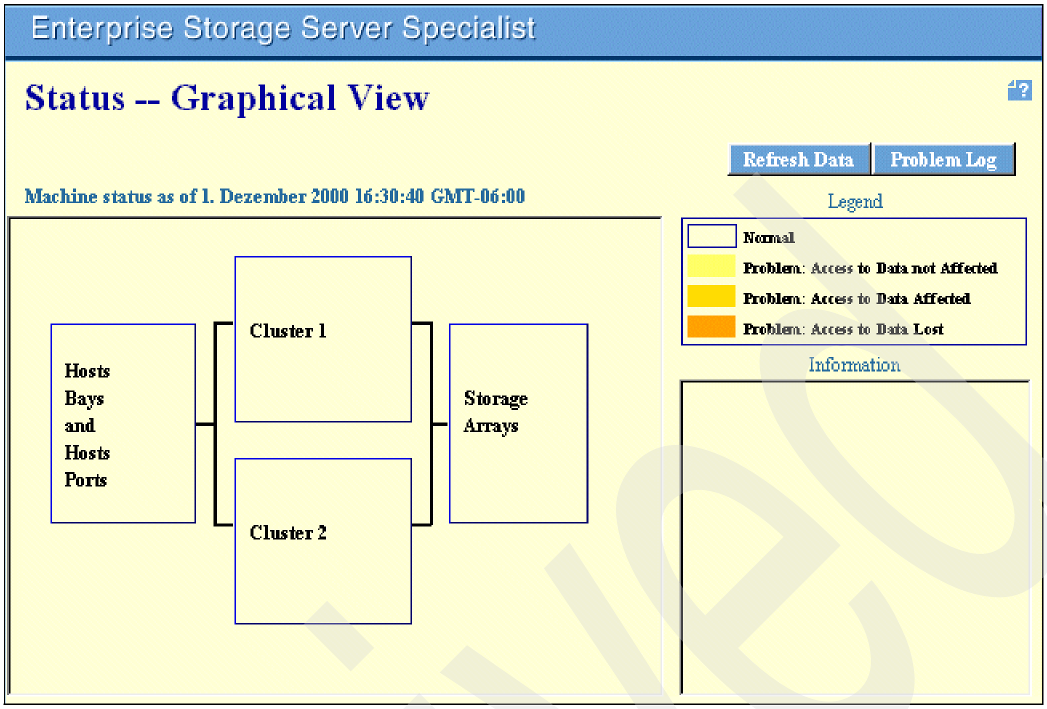

Figure 6-7 Graphical view of the Status panel

The connecting lines represent the physical data paths between the components:

- Host adapters
- Cluster controllers
- Storage arrays

If any component is failing, the server alerts the user by highlighting the element that corresponds to the location of the failed component. The color indicates the severity of a problem and shows its status:

- **White**: Normal status. No failure has occurred.
- **Yellow**: Problem. Access to data is not affected.
- **Light orange**: Problem. A component has failed and access to data is affected but not lost. Performance may be degraded.
- **Dark orange**: Problem. A component has failed, and access to data is lost. One or more hosts have lost access to one or more logical volumes defined in the storage facility.

Color can also indicate that service actions are occurring on the machine or that logical configuration is in progress. See the Information box for more details about a service action. For details about a particular problem, click the **Problem Log** button to access the Status - Problem Log panel.

6.3.2 Status problem log

This section explains the information in the Status - Problem Log panel (Figure 6-8).

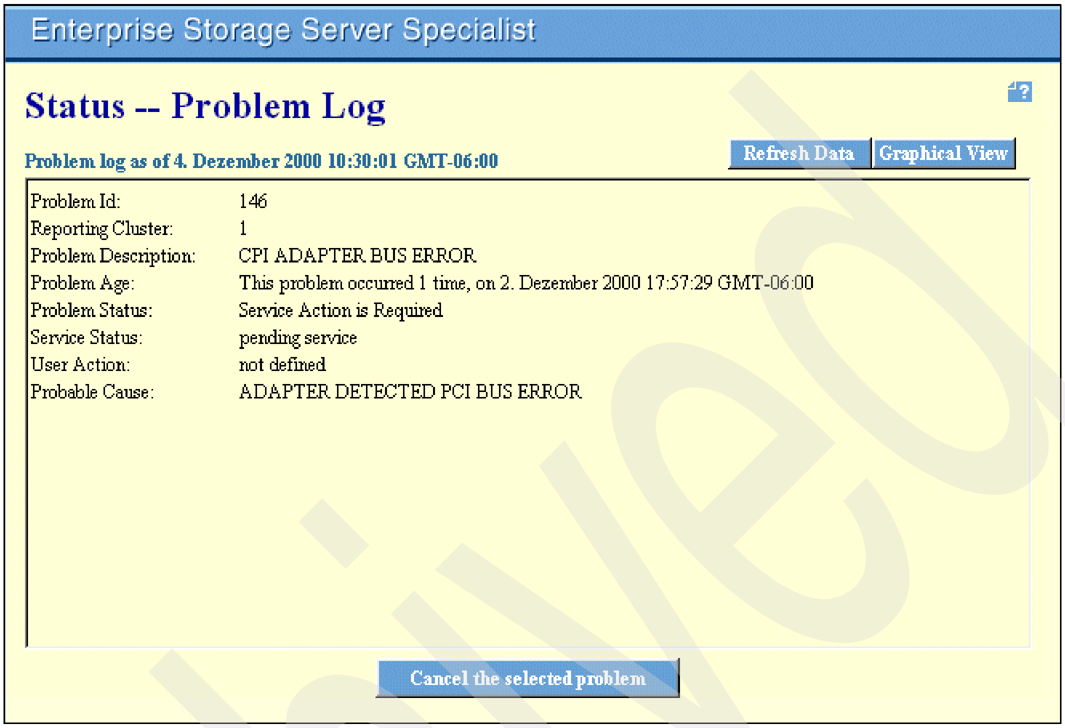

Figure 6-8 Status - Problem Log panel

Use the problem log to help you to identify any components or functions that are degraded or temporarily unavailable due to problem conditions or service actions. You should not proceed with configuration changes until any problems that are listed are resolved by IBM or your service provider. Follow these steps:

1. Click **Problem log** to access the problem log from the Status Graphical View panel.
2. Click **Refresh Data** at the top of the panel to display the most current problem log.

The log lists all problems that are currently open for the server. It offers the following information:

- **Problem ID**: The unique problem identification number
- **Reporting cluster**: Cluster 1 or 2
- **Problem description**: A description of the problem
- **Problem age**: The date, time, and frequency that the problem occurred
- **Problem status**: The problem severity
- **Service status Status** of the problem resolution
- **User action**: Any corrective action that you can perform
- **Probable cause**: A list of the probable causes of the problem
- **Additional information:** Additional information about this problem, if any

To remove a resolved problem from the log, follow these steps:

1. Select the problem by clicking the appropriate **Problem ID**.
2. Click **Cancel the selected problem** at the bottom of the page.

> **Note:** Your service support representative (SSR) can remove a problem from the Problem Log when they perform the required service.

3. To return to the Status Graphical View panel, click **Graphical View** in the upper right corner of the panel.

6.4 Allocating storage

ESS Specialist displays the Storage Allocation panel after you click Storage Allocation from the navigation frame. The Storage Allocation panel contains a logical view of the components of the ESS storage facility. It shows the server's logical configuration of host systems, host ports, storage adapters, arrays, and volume assignments. Your user ID, along with configuration or administration authority, allows you to select options that modify the server configuration.

The Storage Allocation panel consists of two views, a graphical view and a tabular view. The graphical view (Figure 6-9) shows the server hardware components with interconnecting lines to indicate a particular assignment or logical connection.

6.4.1 Using the graphical view of the Storage Allocation panel

The graphical view is interactive. It allows you to explore the path of a particular host to a particular adapter.

The top row contains icons for the installed host system. ESS Specialist adds the anonymous host whenever you configure the Fibre Channel LUN access mode to "access any". The anonymous host represents all hosts on the Storage Area Network that are not explicitly defined to the ESS, but which have access to one or more volumes configured on the ESS. See 6.7.1, "Volume access" on page 155, for more information about setting access modes.

When you click a host icon to select it, the Information box on the right side of the panel displays information about the host nickname and operating system.

The second row contains icons for the host adapters. Each SCSI adapter has two selectable ports. Each Fibre Channel adapter has one port. When you select a host, the configured ports are highlighted in yellow, and lines appear between the host icon and the port icons.

When you select a port or adapter, the ESS Specialist highlights the arrays with volumes that you can access from the host that you selected. The Information box displays information about the adapter and its location. For a Fibre Channel adapter, it also displays information about whether the Fibre Channel uses point-to-point or arbitrated loop protocol.

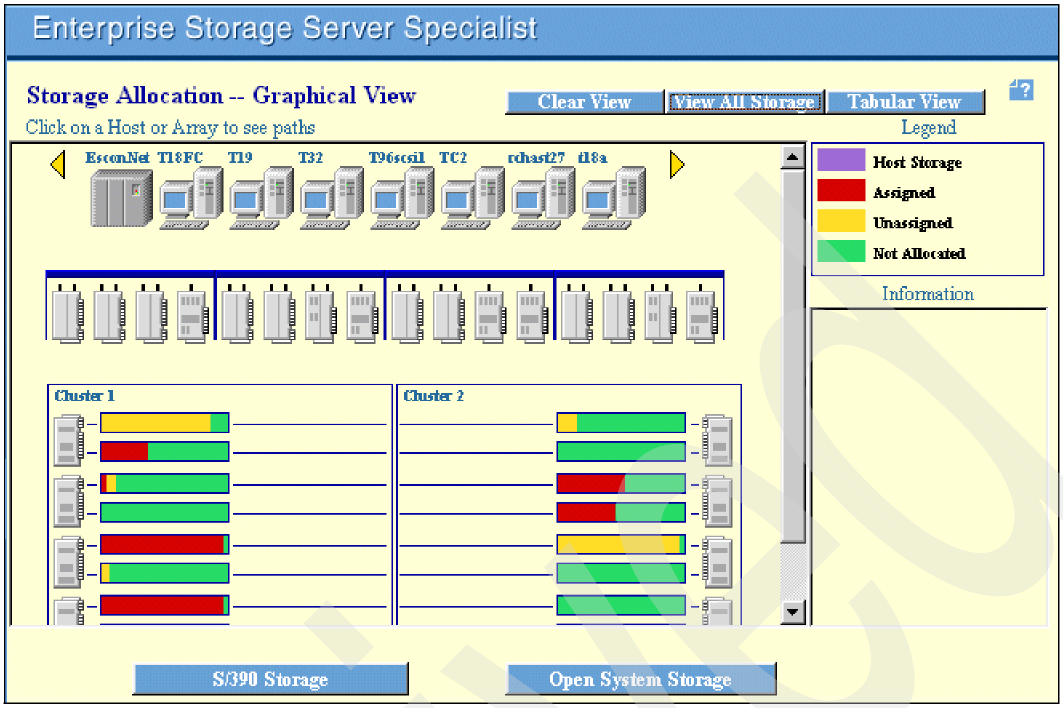

Figure 6-9 Graphical view of the Storage Allocation panel

The bottom of the Storage Allocation - Graphical View panel has two buttons:

- **S/390 Storage**: Click this button to go to the main panel for S/390 hosts (ESCON attached).
- **Open System Storage**: Click this button to go to the main panel for iSeries and open systems (SCSI and Fibre Channel-attached hosts).

> **Note:** If the ESS has no ESCON host adapters, ESS Specialist grays out the S/390 Storage button. If the ESS has no SCSI or Fibre Channel Host Adapters, ESS Specialist grays out the Open System Storage button.

6.4.2 Using the tabular view of the Storage Allocation panel

The tabular view of the Storage Allocation panel (Figure 6-10) presents the same information as the graphical view but in table format. Individual table columns represent the hardware components. The table view lists the information for all host systems and volumes at one time. The information is static, with the exception of the buttons located in the upper right corner:

- Click **Print Table** to print a copy of the storage allocation that the ESS displays in the table.
- Click **Perform Sort** to sort on different columns. See 6.1.2, “Using ESS Specialist tables” on page 129, for more information about sorting.
- Click **Graphical View** to return to the graphical view of the Storage Allocation panel.

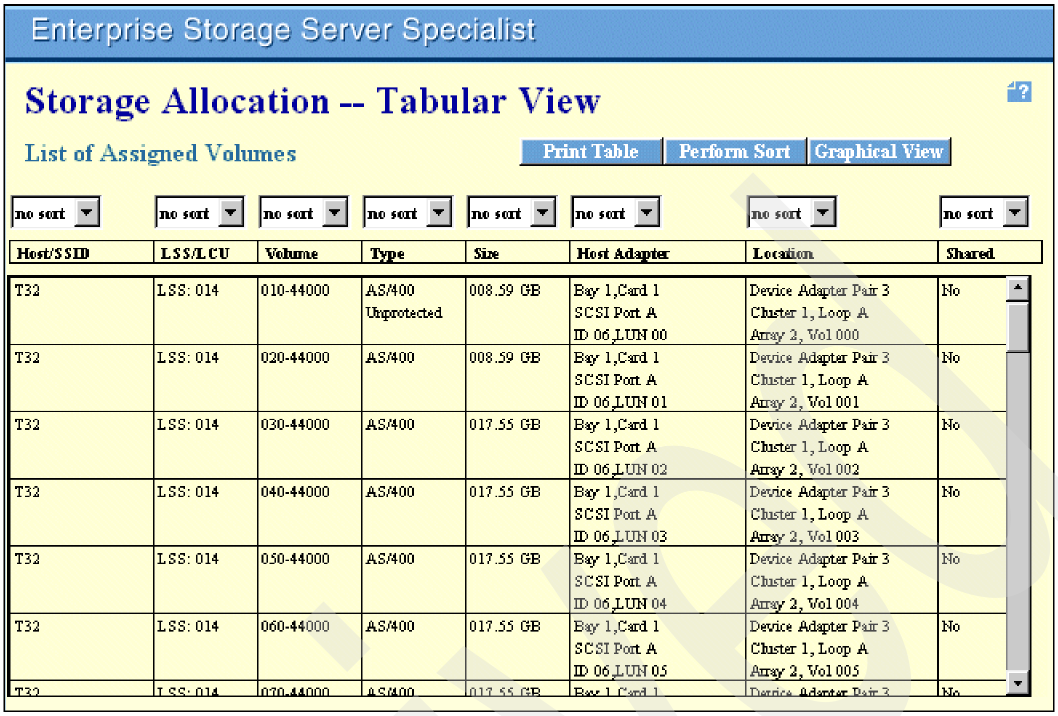

Figure 6-10 Tabular view of the Storage Allocation panel

Perform the following steps to allocate storage:

1. Access ESS Specialist. You need configuration-level access to allocate storage.

> **Note:** IBM install your ESS and performs an initial standard configuration as defined by you in the worksheets from the IBM Enterprise Storage Server Configuration Planner. If you are using Specialist to perform a custom configuration, use these worksheets.

2. Select the graphical view of the Storage Allocation panel to perform the configuration tasks.
3. From the Storage Allocation panel, click **Open System Storage**. This panel identifies the configured host systems with their attached storage. Use the buttons at the bottom of this panel to configure new host systems and storage or to modify existing configurations:
   - **Modify Host Systems**: Use this panel to identify host systems that you are attaching to the storage server.
   - **Configure Host Adapter Ports**: Use this panel to identify which host ports and IDs are used to attach host systems to the storage server by SCSI cables, to set the topology for Fibre Channel ports, and to see the Fibre Channel access modes of the ESS.
   - **Configure** Disk Groups: Click this option to go to the Fixed Block Storage panel. Use this panel to define fixed-block disk groups as RAID or on-RAID to prepare them for volume assignment. This panel displays all disk groups that are undefined or defined as fixed block. Disk groups that you have defined as count-key-data (CKD) are not displayed. After you select and define the disk groups, ESS Specialist displays the group on the Storage Allocation panel.
   - **Add Fixed Block Volumes**: Use this panel to allocate fixed-block volumes in the available disk groups and to assign the volumes to a selected SCSI port storage map or to Fibre Channel hosts.

– **Modify Volume Assignments**: Use this panel to:

- Share an assigned volume with another SCSI port or a Fibre Channel host
- Remove a volume assignment from a SCSI port or Fibre Channel host
- Add an unassigned volume to an eligible SCSI port or Fibre Channel host

6.5 Configuring open system storage for iSeries or AS/400

Use the Open System Storage panel to assign physical storage units for Fibre Channel- and SCSI-attached host systems.

Click **Open System Storage** on the Storage Allocation panel. The information on the Open System Storage panel (Figure 6-11) is presented in the form of two tables.

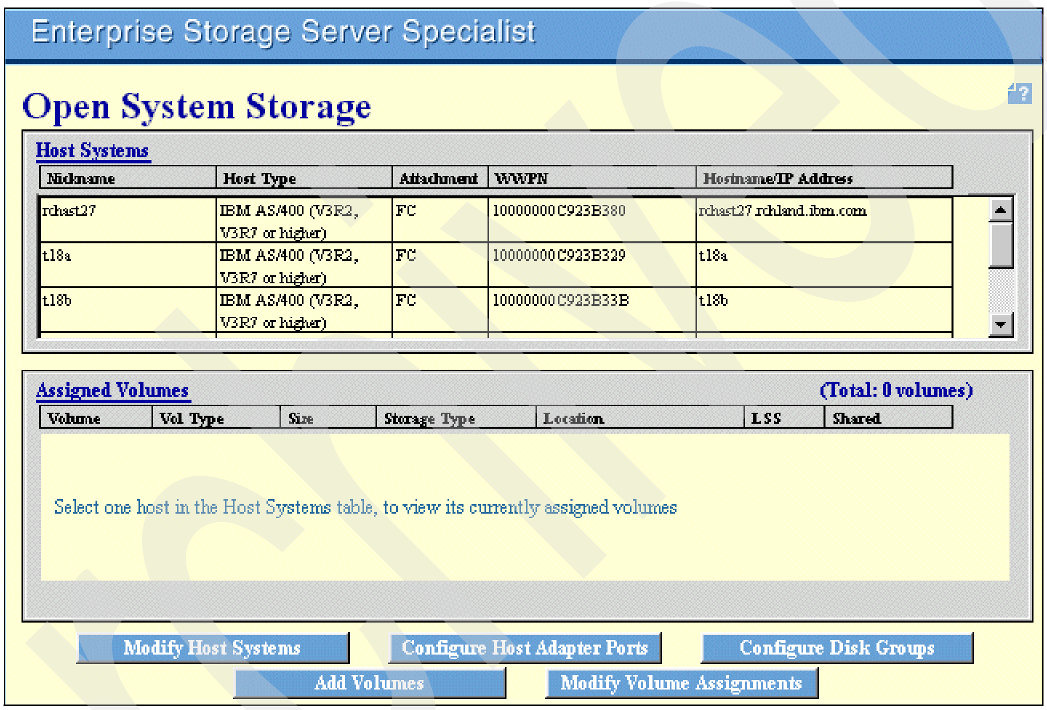

Figure 6-11 Open System Storage panel

Host Systems

The Host Systems table displays all host systems that are defined to the ESS Specialist. If no host systems are defined, the following message is displayed in the table area:

```
No host systems are currently defined
```

The following columns are displayed in this table:

- **Nickname**: The nickname of this host system. This name uniquely identifies the particular host or host attachment within the ESS.
- **Host Type**: The hardware or operating system type of the host.
- **Attachment**: The type of host attachment used by this host; either SCSI attached or Fibre Channel-attached.
- **WWPN**: The world-wide port number (WWPN) associated with the Fibre Channel adapter installed in the host system. This field applies only to Fibre Channel-attached hosts.

- **Hostname/IP address**: An optional Internet Protocol dotted decimal address or host name alias.

Assigned Volumes

The Assigned Volumes table displays all volumes that are assigned to the host system that you selected in the Host Systems table at the top of the screen. If the selected host has no assigned volumes, the following message is displayed in the table area:

```
There are no volumes assigned to the currently selected host system
```

The following columns are displayed in this table:

- **Volume**: This column contains the LUN serial number associated with this volume. The LUN serial number of an ESS volume incorporates a portion of the ESS serial number. This allows LUN serial numbers to uniquely identify volumes within a particular ESS and volumes among all ESS machines that are installed in an enterprise.
- **Vol Type**: This column identifies the type of the volume. This can be either “AS/400” for volumes allocated for AS/400 host systems, or can be “open system” for volumes allocated to UNIX-based host systems or Intel-based personal computers (PCs).
- **Size**: This column displays the size of the volume, in gigabytes. The value is rounded to two decimal places.
- **Storage Type**: This column displays relative fault tolerance of the storage where the volume resides. The possible types are RAID Array (fault tolerant), and non-RAID (non-fault tolerant).
- **Location**: This column provides an identification of the physical location of the volume. This includes the number of the device adapter pair, the number of the cluster, the ID of the SSA loop, and the number of the disk group associated with the RAID array which contains the data for this volume. In case of a non-RAID volume, it also includes the number of the non-RAID disk within the disk group that contains the volume. In addition, the volume number is included to differentiate between different volumes on the same array or disk
- **LSS**: This column contains the number of the logical subsystem (LSS) that contains this volume. This value is important for performing Copy Services operations on the volume. The ESS displays the LSS number as hexadecimal.
- **Shared**: This column indicates whether this volume is shared. A volume is considered shared if multiple host systems have access to the volume. The only exception to this rule is if multiple hosts are connected to the same SCSI bus. The volumes connected to this bus are not considered shared by these hosts because access to these volumes uses the same host interface connection. iSeries servers *do not* support shared volumes.

At the bottom of the panel, follow these steps:

1. Click **Modify Host Systems** to link to the Modify Host Systems panel.
2. Click **Configure Host Adapter Ports** to link to the Configure Host Adapter Ports panel.
3. Click **Configure** Disk Groups to link to the Fixed Block Storage panel.
4. Click **Add Volumes** to link to the Add Volumes (1 of 2) panel.
5. Click **Modify Volume Assignments** to link to the Modify Volume Assignments panel.

6.5.1 Defining host systems attached to the storage server

The Modify Host Systems panel allows the user to modify the list of host systems that is defined in the ESS Specialist. Use the Modify Host Systems panel to:

- Add a host system to the list of host systems that are available to be attached to the storage server.

- Identify a SCSI or Fibre Channel-attached host system by type and name.
- Assign a world-wide node name.
- Remove a selected host from the Host Systems List.
- Modify the attributes for an existing host system.

For more information regarding the hosts, operating systems, and host adapters and cables that the ESS supports, see the ESS Web site at:

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

The left side of the Modify Host Systems panel displays the Host Attributes box with the characteristics of individual host systems. You can modify the following information:

- The Nickname field represents the name of the host system, as it is displayed in the ESS Specialist user interface.
- The Host Type field indicates the type of this host system. The scroll bar allows you to scroll to hosts that may not be displayed in the table. The ESS supports the following types:
  - IBM S/390 with MVS™, VM, and VSE operating systems
  - IBM RS/6000®, including RS/6000 iSeries Parallel Complex, running AIX®
  - IBM AS/400, AS/400 Advanced iSeries, and iSeries

    For AS/400 hosts, an additional distinction is made based on the operating system that is installed:

    - AS/400 hosts that run OS/400 Version 3 Release 1 or Version 3 Release 6 only support one volume size – the emulated 9337-580 (4.19 GB).
    - AS/400 hosts that run OS/400 Version 3 Release 2,Version 3 Release 7, or a higher level support all AS/400 volume sizes, including the 9337-580 (4.19 GB), the 9337-590 (8.59 GB), the 9337-5AC/5AA (17 GB), the 9337-5CC/5CA (35.16 GB), and the 9337-5BC/5BA (36 GB).
    - iSeries hosts that run OS/400 Version 5 Release 1 support additional Fibre Channel connected AS/400 volume sizes, including the 2105-A01/A81 (8.59 GB), 2105-A02/A82 (17.54 GB), 2105-A03/A83 (36 GB), the 9337-A05/A85 (35.16 GB), and 2105-A04/A84 (70.56 GB).

    > **Note:** AS/400 model designations that end in C are protected models. Models that end in A are unprotected models. Software mirroring is supported on the AS/400 unprotected models.

  - Sun Sparc or Ultra: For Sun hosts, an additional distinction is made based on the operating system installed:

    - Sun hosts that run the Solaris operating system at level 2.5.1 or that run the Solaris 2.6.0 base level only support a maximum of 8 LUNs per SCSI device. This limits the number of volumes that the Enterprise Storage Server can assign to those hosts.
    - Later levels of Solaris support up to a maximum of 32 LUNs per device. This 32 LUN support is included in the Solaris 2.6.0 level if the following patches are installed:
      - SunOS 5.6 kernel update patch 105181-06 or later
      - SunOS 5.6 kernel/drv/isp patch 105600-06 or later
    - HP 9000 iSeries 800
    - Compaq Alpha Server, running Tru64 UNIX or OpenVMS

- Data General Server, running DG-UX
- IBM NUMA Servers
- All Intel-based servers running Novell NetWare
- All Intel-based servers running the Windows NT operating system level 4.0 or later

- The Host Attachment field shows the type of interface used to connect the host to the storage server. The available choices are SCSI or Fibre Channel. The choice depends on which host adapter cards you installed in the ESS.
- The Hostname/IP Address field displays the hostname alias or dotted decimal IP address, if the host is attached to the IP network. This field is optional.
- The World-Wide Port-Name field represents the port name of the Fibre Channel adapter installed in the host system. This field applies only if the host system is Fibre Channel-attached to the ESS.
- The Fibre Channel Ports field lists all of the Fibre Channel ports that are installed in the ESS. The highlighted entries indicate which ports this particular host system can use to access its assigned volumes in the ESS. If the first entry, All installed ports, is highlighted, then the host system can use all Fibre Channel ports in the ESS to access its assigned volumes. The default value for this field is All installed ports.

On the right side of the screen, the Modify *Host Systems* panel displays the Host Systems List box. The columns in this box display all the host attributes. You can select a host from the list to move it to the Host Attributes box for modification. Perform the steps in the following section to add a host system.

6.5.2 Adding a host system

To add a host system, follow these steps:

1. Click **Modify Host** from the Open Systems Storage panel (Figure 6-11 on page 144). The Modify Host Systems panel opens (Figure 6-12).

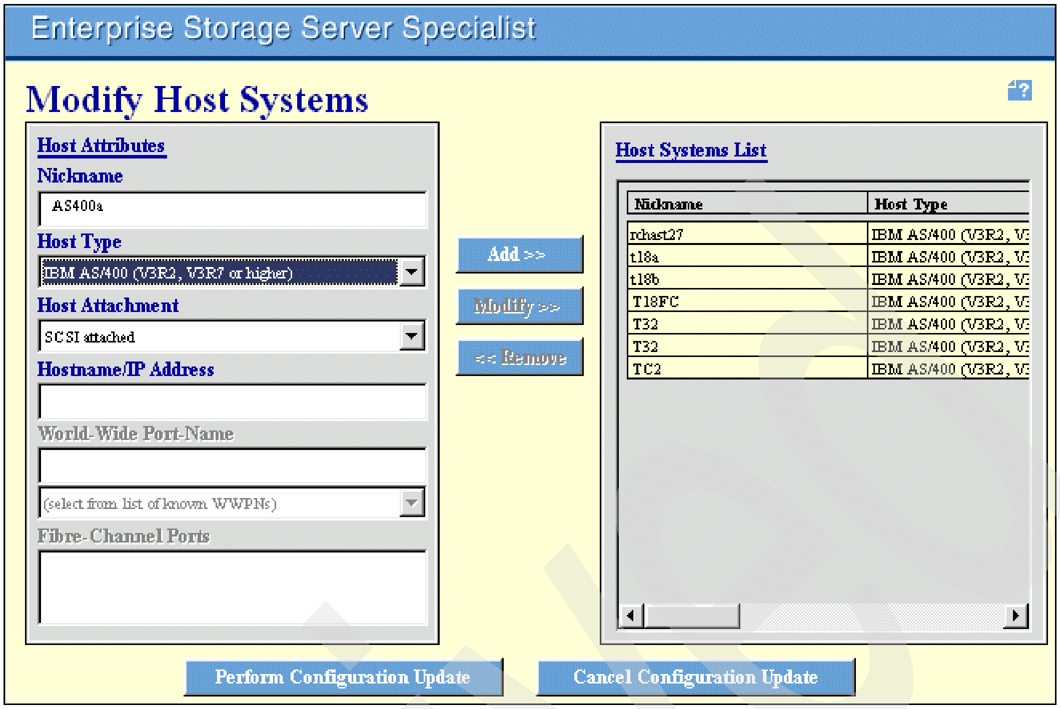

Figure 6-12 Modify Host Systems panel SCSI connection

2. To add a host system, enter the following information in the Host Attributes box on the Modify Host Systems panel:

   a. Enter the host name in the Nickname field. A host nickname must be unique and must not exceed 29 characters.

   b. Select a host type from the drop-down list in the Host Type field. See the list of supported host types above. For the iSeries or AS/400, select **IBM AS/400** (V3R2/V3R7 or later).

   c. From the Host Attachment drop-down list, select either **SCSI attached** or **Fibre Channel attached**. SCSI attached is the default selection.

   > **Note:** If a host has multiple SCSI adapters installed, you need to enter only one nickname in the host system list. ESS Specialist uses the same name for displaying the different SCSI connections. However, if a host has *both* SCSI and Fibre Channel adapters or *multiple* Fibre Channel adapters installed, you must enter a separate nickname in the Host System List for each Fibre Channel connected to the ESS. This is because a unique world-wide port-name is required for each Fibre Channel connection.

   For example, if an iSeries host needs more than one FC adapter, you must define a host (host s1, host s1a, host s1b) for each FC adapter.

   Figure 6-13 shows the options for Fibre Channel-attached hosts.

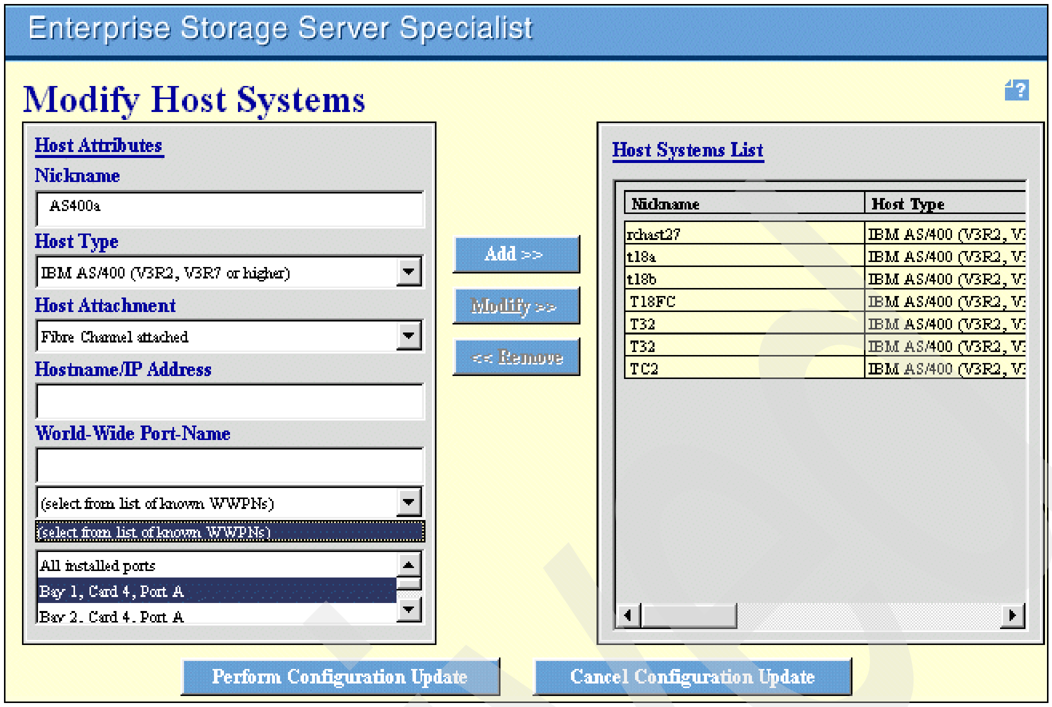

Figure 6-13 Modify Host Systems panel Fibre Channel connected

d. Optionally, type the host name or the dotted decimal IP address in the Hostname/IP Address field, if the host is connected to the IP network. We recommend that you enter information in this field if you are using the IBM StorWatch Expert software package.

e. If the host is Fibre Channel-attached, select the world-wide port name of the Fibre Channel adapter that is installed in the host. Select the name from the drop-down list in the World Wide Port Name field, or enter the WWPN manually.

The WWPN field is mandatory for Fibre Channel hosts. It is enabled only when you select Fibre Channel from the Host Attachment field drop-down list. If the host is SCSI attached, this field is disabled.

The WWPN drop-down list contains the WWPNs of all host Fibre Channel adapters that are currently connected but not yet defined to the ESS. The ESS “discovers” the WWPN from your host Fibre Channel adapter when you attach your host system to the ESS.

The Enterprise Storage Server supports only IEEE compliant Fibre Channel adapters.

> **Note:** In order for ESS Specialist to present the correct WWPN, you must have the host Fibre Channel adapter connected to the ESS. Either reboot or run the equivalent **`cfgmgr`** command so that the host Fibre Channel adapter initiates the log into the ESS Fibre Channel adapter.

If the connection is made via a single ESS Fibre Channel port, that port identifier is also listed. If the connection is made with multiple ESS Fibre Channel ports, the ports are not indicated in the drop-down list.

See the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide,* SC26-7446, for instructions on locating the WWPN for the fibre-adapter on any of the supported operating systems.

f. If the host is Fibre Channel-attached, optionally select one or more entries from the list in the Fibre Channel Ports field. If you do not make a selection from this field, the default is used. This allows this host system to access the ESS through all Fibre Channel ports installed in the ESS. To limit the ports through which this host can access the ESS, select one or more individual ports in the list.

> **Note:** If you enter the WWPN manually, be aware that a valid WWPN consists of exactly 16 hexadecimal digits. These digits uniquely identify that particular host Fibre Channel adapter. The hexadecimal characters are 0 -9, a -f, and A-F.

3. Click **Add**. The host nickname and type are added to the Host System List on the right side of the Modify **Host Systems** panel.
4. Repeat the above steps for each host that you are adding.
5. Click **Perform Configuration Update** to apply the configuration changes to the ESS. A progress bar is displayed indicating the progress of the configuration action.

   Click **Cancel Configuration Update** at the bottom of the panel to cancel the information that you entered.

6.5.3 Removing a host system

Follow these steps to remove a host system from the Host System List:

1. Click the system to be removed in the Host Systems List on the right side of the panel. Then click **Remove**.

> **Notes:**
>
> - The action of removing a system from this list does not affect machine operation, until you click the Perform Configuration Update button at the bottom of the panel.
> - Remove a host system from the list only after you disconnect it physically from the storage server. When you select a host system that has assigned volumes, the ESS displays the following warning message:
>
> ```
> Host has assigned volumes. Removal of host xxxx will cause it to lose access to
> all nn volumes...Are you sure that you want to make this change? (Yes/No)
> ```

2. Repeat Step 1 for each host that you are removing.

Follow these steps to modify the characteristics for an existing host system:

1. Click to select the row in the table of the host system to be modified in the Host Systems List. Attribute fields that you cannot change are disabled. The characteristics of the selected host are displayed in the Host Attributes fields.
2. Modify the host's nickname, type, attachment, WWPN, and hostname alias, as needed in the Host Attributes fields.

> **Note:** If the host that you selected already has one or more volumes assigned to it, modification of the host type and the host attachment may result in one or more of those volume assignments being lost. The ESS issues an appropriate warning message in this case.

3. Click **Modify**.

   If you modified the nickname, click either **Add** or **Modify**:

   – Clicking Add adds the host as a new host to the list, keeping the original host unchanged. Volume assignments are *not* copied.
   – Clicking Modify modifies the existing entry in the list. This is especially important when the selected host already has one or more volumes assigned, because the volume assignments are copied for a modified host.

4. The Specialist displays the selected host with the changed attributes in the Host System List.
5. Repeat the previous four steps for each host that you are modifying.
6. When you complete all add, remove, or modify actions, perform these steps:
   a. Click **Perform Configuration Update** to apply the configuration changes to the ESS. A progress bar appears, indicating the progress of the configuration action, as it is being processed by the ESS.
   b. Click **Cancel Configuration Update** at the bottom of the panel to cancel the information you entered.

The Specialist returns to the Open System Storage panel. Updates to the set of host system icons take effect after the storage server completes the internal configuration change.

6.6 Configuring host adapter ports

The Configure Host Adapter Ports panel (Figure 6-14) identifies the attachment of one or more host systems to a storage server host port. The icons at the top of the panel represent the host adapters. The Fibre Channel adapter icon has one port. This single port distinguishes the Fibre Channel adapter from the SCSI and ESCON adapters that have two ports.

Next you see the Host Adapter Port drop-down list. The port adapter displays information about port identification, for example, Bay 2, Adapter 3, Port A (SCSI).

The Storage Server Attributes and FC Port Attributes boxes display options for the adapter type you select. See 6.6.2, “Configuring a SCSI port” on page 153, and 6.6.1, “Configuring a Fibre Channel port” on page 151. You do not need to perform configuration for ESCON adapter ports.

When you access the ESS through several adapters from one host, spread the adapters evenly over the ESS adapter bays to achieve better workload balance. Volumes assigned to one adapter may be from a single array or from multiple arrays controlled by multiple device adapters and clusters. For best performance, the volumes assigned to an adapter should be from multiple arrays that are controlled by as many device adapters and clusters as possible.

6.6.1 Configuring a Fibre Channel port

Click **Host Adapter Ports** from the Open System Storage panel. Figure 6-14 shows an example of the options for a Fibre Channel adapter type.

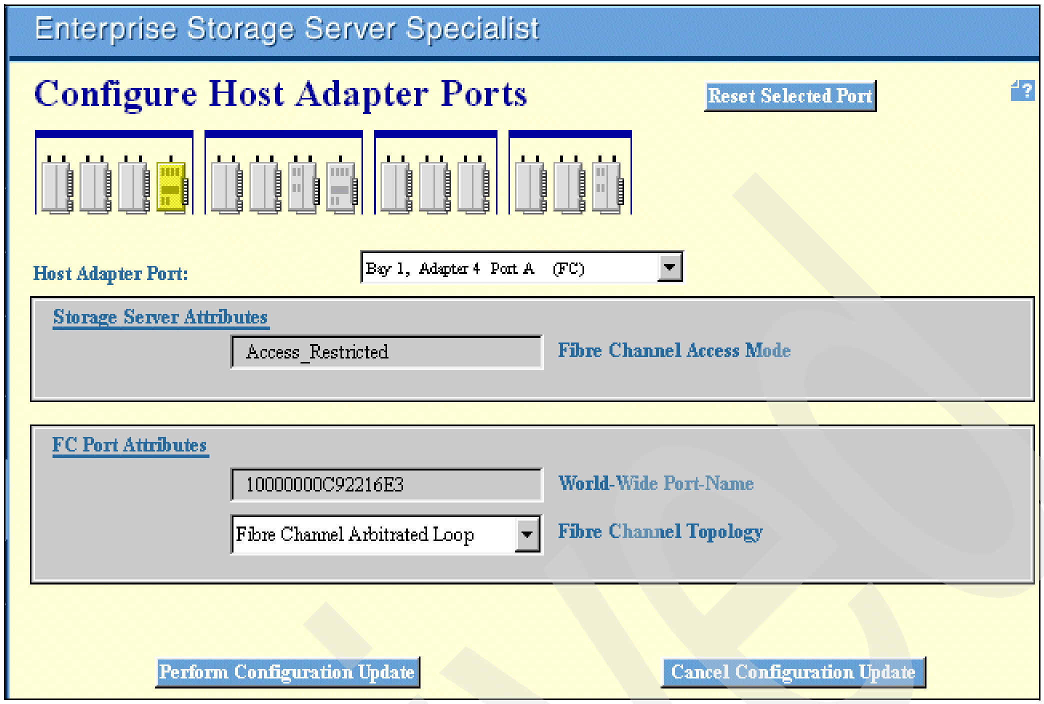

Figure 6-14 Configure Host Adapter Ports panel: Example of a Fibre Channel configuration

On this panel, perform the following actions:

1. Select the **host adapter port** you want to work with. A small red line is displayed on the panel under each Fibre Channel port after you configure or modify them.
2. The Fibre Channel Access Mode field in the Storage Server Attributes box shows the current Fibre Channel access mode for the port that you selected. The Specialist displays one of the following modes:
   – Access any
   – Access restricted

   For iSeries: We strongly recommend you to have access mode access restricted. See 6.7.1, “Volume access” on page 155.

3. The FC Fabric Topology field in the FC Port Attributes box shows the current Fibre Channel topology for the port you selected. If the topology is undefined, you can use the drop-down menu to select either:
   – Point-to-Point
   – Arbitrated Loop

   For iSeries: Select **Arbitrated Loop** for a hub connection and switch with QuickLoop support.

   If the topology is defined, you must first change the setting to **Undefined** before the ESS can make an alternate setting available for configuration.

   If you want to change the port attributes, an IBM SSR can perform this change during a service action or installation of a new host adapter. See 6.7.1, “Volume access” on page 155, for more information.

4. At any time, you can click **Reset Selected Port** to cancel any pending configuration changes made to the currently selected port.
5. When all entries are complete for the Fibre Channel port, click **Perform** Configuration Update to apply or **Cancel Configuration Update** to cancel the configuration step for all modified ports.

6.6.2 Configuring a SCSI port

Figure 6-15 shows an example of the SCSI-port configuration options.

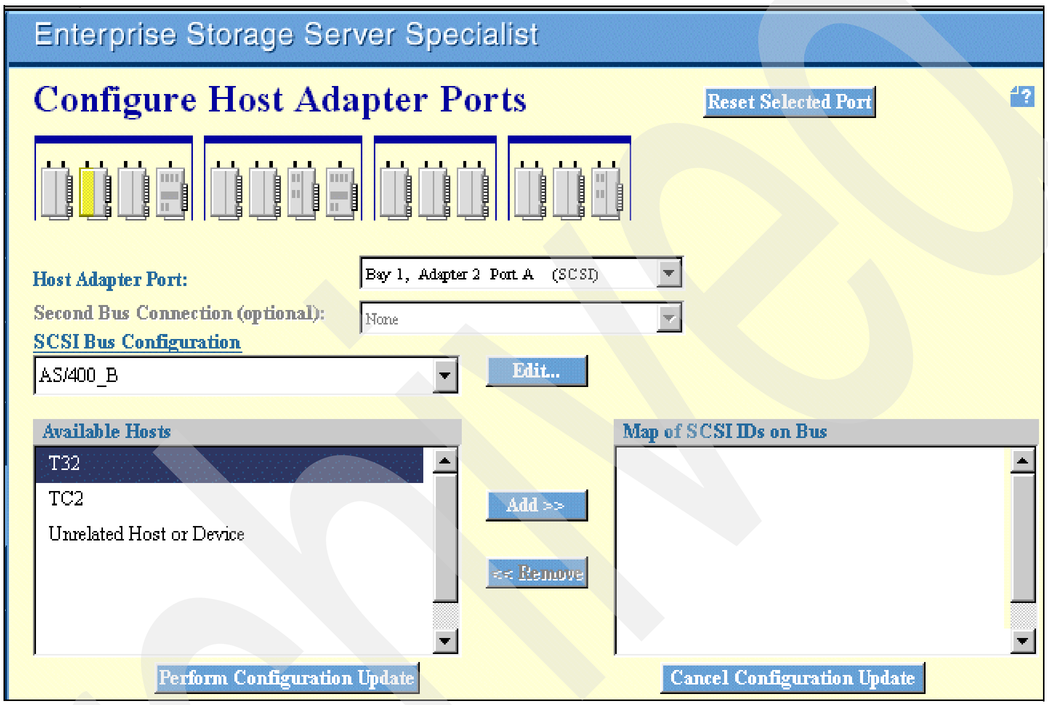

Figure 6-15 Configure Host Adapter Ports panel: Example of a SCSI configuration

The options for the SCSI adapters include:

- Reserve SCSI bus SCSI IDs for use by a second host port that may be attached to the SCSI cable
- Reserve an ID for another SCSI device

To configure a SCSI port, follow these steps:

1. From the Configure SCSI Ports panel, select a port either by clicking a port icon at the top of the panel or by selecting the port number from the Host Adapter Port field. The selected port is highlighted in yellow and the port number is identified in the Host Adapter Port field.

   When you select a port, the remainder of the panel displays information about the selected port:
2. The SCSI Bus Configuration field identifies the set of configuration parameters in effect for the SCSI bus that is associated with the selected port. If you have not yet configured the selected port, then the bus configuration displays none.

   Refer to the Bus Configuration field to identify the configuration parameters to be used for the selected SCSI port.

In the SCSI Bus Configuration field, select **IBM AS/400_A** for an AS/400 host running OS/400 V3R1 or V3R6 or **IBM AS/400_B** for an AS/400 host running OS/400 V3R2, V3R7, or a later level. Figure 6-15 shows that **AS/400_B** is selected for iSeries or AS/400.

A warning message is displayed if you already configured the port or if you select a conflicting host type.

3. The Available Hosts box lists the host systems that are compatible with the selected bus configuration.

   Select a host system that you want to attach to the port you selected, and click **Add** to list the host system in the map of SCSI IDs.

4. The Map of SCSI IDs on Bus box displays a list of SCSI IDs that are assigned to the attached host systems (maximum of four) and the SCSI IDs on the bus that are reserved for unrelated device attachments. The map of SCSI IDs also identifies any SCSI ID that currently has LUNs assigned. Any SCSI ID that is not listed in the map is available to be used for future LUN assignments.

   Select a proper numeric value for each SCSI ID. This value is the SCSI ID numeric value for the host system or unrelated device SCSI adapter. All SCSI ID values in the list must be unique. An iSeries or AS/400 will have the default SCSI ID 7.

   The default values in the SCSI ID selection field are in accordance with SCSI protocol, which defines SCSI ID 7 as having the highest priority. (Refer to 6.7, “Performance guidelines for configuring SCSI ports” on page 154.) The first host system that you add is assigned to SCSI ID 7, and the second is assigned to SCSI ID 6. You must verify that these assignments match the SCSI ID setting in each host system SCSI adapter card, and make adjustments to the map of SCSI IDs as necessary.

5. At any time, click **Reset Selected Port** to cancel any pending configuration changes made to the currently selected port.

6. When all entries are complete for each SCSI port, click **Perform** Configuration Update to apply or **Cancel Configuration Update** to cancel the configuration step for all modified ports.

6.7 Performance guidelines for configuring SCSI ports

For the best performance, we recommend that you follow these configuration guidelines:

- The SCSI ID priority is assigned in the order of 7 – 0 and then 15 – 8, where ID 7 has the highest priority and ID 8 has the lowest.
  a. Assign host systems to the highest priority SCSI IDs, with the highest priority system being assigned to ID 7. An iSeries or AS/400 will have the default SCSI ID 7.
  b. After the host system assignments are made, assign the highest priority volumes to the highest priority SCSI IDs. This is done automatically if you use the Add Fixed Block Volumes panel to create the highest priority volumes first. The highest priority volumes need to be at the top of the Volumes To Be Added list.
  c. If two host ports are connected to a common SCSI bus cable, arrange the SCSI IDs allocated to each port to give each port equivalent priority for volume access.
- For best host-bay performance, spread volumes assigned to host ports across device adapters in both clusters 1 and 2.
- For best host-port performance, assign only volumes from a common cluster (cluster 1 or cluster 2) to each port.

When you establish several connections to the ESS, first use one of the ports, of all available adapters. When there is a need for more connections, start using the second port. Be careful to spread all adapters evenly across the bays.

6.7.1 Volume access

Fibre Channel architecture allows any Fibre Channel initiator to access any open system device, without access restrictions. You can restrict this access when IBM sets the access mode for your ESS during initial configuration. You can select from the following options for access:

- **Access-any mode**: This mode allows all Fibre Channel-attached host systems, with no defined access profile, to access all logical volumes on the ESS. This applies only if you have not defined a host profile on ESS Specialist. After you define a profile for a host, that host has access only to those volumes that are assigned to the host’s WWPN.
- **Access-restricted mode**: This mode allows access only to host systems for which you have defined a profile. The profile limits the access of the host system to only those volumes assigned to the profile. This is the default mode. An access profile requires that:
  - The host system is identified to the ESS with a world-wide port number identity.
  - The logical volumes to which the host has access are identified to the ESS by associating the logical volumes with the access profile. Each WWPN can access 256 or 4,096 logical volumes, depending on the LUN addressing mode supported by the host operating system.
  - The iSeries supports up to 32 logical volumes per FC host adapter. The ESS Specialist updates the access profiles of a WWPN when you define the Fibre Channel host and assign volumes to the host.

Your IBM SSR can change the access mode on the ESS for you. Changing the access mode requires that you perform a power cycle (reboot your ESS).

6.7.2 Configuring fixed block disk groups

To configure a fixed block disk group, follow these steps:

1. At the bottom of the Open Systems Storage panel (Figure 6-11 on page 144), click **Configure Disk Groups**. This selection links you to the Fixed Block Storage panel (Figure 6-16).

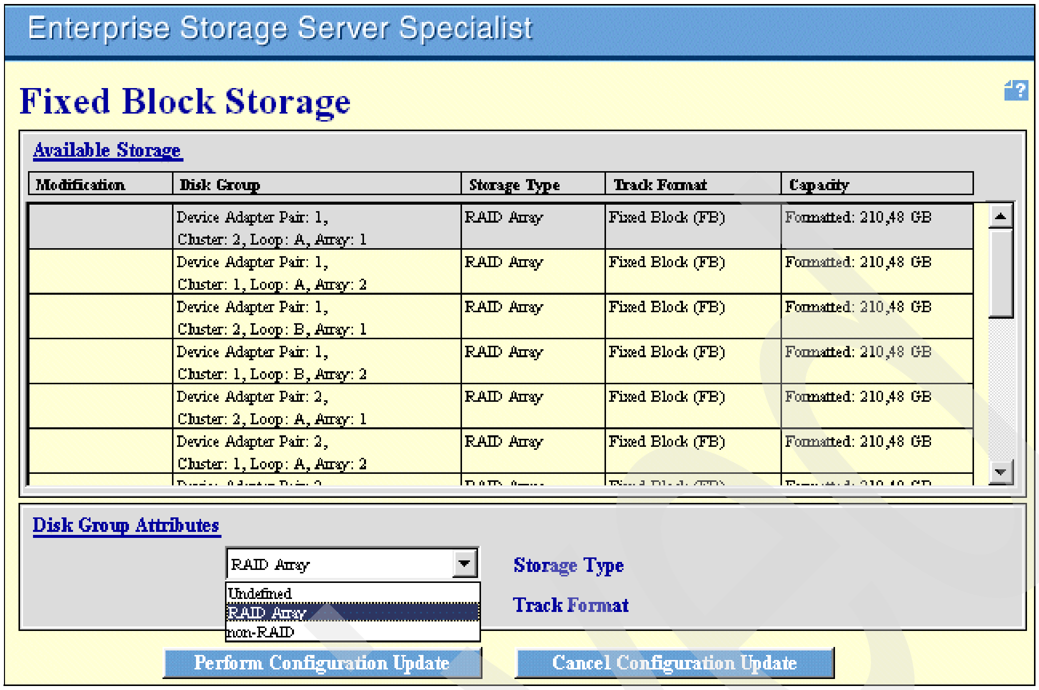

Figure 6-16 Fixed Block Storage panel

2. On the Fixed Block Storage panel, click a disk group in the list of disk groups to select it. See 6.7.3, “Assigning disk groups to logical subsystems (CE task)” on page 157, for information and recommendations on configuring disk groups.
3. Select one of the following options from the Storage Type drop-down list.
   – RAID array for iSeries, AS/400
   – Non-RAID disk group (not supported for iSeries)
   – Undefined

   If you select non-RAID, then the selected row expands to seven or eight rows one for each disk in the group.

   For non-RAID, click a disk row to define its track format as fixed block, or let it default to none. If you leave the default, you can define these disks at a later time to either count-key-data (CKD) or fixed block (FB).
4. Click each disk row to define its track format as **Fixed Block** (FB, for iSeries or AS/400) to define the disk, or select **None** (unused disk) to leave the disk undefined. You can define undefined disks at a later time to either fixed block (FB) track format (from this panel) or to count-key-data (CKD) track format (from the Configure CKD Disk Groups panel).
5. Repeat steps 2 through 4 as needed to define, undefine, or redefine the wanted disk groups. The modification column in the table indicates which table entries will be affected when you click **Perform** Configuration Update.
6. When all entries are complete, click **Perform Configuration Update** to apply the changes, or click **Cancel Configuration Update** to cancel the configuration step for the FB groups.

   Go to 6.7.4, “Adding volumes” on page 157, for instructions about how to convert array capacity for FB storage to one or more volumes.

6.7.3 Assigning disk groups to logical subsystems (CE task)

The default ESS Control Switch setting for Fixed Block LSSs is 16. Four LSSs are allocated to each Device Adapter pair, two of which are allocated to Cluster 1 and two to Cluster 2.

ESS Specialist assigns a disk group to an LSS when the disk group is defined as a RAID array or as a set of non-RAID disks. Even numbered disk groups are assigned to an LSS belonging to Cluster 1 and odd numbered disk groups are assigned to an LSS belonging to Cluster 2.

ESS Specialist assigns disk groups to only one of the two LSSs belonging to a cluster. It assigns disk groups to the second LSS only if the disk groups already assigned to the first LSS contain a total of 192 or more volumes. Each LSS can contain a maximum of 256 volumes.

For large numbers of high capacity disk groups with volume sizes of 8 GB or smaller, you may want to add volumes to half of the disk groups before you define the second half as RAID arrays or non-RAID disks. This avoids the possibility of using up all 256 LSS volumes before the disk group capacity is fully configured.

6.7.4 Adding volumes

Use the two Add Volumes panels to create logical volumes from the capacity on your storage server.

Selecting resources

The Add Volumes (1 of 2) panel (Figure 6-17) contains icons for all attached host systems along the top of the panel. Beneath these icons, you see the icons for the installed host adapters. Each SCSI adapter contains two selectable ports. A Fibre Channel adapter has one selectable port. Then beneath those icons, you see icons for the installed device adapters for Cluster 1 and Cluster 2.

To select resources for the logical volumes, perform the following steps:

1. Click **Add Volumes** on the Open Systems Storage panel (Figure 6-11 on page 144). ESS Specialist displays the Add Volumes (1 of 2) panel.

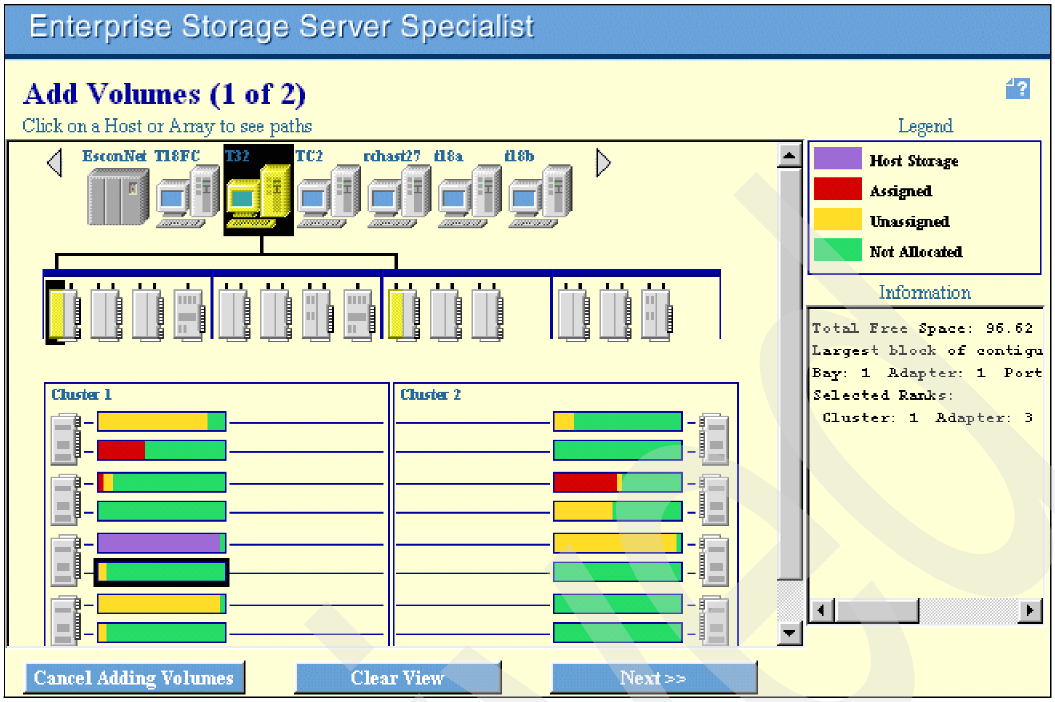

Figure 6-17 Add Volumes panel (1 of 2)

2. Click a host system in the first row to select it. You may select only one host system for an add volumes sequence. Click the **Clear View** button at the bottom of the panel at any time to start a new sequence.

   We recommend that you do not add volumes to an anonymous host. Anonymous host icons represent all Fibre Channel-attached host systems that are explicitly defined to the ESS. If you add volumes to the Anonymous host, the volumes are defined in specified storage, but no assignment to a particular host system is performed. If the storage server is configured to run in "Access Any" mode, then all Fibre Channel-attached hosts not explicitly defined to the ESS can access these volumes.

   The host system that you select is highlighted in yellow and marked with a black background to indicate selection. The name of the host system appears in the Information box at the bottom right of the panel. A thin black line is drawn from this host system to the one or more host ports in the second row that are attached to this system.

   Attention: Do not add a volume or modify a volume assignment on an active port. Some host systems react adversely to this action.

3. Click a highlighted port to select it. The background of the port you selected changes to black to indicate that it is selected.

   a. For SCSI-attached hosts, select only one port for the current sequence.

   b. For a Fibre Channel-attached host, click the port. If you defined more than one ESS FC host adapter, click any port. Volumes are assigned to the selected host, not to the port.

4. Click one or more FB groups (optional) to select them. Only these groups are used for the add-volume process that continues on the next panel. If no groups are specifically highlighted before you continue to the next panel, *all* available disk groups are used for the add-volume process. The Information box on the right side of the panel lists all RAID arrays and non-RAID disks that are currently selected.

5. Click **Next**. ESS Specialist displays the Add Fixed-Block Volumes Panel (2 of 2). See Figure 6-18.

Defining volume sizes and numbers

Use the Add Volumes (2 of 2) panel (Figure 6-18) to define volume sizes and numbers of volumes. See 6.8, “Recommendations for volume size and placement” on page 161, for more information on setting sizes.

> **Note:** We recommend that you do not mix AS/400 volumes within the same array with other open systems volumes.

The Free Storage Space Information at the top of the panel identifies the total free space available to allocate. It also indicates the largest possible volume that you may define, which is governed by the largest amount of free storage space within a single array or disk. This is due to the fact that a volume is not allowed to span multiple RAID arrays or non-RAID disks. It must be wholly contained in a single storage area.

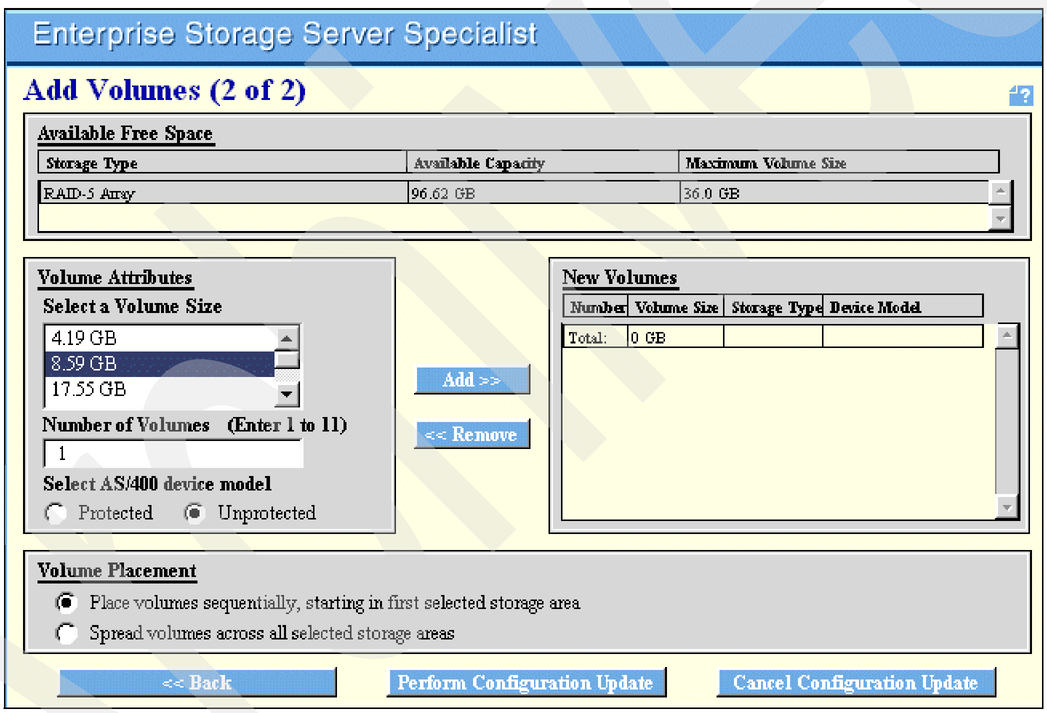

Figure 6-18 Add Volumes: Example of an unprotected AS/400 device model (2 of 2)

To define volume sizes, perform the following steps:

1. Under Available Free Space, click one ore more RAID-5 arrays (turn gray when selected).
2. Under Select a Volume Size, choose the volume size, in gigabytes (GB), that you want. You cannot select a volume size that exceeds the maximum size that may be configured in the selected space. You also cannot select a combined volume size and quantity that exceeds the free space of the selected space.
3. In the Number of Volumes box, select the number of volumes you want of the sizes you selected in step 1.

   For Select AS/400 device model, choose **Protected** or **Unprotected**. For example, choose unprotected if you are defining this device model as a remote load source device or as a device for mirroring any other drives. See Figure 6-18. Select protected (RAID) if you are defining non-mirrored Volumes for iSeries or AS/400. See Figure 6-19.

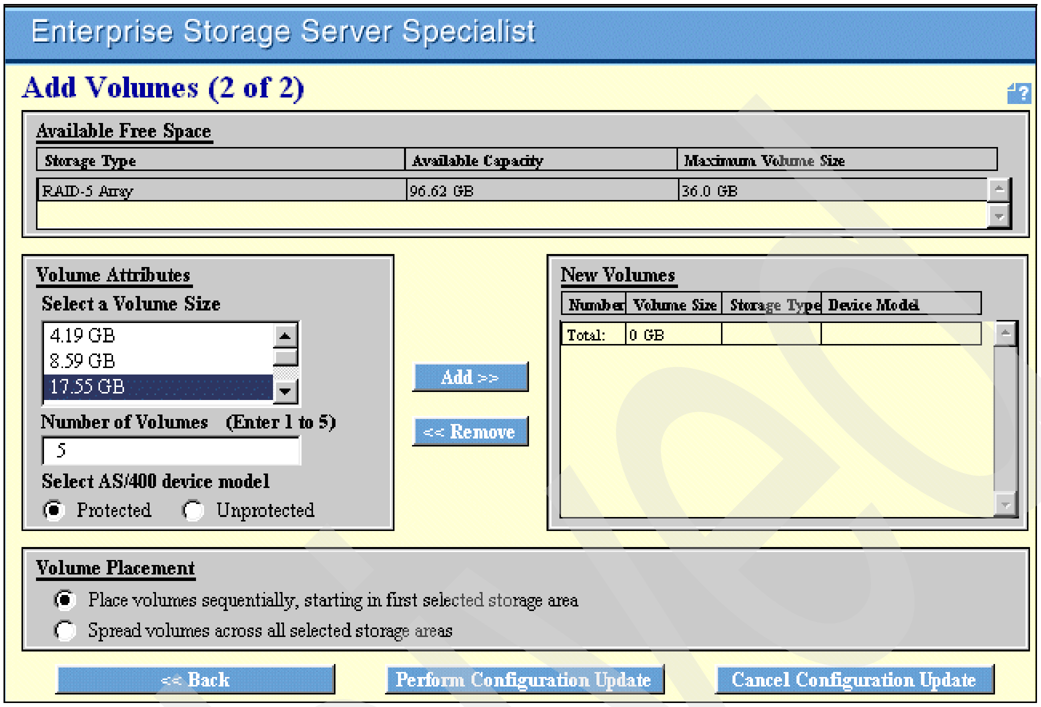

Figure 6-19 Add Volumes Example of a protected AS/400 device model (2 of 2)

4. Click **Add** to add the volumes to the New Volumes list. You may add more volume size or type selections to the list until you achieve the desired capacity for the selected port.
5. To remove a volume from the list, click the volume to select it and click **Remove**. The remaining capacity tabulations are updated for each add or remove action.
6. Select a volume placement option:
   – Click **Place volumes sequentially starting in the first selected storage area**, which is the default.
   – Click **Spread volumes across all selected storage area**, which option spreads the new volumes across the selected arrays and non-RAID disks.
7. When you are finished adding or deleting volumes, click **Perform** Configuration Update to apply the changes or **Cancel Configuration Update** to cancel the changes to the FB volumes.

When volumes are created and assigned to a SCSI port and FC port, the ESS initiates a formatting process that can exceed one hour for one or more volumes after the ESS is given the OK signal. A volume is not available to the host system while the ESS is formatting it.

Note: You cannot delete a volume. You *must* redefine the disk group.

Performance considerations

When adding volumes, consider these points in regard to performance:

- For SCSI-attached hosts, the volume at the top of the Volumes to be Added list is assigned to the highest priority SCSI ID that is available for new LUN assignments.

- Volumes are assigned to selected arrays according to a “first-fit” algorithm. This algorithm sometimes works best if the Volumes to be Added list is ordered with the largest capacity volumes first and the least capacity volumes last.
- Select the **Spread volumes across all selected storage areas** option if you selected disk groups attached to multiple SSA adapters. This optimally places the volumes for minimal access contention.
- For best host-bay performance, spread volumes assigned to host ports across device adapters in both clusters 1 and 2.
- For best host-port performance, assign only volumes from a common cluster (cluster 1 or cluster 2) to each port.

6.8 Recommendations for volume size and placement

This section offers IBM recommendations for volume sizes and placement. For information and recommendations on configuring disk groups, see 6.7.3, “Assigning disk groups to logical subsystems (CE task)” on page 157.

We recommend that you do not mix iSeries or AS/400 volumes within the same array with other open systems volumes. For example, the lowest volume number within an array (such as 000) has the best performance within an array. A bigger volume can be a better choice for operating systems that have the ability to divide a physical volume into smaller parts (partitions or slices). However, if you decide to use several smaller volumes, place them on several arrays to achieve better performance.

6.8.1 Modifying volume assignments

Use the Modify Volume Assignments panel to:

- Modify a volume assignment.
- Remove a volume from a target host. This does not affect the volume definitions, nor does it affect any of the data on the volume.
- Assign a volume that is assigned to a host to another SCSI host or to another Fibre Channel host. This allows shared access to a volume by two or more hosts. The iSeries or AS/400 does not support shared volumes.

A volume is not available to an attached-host system and is not displayed while formatting is in process. Depending on volume size, formatting may take one hour or more.

The Modify Volume Assignments panel has the following elements:

- A Volume Assignments Table of eight columns and multiple rows, with a minimum of one row for each defined volume. If a volume is assigned to a host, the row identifies the assigned host port and host system. If a volume is assigned to two or more hosts, the table contains an additional row for each additional host assignment.
- The table columns display the following information:
  - **Volume:** The Volume column contains the LUN serial number associated with this volume. The LUN serial number of an ESS volume incorporates a portion of the ESS serial number. This allows the LUN serial numbers to uniquely identify volumes within a particular ESS and volumes among all ESS machines installed in an enterprise.

    In addition, this column lists any volume formatting status. When a volume is first created, it goes through a one-time formatting process that must complete before the volume becomes available to the assigned host systems. If the formatting is still in progress, or has failed, the appropriate status is listed in this column. The format status for a volume can be one of the following types:

- Formatting (xx%) indicates that formatting is in progress. The percentage indicates approximately how much of the formatting is complete.
- Formatting (busy) indicates that formatting is in progress. The completion percentage cannot be obtained at this time. Wait a few minutes and then click the **Refresh Status** button to obtain the latest formatting status.
- Format failed (xx%) indicates that formatting did not complete successfully. Call your service representative to determine and correct the cause of the failure.
- Format error indicates that a hardware or microcode internal error caused the format to fail. Call your service representative to determine and correct the exact cause of the failure.

The ESS only updates this information when you click the Refresh button on the Modify Volume Assignment panel.

– **Location:** An identification of the physical location of the volume. This includes the number of the device adapter pair, the number of the cluster, the ID of the device loop, and the number of the disk group associated with the RAID array that contains the data for this volume. In case of a non-RAID volume, it also includes the number of the non-RAID disk within the disk group that contains the volume. In addition, the volume number is included to differentiate between different volumes on the same array or disk.

– **LSS:** The number of the logical subsystem that contains this volume. This value is important only for performing Copy Services operations on the volume. The ESS displays the LSS number as hexadecimal.

– **Volume Type:** The volume type can be either “AS/400” for volumes allocated for AS/400 host systems, or it can be “open system” for volumes allocated to UNIX-based host systems or Intel-based personal computers (PCs).

– **Size:** Displays the volume size in gigabytes. The value is rounded to two decimal places.

– **Storage Type:** The relative fault tolerance of the storage where the volume resides. The possible types are RAID Array (fault tolerant) and non-RAID (non-fault tolerant).

– **Host Port:** Provides the bay number and adapter information (SCSI or Fibre Channel). The SCSI ID is also displayed for SCSI ports.

– **Attention:** Do not add a volume or modify a volume assignment on an active port. Some host systems react adversely to this action.

– **Host Nicknames:** Displays the host system or host system nicknames that are associated with the volume and port number combination.

- The panel has three buttons in the upper-right corner:

  – Click **Perform Sort** after you use the drop-down list at the top of each column to set sort priorities. See 6.1.2, “Using ESS Specialist tables” on page 129, for details on sorting.

  – Click **Refresh Status** to update the autoformat information in the Volume column. This refresh can take several seconds to return and display updated information.

  – Click **Print Table** to capture and print the information on volume assignments. These buttons and the Sort button are in the upper right corner of this panel.

- The lower section of the panel includes:

  – The Action Box, which has two actions that you can select after you select a volume:

- The Assign selected volume(s) to target hosts option lists all eligible hosts in the Target Host box. Selecting this action enables the optional Use same ID/LUN in source and target action.

  Selecting the Use same ID/LUN in source and target keeps the same volume ID and LUN for source and target hosts. The display changes to show only hosts in the Target Host box that are compatible for sharing the selected volume.
- The Unassign selected volume(s) from target host option removes selected volume assignments. Removing a volume assignment does not affect the volume definition. The volume is not deleted; only its assignment to the target host is deleted.

– The Target Hosts box is populated when you select the Assign selected volume(s) to target hosts action. This box displays the host nickname and the host port or location (bay, card, and port).

To modify volume assignments, perform the following steps:

1. Select **Modify Volume Assignments** from the Open System Storage panel (Figure 6-11 on page 144). The ESS Specialist displays the Modify Volume Assignments panel (Figure 6-20).

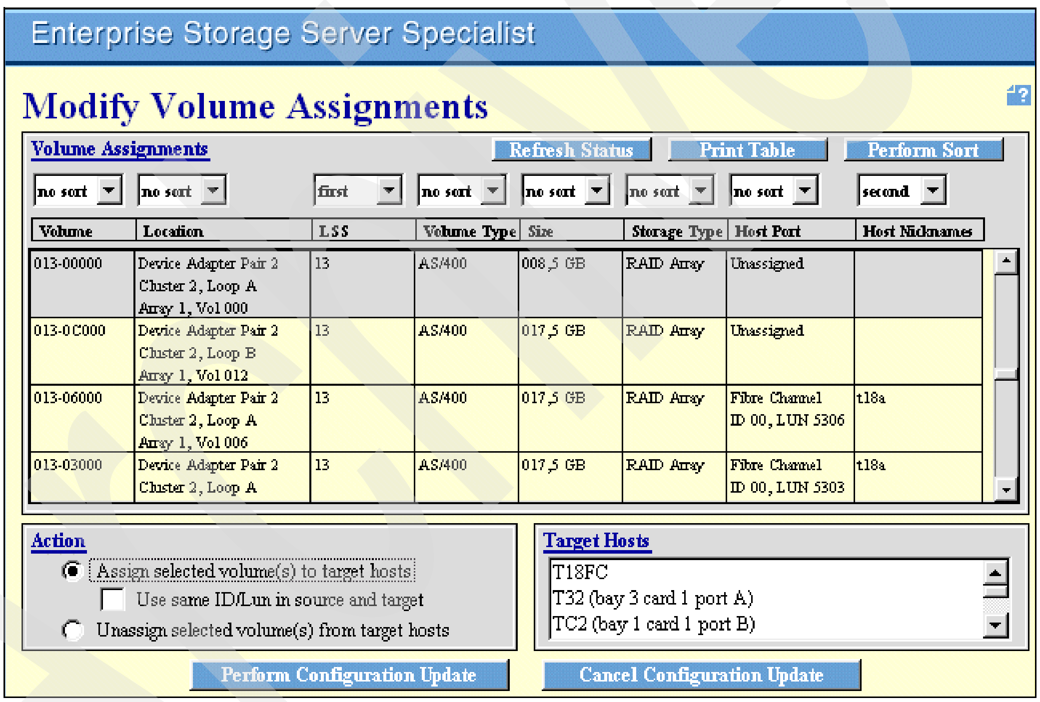

Figure 6-20 Modify Volume Assignments panel

2. Click one or more rows in the table to select the volumes you want to modify (see 6.1.2, "Using ESS Specialist tables" on page 129). For the iSeries or AS/400, click *all* rows in the table to select the volumes you want to modify.

   If you select a row a second time, it becomes deselected. When you select a row, the two action buttons in the Action box are enabled.
3. Select an action in the Action box. Click the **Assign selected volume(s)** to target hosts action. When you select this action, the following actions occur:

   a. The Target Host box is populated with all the hosts that are connected by SCSI or Fibre Channel adapters to the ESS. Use the scroll bar to see the complete list. Select a target (iSeries or AS/400 server) from the Target Host box.

b. The optional Use same ID/LUN in source and target action is enabled. Select this option to keep the same volume ID and LUN for source and target hosts. The display changes to show only hosts in the Target Host box that are compatible for sharing the selected volume.

4. Remove a volume assignment:

> **Attention:** For iSeries, you must power off the system if you plan to unassign removed units (units with a status of “non-configured units”).

a. Select the volume in the Volume Assignment table.

b. Click the **Unassign selected volume(s) from target hosts** option to remove the selected volume-to-host assignment.

c. Select one or more targets in the Target Hosts box.

5. When you finish modifying the volume assignments, click **Perform** Configuration Update to apply the modifications. Or click **Cancel Configuration Update** to cancel the modifications to the volume assignments. See Figure 6-21.

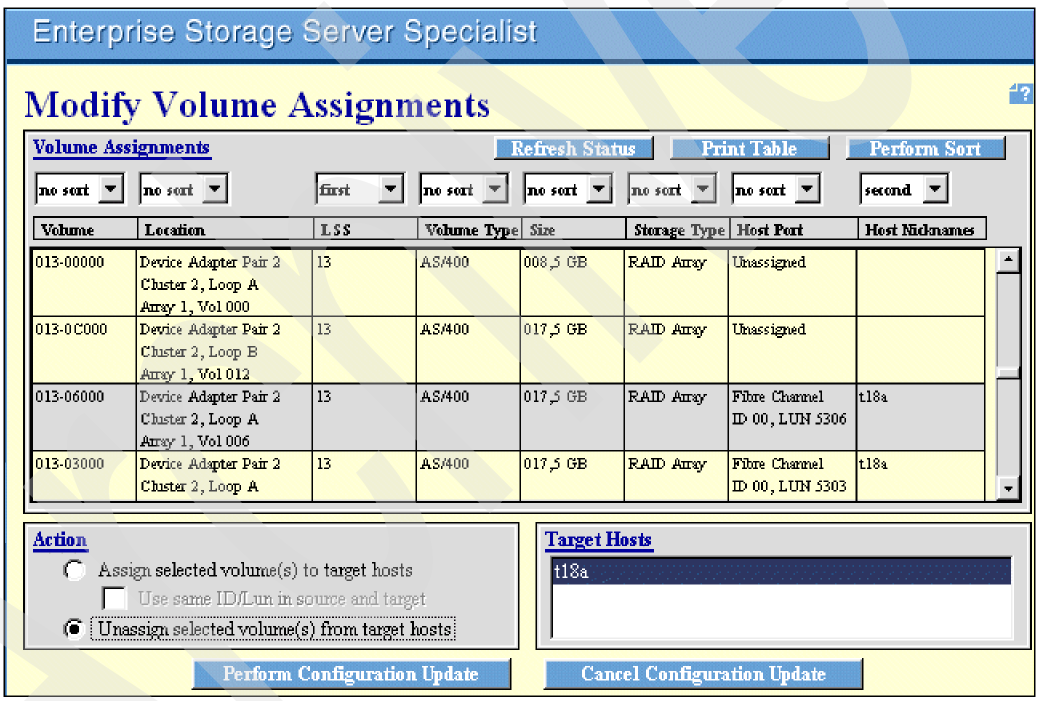

Figure 6-21 Modify Volume Assignments panel (example of unassigning a volume)

6.8.2 Signing off ESS Specialist and closing the browser window

The browser retains previously entered passwords in memory. Therefore, you must shut down the ESS Specialist and all browser windows completely to prevent unauthorized users from accessing the ESS Specialist. Access through an open browser window may give unauthorized users access to your IBM Enterprise Storage Server.

7

Mirroring to an ESS

Mirroring is an OS/400 facility that can increase levels of availability by maintaining a duplicate copy of data on separate disk units. In addition to this, further levels of redundancy can be provided by mirroring at the input/output processor (IOP), bus, or input/output (I/O) tower level.

This chapter looks at how you can use OS/400 mirroring in conjunction with the Enterprise Storage Server (ESS). First, it looks at setting up the ESS to allow mirroring and then shows you how to set up OS/400 for remote load source mirroring to support the ESS Copy Services PPRC and FlashCopy functions. Next, this chapter considers the case when using an ESS as a remote disk facility to provide an off-site mirrored copy of data that is controlled by OS/400 (not ESS Copy Services). Finally, this chapter looks at how to recover using a mirrored load source unit, irrespective of whether you were using remote load source mirroring with Copy Services or using OS/400 to mirror the entire local disk units into an ESS.

Until recently, mirroring was only available for internal disk units. Support has been added to OS/400 to allow internal disk, such as the load source unit (LSU), to be mirrored to external disk. This requires the external disk to report into OS/400 as unprotected, even though in practice, it is actually protected in a RAID-5 array within the Enterprise Storage Server.

OS/400 V5R1 and later releases include this support in the base System Licensed Internal Code (SLIC) and operating system.

We include PTF information for earlier releases, these releases are no longer supported, so obtaining these PTFs may be difficult. The facility for SCSI-attached external disk was added with the following program temporary fixes (PTFs):

- #6501 IOP PTFs:
  - V4R3: MF24552
  - V4R4: MF24558
  - V4R5: MF24559
- SLIC PTFs:
  - V4R3: MF24930
  - V4R4: MF24931
  - V4R5: MF24929

- OS/400 PTFs:
  - V4R3: SF57394
  - V4R4: SF57395
  - V4R5: SF62704

Before you implement mirroring to an ESS, check to see if these PTFs have been superseded. In addition to these PTFs, the Enterprise Storage Server must be at the GA driver level or higher. To use this support with Peer-to-Peer Remote Copy (PPRC)/Small Computer Systems Interface (SCSI) attach, refer to the ESS support site for more information regarding driver level support:

`http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/2105`

Table 7-1 shows the protected and unprotected model numbers for the possible logical unit number (LUN) sizes on the Enterprise Storage Server.

Table 7-1 LUN sizes and equivalent model numbers

Size (GB)	Interface	Type	Protected	Unprotected	Logical Block Addressing (decimal)
4.19	SCSI-3	9337	48C	48A	8,192,000
8.59	SCSI-3	9337	59C	59A	16,777,216
17.55	SCSI-3	9337	5AC	5AA	34,275,328
35.17	SCSI-3	9337	5CC	5CA	68,681,728
36.00	SCSI-3	9337	5BC	5BA	70,320,128
8.59	SCSI-FCP	2105	A01	A81	16,777,216
17.55	SCSI-FCP	2105	A02	A82	34,275,328
35.17	SCSI-FCP	2105	A05	A85	68,681,728
36.00	SCSI-FCP	2105	A03	A83	70,320,128
70.56	SCSI-FCP	2105	A04	A84	137,822,208

When mirroring to an ESS, normal OS/400 rules for mirroring apply. Both the source and target disk drives must be the same size, although they can be of different drive types and spin speeds. It is simply the capacity that must match.

7.1 Setting up the ESS for mirroring

To allow mirroring to the ESS, the target disk or disks must be unprotected, because OS/400 does not allow you to mirror any disk to another disk that is already protected. This must be done when you first define the LUN on the Enterprise Storage Server. Once a LUN is specified, you cannot change the designated protection. When using PPRC or FlashCopy, you only define the LUN for the LSU mirror as being unprotected. All other LUNs are defined as protected to reflect their true status to OS/400. With OS/400 Remote Disk Mirroring, you must define all LUNs in the ESS as unprotected.

To define the LUN or LUNs as “unprotected” on the ESS, refer to Chapter 6, “Creating iSeries storage in ESS” on page 127, and select **Unprotected** under Select AS/400 device model in the Volume Attributes box as shown in Figure 7-1.

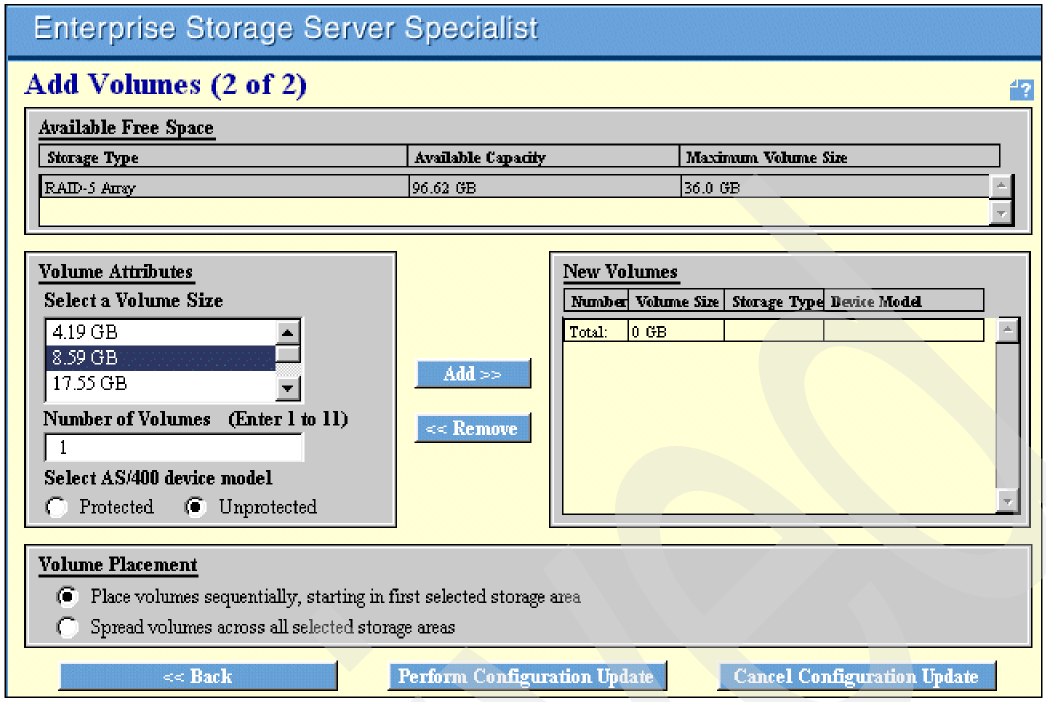

Figure 7-1 Defining an unprotected LUN

Once the unprotected LUNs are defined, the *only* place you can see this in the tabular view on StorWatch Specialist. It is a good idea to print the tabular view for the system you are setting as shown in Figure 7-2. It is important to remember the serial number of the unprotected LUN when using remote load source mirroring. Other LUNs do not matter because they are not important to SLIC. They are only assigned to single-level storage.

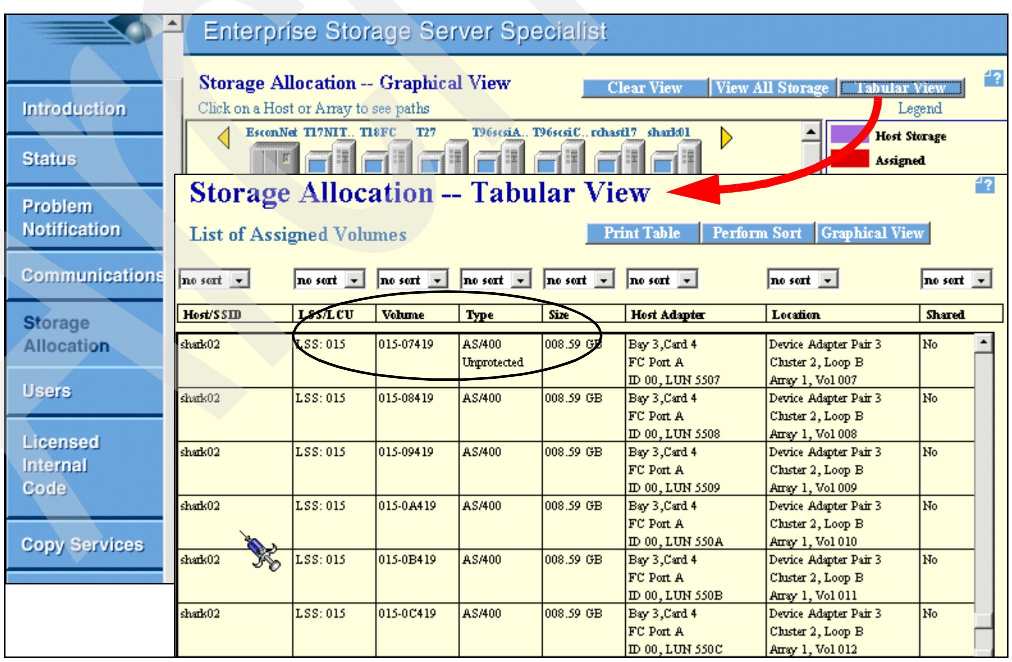

Figure 7-2 StorWatch tabular view showing an unprotected-type LUN

7.2 Remote load source mirroring

The load source unit is a special disk in the iSeries. This section relates to the implementation of remote load source mirroring function to maintain a copy of the LSU in the ESS to allow Copy Services (PPRC and FlashCopy) to provide protection for the entire AS/400 or iSeries disk capacity. It does not use Remote Disk Mirroring to protect the entire disk capacity. This is covered in 7.3, “Using OS/400 mirroring for data protection” on page 172.

You should recognize that using remote load source mirroring in this manner is not the same as when mirroring to internal drives. Although it uses OS/400 mirroring functions, the primary reason for using remote load source mirroring is to obtain a copy of the LSU into the Enterprise Storage Server so that the entire disk space in single-level storage can be duplicated by the hardware facilities, as done by PPRC and FlashCopy.

Note: In a configuration with a mirrored load source unit, the system can only IPL from, or perform a main storage dump to, the internal LSU. If the internal LSU fails, the system can continue to run on the other disk unit of the LSU mirrored pair but the system is not be able to IPL or perform a main storage dump until the LSU is repaired and usable. Therefore, if the system crashes after the internal LSU fails, the system cannot use the main storage dump to reduce recovery time nor to diagnose the problem that caused the system to end abnormally. This is not a restriction due to the Enterprise Storage Server. Rather it is due to OS/400 support for load source mirroring.

The LSU is the disk that is used to IPL the system (among other things) and is similar to a boot drive. All other “user data” can be located on external disk units, but the LSU must be an internal drive. This is because you cannot IPL the system from an I/O adapter (IOA) that supports external drives.

Due to the nature of iSeries single-level storage, it is necessary to consider the LSU as a special case. On other platforms, such as UNIX or Windows NT, each disk volume can be identified with its contents. The iSeries is different because all storage is considered as a single large address space. The LSU is within this address space. Therefore, if you want to use facilities such as PPRC or FlashCopy to perform a hardware copy of the disks attached to the iSeries, you *must mirror the LSU* from the internal drive into the Enterprise Storage Server to ensure the entire single-level storage is copied.

You must not have any other internal disks apart from the LSU when using PPRC or FlashCopy with OS/400. This is because the internal disks cannot be copied by PPRC or FlashCopy. It would negate the remote copy because not all single-level storage would be copied.

Once the LUN to be used as the remote LSU is created (see 7.1, “Setting up the ESS for mirroring” on page 166), this and any other LUNs that are allocated to your system are identified by SLIC and displayed under non-configured units in Dedicated Service Tools (DST) or SST.

You must then use DST to set up load source mirroring on the iSeries as explained in the following steps:

1. From the DST menu, select option 4 (Work with disk units) as shown in Figure 7-3.

```
                  Use Dedicated Service Tools (DST)
                                                    System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Work with Licensed Internal Code
     4. Work with disk units
     5. Work with DST environment
     6. Select DST console mode
     7. Start a service tool
     8. Perform automatic installation of the operating system
     9. Work with save storage and restore storage
    10. Work with remote service support


Selection
     4

F3=Exit   F12=Cancel
```

Figure 7-3 Use Dedicated Service Tools display

2. From the Work with Disk Units menu (Figure 7-4), select option 1 (Work with disk configuration).

```
           Work with Disk Units

Select one of the following:

     1. Work with disk configuration
     2. Work with disk unit recovery






Selection
     1

F3=Exit         F12=Cancel
```

Figure 7-4 Work with Disk Units display

3. From the Work with Disk Configuration menu (Figure 7-5), select option 4 (Work with mirrored protection).

```
        Work with Disk Configuration

 Select one of the following:

      1. Display disk configuration
      2. Work with ASP threshold
      3. Work with ASP configuration
      4. Work with mirrored protection
      5. Work with device parity protection
      6. Work with disk compression




Selection
     4

 F3=Exit         F12=Cancel
```

Figure 7-5 Work with Disk Configuration menu

4. From the Work with mirrored protection menu (Figure 7-6), select option 4 (Enable remote load source mirroring). This does not actually perform the remote load source mirroring but tells the system that you want to mirror the load source when mirroring is started.

```
         Work with mirrored protection

 Select one of the following:

      1. Display disk configuration
      2. Start mirrored protection
      3. Stop mirrored protection
      4. Enable remote load source mirroring
      5. Disable remote load source mirroring




 Selection
      4

 F3=Exit         F12=Cancel
```

Figure 7-6 Setting up remote load source mirroring

5. You see the Enable Remote Load Source Mirroring confirmation display (Figure 7-7). Press Enter to confirm that you want to enable remote load source mirroring.

```
      Enable Remote Load Source Mirroring

 Remote load source mirroring will allow you to place the two
units that make up a mirrored load source disk unit (unit 1) on
two different IOPs.  This may allow for higher availability
if there is a failure on the multifunction IOP.


Note: When there is only one load source disk unit attached to
the multifunction IOP, the system will not be able to IPL if
that unit should fail.


This function will not start mirrored protection.

Press Enter to enable remote load source mirroring.
```

Figure 7-7   Enable Load Source Mirroring confirmation display

6. The Work with mirrored protection menu (Figure 7-8) is displayed. It shows a message, indicating that remote load source mirroring enabled. Select option 2 to start mirrored protection of the LSU.

```
         Work with mirrored protection

 Select one of the following:

      1. Display disk configuration
      2. Start mirrored protection
      3. Stop mirrored protection
      4. Enable remote load source mirroring
      5. Disable remote load source mirroring




 Selection
      2

 F3=Exit         F12=Cancel
   Remote load source mirroring enabled successfully.
```

Figure 7-8   Confirmation that remote load source mirroring is enabled

7. On the Work with mirrored protection menu, select option 1 (Display disk configuration) and on the next display, select option 1 (Display disk configuration status).

   The Display Disk Configuration Status display (Figure 7-9) shows the internal 6717-050 LSU as disk serial number 68-09011 and our target LUN in the ESS is shown as 9337-59A, serial number 75-1044000. Unprotected Fibre Channel attached LUNs would be shown as 2105-A8x, depending on LUN size.

   Press Enter to continue.

```
         Display Disk Configuration Status

          Serial                Resource
ASP  Unit  Number      Type Model Name        Status
  1                                           Mirrored
       1  68-09011    6717  050  DD023       Unprotected
       2  75-1044000  9337  59A  DD037       Unprotected
       3  75-2044000  9337  59C  DD038       DPY/Active
       4  75-3044000  9337  59C  DD039       DPY/Active
       5  75-4044000  9337  59C  DD040       DPY/Active
       6  75-5044000  9337  59C  DD041       DPY/Active
       7  75-6044000  9337  59C  DD030       DPY/Active
       8  75-7044000  9337  59C  DD031       DPY/Active







 Press Enter to continue.

 F3=Exit         F5-Refresh        F9-Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 7-9 Unprotected LSU ready to start remote load source mirroring

8. When the remote load source mirroring task is finished, you must IPL the system to mirror the data from the source unit to the target. This is done during the database recovery phase of the IPL.

Starting the mirroring process can be a long running job, depending on how many disk units are being mirrored.

7.3 Using OS/400 mirroring for data protection

As well as using OS/400 mirroring to mirror the LSU when using Copy Services, you may also use mirroring in a more “standard” way. OS/400 mirroring is designed to provide high levels of disk protection. Disk-, IOP-, and bus/tower-level protection are available. If you have an ESS attached to your iSeries or AS/400, you can mirror the internal drives to the LUNs in the ESS. This is especially useful if the ESS is situated remotely from the system unit.

Figure 7-10 shows an example of this. Note that it is also possible to connect two hubs via long wave fibre. This allows the remote ESS to be situated up to 10km away from the system unit and internal disk, providing an off-site remote ESS.

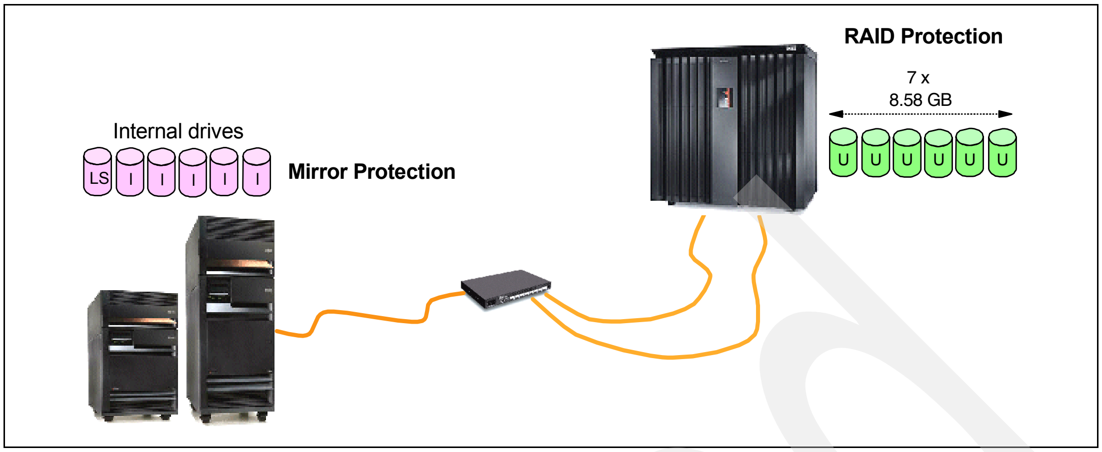

Figure 7-10   Internal disks mirrored to ESS

Since OS/400 treats ESS disk in exactly the same way as it does internal disk, you can also use this concept if you have a local ESS and a remote ESS as shown in Figure 7-11.

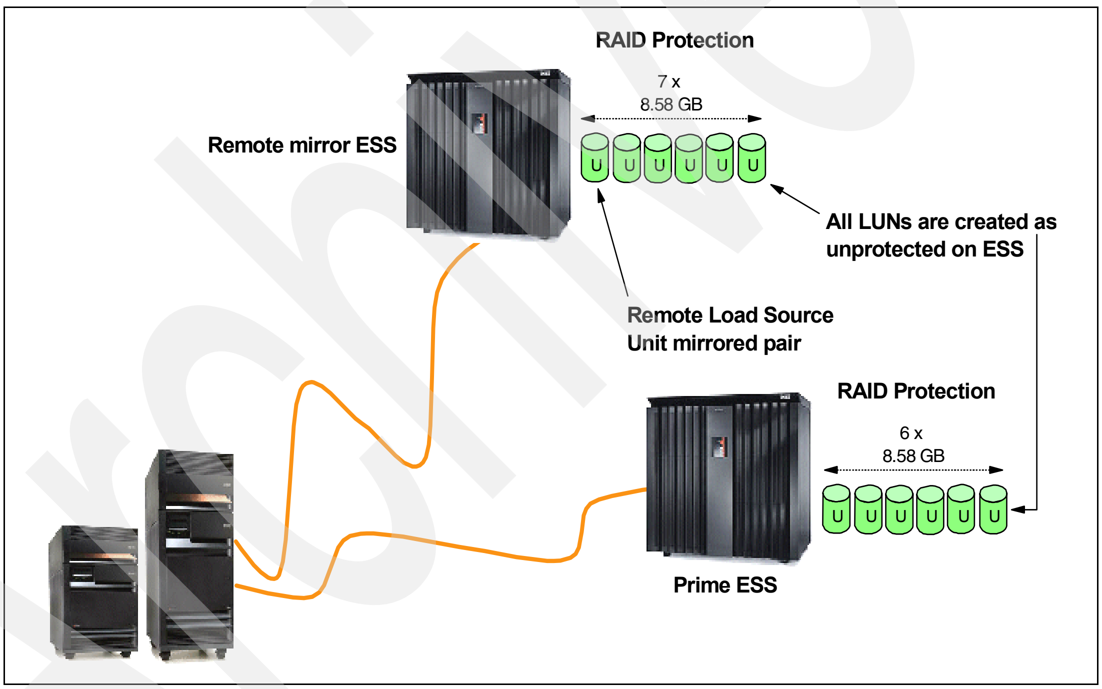

Figure 7-11   Local ESS mirrored to remote ESS

7.3.1 Setting up ESS mirroring

The tasks necessary to set up this remote mirroring are the same for both examples, whether you have internal disks or local ESS disks that are to be mirrored to a remote ESS. To implement OS/400 mirroring, you must have the same size disks for both “halves” of the mirror, and you must have an even number of disk units. The model, type, or spin-speed is not relevant. You must define the ESS LUNs as unprotected (even though they are really protected by RAID-5 in ESS). Otherwise OS/400 does not recognize them as candidates for mirroring. In this case, all the ESS disks should be defined as unprotected.

To ensure that local disks are not mirrored to other local disks or that the remote ESS LUNs are not mirrored to other LUNs in the same ESS, you need to “tag” the buses that control the remote disks. This involves renaming the bus controlling the “remote” disks. When mirroring is started, OS/400 attempts to mirror to the highest possible level (tower, then bus, then IOP, and finally disk unit). When OS/400 sees a bus that begins with an “R”, it assumes this is a remote bus. By default, local buses are named “Lxxx”.

Renaming a bus

Hardware resource names can be changed in System Service Tools (SST). If you do not have sufficient authority to use SST, see your supervisor. Follow these steps to rename a bus:

1. On the System Service Tools main menu (Figure 7-12), select option 1 (Start a service tool).

```
                         System Service Tools (SST)

 Select one of the following:

      1. Start a service tool
      2. Work with active service tools
      3. Work with disk units
      4. Work with diskette data recovery
      5. Work with system partitions
      6. Work with system capacity









 Selection
      1

 F3=Exit         F10=Command entry         F12=Cancel
```

Figure 7-12 Starting a service tool

2. The Start a Service Tool menu (Figure 7-13) is displayed. It shows the many service tools that are available. We need to work with the Hardware Service Manager, so choose option 7 (Hardware service manager).

   Attention: SST is a powerful tool, and you should use it with caution. Do not use SST if you are not familiar with the facilities provided. See the warning on the Start a Service Tool display shown in Figure 7-13.

```
                           Start a Service Tool

Warning: Incorrect use of this service tool can cause damage
 to data in this system.  Contact your service representative
for assistance.

Select one of the following:

     1. Product activity log
     2. Trace Licensed Internal Code
     3. Work with communications trace
     4. Display/Alter/Dump
     5. Licensed Internal Code log
     6. Main storage dump manager
     7. Hardware service manager

Selection
     7

F3=Exit         F12=Cancel         F16=SST menu
```

Figure 7-13 Starting the Hardware Service Manager

3. The Hardware Service Manager display (Figure 7-14) appears. This tool provides many facilities that can severely impact your system. We are only interested in changing the resource details. If you know the name of the resource you want to change, select option 3 (Locate resource by resource name). If you are not sure about the name, you can view all resources and choose the correct one from those shown. To do this, select option 2 (Logical hardware resources) as we do in this example.

```
                            Hardware Service Manager

Attention: This utility is provided for service representative use only.

   System unit . . . . . . . :   9406-820 10-5GG8M
   Release . . . . . . . . . :   V5R1M0 (1)

Select one of the following:
    1. Packaging hardware resources (systems, frames, cards,...)
    2. Logical hardware resources (buses, IOPs, controllers,...)
    3. Locate resource by resource name
    4. Failed and non-reporting hardware resources
    5. System power control network (SPCN)
    6. Work with service action log
    7. Display label location work sheet
    8. Device Concurrent Maintenance

Selection
     2

F3=Exit      F6=Print configuration      F9=Display card gap information
F10=Display resources requiring attention      F12=Cancel
```

Figure 7-14 Hardware Service Manager display

4. The Logical Hardware Resources display (Figure 7-15) appears. The resources we want to work with are on the system bus. Select option 1 (System bus resources).

```
                           Logical Hardware Resources


Select one of the following:

     1. System bus resources
     2. Processor resources
     3. Main storage resources
     4. High-speed link resources








Selection
     1

F3=Exit      F6=Print configuration      F12=Cancel
```

Figure 7-15 Selecting the system bus resources

5. You see the Logical Hardware Resources on System Bus display (Figure 7-16), which shows the logical hardware resources. Scroll down this list until you see the IOP that supports your tape adapter and tape devices.

   In the example, two system buses are shown, LB09 and LB10. Use option 2 (Change detail) to select the bus you want to "tag" as remote.

```
                Logical Hardware Resources on System Bus

System bus(es) to work with . . . . . .   *ALL  *ALL, *SPD, *PCI, 1-511
Subset by . . . . . . . . . . . . . . .   *ALL  *ALL, *STG, *WS, *CMN, *CRP

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=I/O Debug
  7=Display system information
  8=Associated packaging resource(s)     9=Resources associated with IOP

                                                                Resource
Opt Description                   Type-Model  Status            Name
      System Bus                   28AB-      Operational       LB09
       Multi-adapter Bridge        28AB-      Operational       PCI10D
     Bus Expansion Adapter             -      Operational       BCC10
  2   System Bus                   28AB-      Operational       LB10
       Multi-adapter Bridge        28AB-      Operational       PCI05D
        Combined Function IOP      2843-001   Operational       CMB05
       Multi-adapter Bridge        28AB-      Operational       PCI12D
                                                                        More...
F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
F9=Failed resources      F10=Non-reporting resources
F11=Display serial/part numbers      F12=Cancel
```

Figure 7-16 Selecting the bus to be renamed

6. The Change Logical Hardware Resource Detail display (Figure 7-17) appears. Local buses are named LB*xx* by default. Type the new name for this resource and press Enter to change the resource name. We recommend that you rename the remote buses as `RBxx`.

```
                 Change Logical Hardware Resource Detail

Description . . . . . . . . . . . . :  System Bus
Type-Model  . . . . . . . . . . . . :  28AB
Status  . . . . . . . . . . . . . . :  Operational
Serial number . . . . . . . . . . . :  B8-1010009

Current resource name . . . . . . . :  LB01

Type changes, press Enter.

New resource name . . . . . . . . .    RB01

F3=Exit      F5=Refresh      F6=Print
F9=Display detail      F12=Cancel
```

Figure 7-17 Changing the resource name

7. Press F3 (Exit) to return to the Logical Hardware Resource Associated with IOP display. You should see the new name of the resource displayed as shown in Figure 7-18.

```
                  Logical Hardware Resources on System Bus

 System bus(es) to work with . . . . . .   *ALL  *ALL, *SPD, *PCI, 1-511
 Subset by . . . . . . . . . . . . . . .   *ALL  *ALL, *STG, *WS, *CMN, *CRP

 Type options, press Enter.
   2=Change detail    4=Remove    5=Display detail    6=I/O Debug
   7=Display system information
   8=Associated packaging resource(s)     9=Resources associated with IOP

                                                              Resource
 Opt Description                   Type-Model  Status          Name
       System Bus                   28AB-      Operational     LB09
        Multi-adapter Bridge        28AB-      Operational     PCI10D
      Bus Expansion Adapter             -      Operational     BCC10
       System Bus                   28AB-      Operational     RB10
        Multi-adapter Bridge        28AB-      Operational     PCI05D
         Combined Function IOP      2843-001   Operational     CMB05
        Multi-adapter Bridge        28AB-      Operational     PCI12D
                                                                      More...
 F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
 F9=Failed resources      F10=Non-reporting resources
 F11=Display serial/part numbers      F12=Cancel
```

Figure 7-18 LB01 renamed to RB10

Ensure that the load source unit is mirrored remotely. Now, set up remote load source mirroring and start mirroring in the normal way as described in 7.1, "Setting up the ESS for mirroring" on page 166. Use option 2 (Start mirrored protection) on the Work with mirrored protection display to perform the mirroring function. This pairs up the disks on local buses (LBxx) with disks on remote buses (RB01).

Important: When you configure the hardware for remote mirroring, you must ensure that the internal or local disks are on different buses from the remote disks. Failure to ensure this can result in some local disks being mirrored to other local disks or remote ESS disks being mirrored to other remote ESS disks. In some circumstances, this may require additional hardware to ensure that the IOP/IOA for the remote ESS is on a different bus to the local disk units.

When using OS/400 mirroring with "remote external disk", such as a remote Enterprise Storage Server, consider this. Reads are issued only to a single unit in a mirrored pair. The unit with the shortest read I/O queue is selected. There is a concept of a "preferred unit" in case of a tie (as is likely during low activity). The preferred unit is initially the "B" unit in an "A" and "B" mirrored pair. The "B" unit is the unit that had the least data on it when mirroring was originally started.

From a mirroring standpoint, reads should be somewhat dynamically self-governing in that, if read responses from the remote ESS are slower than the local drives, the I/O queues to the remote ESS should become longer than queues to the local drives. Also, more reads are directed to the local drives. In any case, the mirroring algorithm distributes the reads accordingly.

It is likely that when mirroring is first started, the ESS drives are empty and, therefore, become the “B” unit or “preferred unit”. Under lighter loads, where the read I/O queue length is zero to both drives, the reads then goes to the ESS. This should not be a concern considering:

- The lighter load
- Near equal distribution (to the **remote ESS** and local units) under heavier loads
- The ESS read cache

In this case, if you have the opportunity, it may be better to add the remote drives to the configuration and populate them with data before you add the local drives and start mirroring. This holds especially true if the two local drives have significantly different characteristics such as internal drives or ESS with more, smaller physical drives. However, in many cases, the local drives may be in place and populated with data before the remote ESS is added to the configuration. Therefore, it would not be worthwhile (or even practical) to populate the remote drives first.

Given this, it seems that, while this can potentially introduce performance variability at times depending on the nature of the I/O activity, the overall effect of the larger drives may be negligible.

7.4 Recovering the load source unit

To use the disks in the Enterprise Storage Server on another system (for example, accessing the remote copies created by PPRC and FlashCopy or recovering the entire disk configuration from the OS/400 remote mirror), you must first recover the load source unit back onto an internal disk in the iSeries. Chapter 12, “Peer-to-Peer Remote Copy” on page 327, and Chapter 13, “FlashCopy” on page 367, look in more detail at how PPRC and FlashCopy are implemented.

To recover the load source unit, it is necessary to perform a D-IPL and do a scratch installation as explained in the following sections.

7.4.1 Performing a D-IPL and reloading SLIC

We advise that you do this from tape (either a SAVSYS, a distribution tape, or a tape used to save the Licensed Internal Code (LIC) in DST) rather than use the distribution CD because it does not have any PTFs loaded. Load the tape (or CD) into the alternate IPL device and make sure the device is ready before you perform the D-IPL. If you are not sure how to perform a D-IPL, refer to *Backup and Recovery Version 5*, SC41-5304.

1. The Select a Language Group display may appear. If so, select your country (region) code (for example, `2924` for US English) and press Enter.
2. When the Install Licensed Internal Code menu (Figure 7-19) is displayed, select option 1 (Install Licensed Internal Code).

```
                   Install Licensed Internal Code
                                                       System:  S101880D
Select one of the following:

     1. Install Licensed Internal Code
     2. Work with Dedicated Service Tools (DST)
     3. Define alternate installation device




Selection

     1
Licensed Internal Code - Property of IBM 5769-999 Licensed
Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
rights reserved.  US Government Users Restricted Rights -
Use duplication or disclosure restricted by GSA ADP schedule
Contract with IBM Corp.

Tape unit available
```

Figure 7-19   Beginning to install LIC

3. At this stage, you may be prompted to specify a device from which to load the LIC on the Work with Tape Devices display (Figure 7-20). However, this display may not always appear. Type option 1 next to the tape drive to select it and press Enter.

```
                          Work with Tape Devices
                                                       System:  S101880D
Type option, press Enter
  1= Select  2=Deselct    5=Display details

            Resource                          Serial
Option      Name            Type      Model     Number          Selected
  1         TAP01           6380      001       00-3088792






F3=Exit F5=Refresh    F12=Cancel
Make the device ready before making a selection
```

Figure 7-20   Selecting a tape drive to load the LIC

4. Because we are planning to recover the LSU from the remote mirror copy, we want to re-initialize the LSU before we install the LIC. To do this, select option 2 (Install Licensed Internal Code and Initialize system) from the Install Licensed Internal Code (LIC) display (Figure 7-21).

```
                 Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
    Serial Number     Type     Model     I/O Bus    Controller    Device
     00-CA71E         6713      030         0           1           0

Select one of the following:

     1. Restore Licensed Internal Code
     2. Install Licensed Internal Code and Initialize system
     3. Install Licensed Internal Code and Recover Configuration
     4. Install Licensed Internal Code and Restore Disk Unit Data
     5. Install Licensed Internal Code and Upgrade Load Source



Selection
     2

F3=Exit          F12=Cancel
```

Figure 7-21 Selecting to install LIC and initialize the system

5. When using ESS Copy Services (PPRC or FlashCopy) or recovering from a remote mirrored ESS, we recover the "remote mirrored" disk to a different system. You should see a warning display similar to the example in Figure 7-22. The actual text of the warning may differ depending on your particular circumstances (for example, logical partition (LPAR) first or subsequent recovery of remote load source).

   Press Enter to continue.

```
                    Install Licensed Internal Code - Warning

Disk selected to write the Licensed Internal Code to:
    Serial Number     Type     Model     I/O Bus    Controller    Device
     00-CA71E         6713      030         0           1           0

Warning:
    The load source disk and its mirrored pair could not be found
    (see disk information below).


Missing load source disk:
    Serial Number     Type     Model     I/O Bus    Controller    Device
     75-2040000       9337      59A         1           1           1


Press Enter to continue the restore or install on the selected
disk.

F3=Exit          F12=Cancel
```

Figure 7-22 Missing load source unit

6. You may see a confirmation display as in Figure 7-23. If so, press F10 to continue the installation process.

```
          Install LIC and Initialize System - Confirmation

 Warning:
     All data on this system will be destroyed and the Licensed
     Internal Code will be written to the selected disk if you
     choose to continue the initialize and install.

     Return to the install selection screen and choose one of the
     other options if you want to perform some type of recovery
     after the install of the Licensed Internal Code is complete.



     Press F10 to continue the install.
     Press F12 (Cancel) to return to the previous screen.
     Press F3 (Exit) to return to the install selection screen.



F3=Exit          F12=Cancel
```

Figure 7-23 Confirming the initialization of LSU and install LIC

7. The new load source unit is being initialized. The status display in Figure 7-24 shows an estimate of the time to initialize the new LSU and the progress.

```
             Initialize the Disk - Status

 The load source disk is being initialized.


 Estimated time to initialize in minutes :    26

 Elapsed time in minutes . . . . . . . . :    14.5




 Please wait.


 Wait for next display or press F16 for DST main menu
```

Figure 7-24 Initializing the disk status

When the LSU is initialized, you see the Install Licensed Internal Code - Status display (Figure 7-25).

```
            Install Licensed Internal Code - Status

Install of the licensed internal code in progress.

            +---------------------------------------------------+
Percent     |             20%                                   |
complete    +---------------------------------------------------+


Elapsed time in minutes  . . . . . . . :      6.5




Please wait.


Wait for next display or press F16 for DST main menu
```

Figure 7-25 Install Licensed Internal Code - Status

After the LIC is installed, the system begins to IPL.

If you have a non-LPAR system, you are notified that there is a different disk configuration (Figure 7-26). However, if you have an LPAR system, you are warned that there is invalid LPAR configuration information on a non-reporting disk. See 7.4.3, “Recovery on an LPAR system” on page 186, to follow the steps to recover from this.

7.4.2 Recovery on a non-LPAR system

To perform recovery on a non-LPAR system, follow the steps as explained here:

1. On the Disk Configuration Attention Report (Figure 7-26), select option 5 (Display Detailed Report).

> **Attention:** *Do not* press F10 to accept the error. You should check the details of the error.

```
               Disk Configuration Attention Report

Type option, press Enter.
  5=Display Detailed Report

Press F10 to accept all the following problems and continue.
The system will attempt to correct them.

Opt  Problem
 5   New disk configuration





F3=Exit   F10=Accept the problems and continue   F12=Cancel
```

Figure 7-26 New Disk Configuration Attention Report

2. There are several possible causes for the problem as shown in the Accept New Disk Configuration report (Figure 7-27 and Figure 7-28). Proceed through these displays until you reach the Disk Configuration Attention Report.

```
                 Accept New Disk Configuration

 The current configuration indicates a single unit system.  You
 can choose to accept it or do one of the following:

 Following are the possible causes and recovery procedures:

   o  You can define a new configuration by adding units.

   o  Press F3 to exit to Dedicated Service Tools (DST) and if
      necessary, take the right option to get to
      'Use Dedicated Service Tools (DST)' display.
      On the 'Use Dedicated Service Tools (DST)' display,
       -  Select option 4, Work with disk units.
       -  Select option 1, Work with disk configuration.
       -  Select option 3, Work with ASP configuration.
       -  Select option 3, Add units to ASPs.

   o  If you are performing 'Recover mirror load source'
      utility, press F3 to exit to Dedicated Service Tools (DST).
      If necessary, take the right option to get to 'Use
                                                                 More...
 F3=Exit                                        F12=Cancel
```

Figure 7-27 Accept New Disk Configuration (Part 1 of 2)

```
                        Accept New Disk Configuration

      Dedicated Service Tools (DST)' display.
      On the 'Use Dedicated Service Tools (DST)' display
       -  Select option 4, Work with disk units.
       -  Select option 2, Work with disk unit recovery.
       -  Select option 16, Recover mirrored load source.

 Press Enter to accept the current configuration and continue.







                                                                  Bottom
 F3=Exit                                        F12=Cancel
```

Figure 7-28 Accept New Disk Configuration (Part 2 of 2)

> **Attention:** We are performing a “Recover mirror load source” task. The report indicates that you should press F3 to Exit from the New Disk Configuration display and return to DST. Do *not* press F10 to accept the new configuration.

Press F3 (exit) on the Disk Configuration Attention Report (Figure 7-29).

```
                    Disk Configuration Attention Report

 Type option, press Enter.
   5=Display Detailed Report

 Press F10 to accept all the following problems and continue.
 The system will attempt to correct them.

 Opt  Problem
      New disk configuration







 F3=Exit   F10=Accept the problems and continue   F12=Cancel
```

Figure 7-29 Exiting the Disk Configuration Attention Report

3. On the IPL or Install the System display (Figure 7-30), select option 3 (Use Dedicated Service Tools (DST)).

```
                       IPL or Install the System
                                                       System:  S101880D
 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Use Dedicated Service Tools (DST)
      4. Perform automatic installation of the operating system
      5. Save Licensed Internal Code



 Selection
      3

 Licensed Internal Code - Property of IBM 5769-999 Licensed
 Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
 rights reserved.  US Government Users Restricted Rights -
 Use duplication or disclosure restricted by GSA ADP schedule
 Contract with IBM Corp.
```

Figure 7-30 IPL or Install the System display

4. Sign on to DST using an acceptable user ID and password as shown in Figure 7-31.

> **Tip:** If you are using a distribution or SAVSYS tape, the password is the same as it was when the tape was created. If you are using the distribution media, the default password for *qsecofr* is QSECOFR (uppercase).

```
                    Dedicated Service Tools (DST) Sign On
                                                          System:  S101880D
 Type choices, press Enter.

   DST user . . . . . . . . . . . .   qsecofr
   DST password . . . . . . . . . .











 F3=Exit   F12=Cancel
 Cannot access user profiles - defaults are in effect
```

Figure 7-31 Signing on to DST

You have now completed the tasks specific to a non-LPAR environment. Now go to 7.4.4, “Continuing to the recover the load source” on page 190.

7.4.3 Recovery on an LPAR system

Use the following steps to correct LPAR configuration errors found when recovering your remote load source:

1. You see a Disk Configuration Error Report as shown in Figure 7-32. You may choose option 5 to view the error report or simply press F12 (Cancel) to return to the IPL or Install the System menu.

```
              Disk Configuration Error Report

 Type option, press Enter.
   5=Display Detailed Report

 Opt  Error
      Unit has incorrect logical partition configuration









 F3=Exit       F12=Cancel
```

Figure 7-32 Incorrect LPAR configuration warning

2. We need to work with Dedicated Service Tools to modify the LPAR configuration details. Select option 3 (Use Dedicated Service Tools (DST)) on the IPL or Install the System display (Figure 7-33).

```
                       IPL or Install the System
                                                        System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Use Dedicated Service Tools (DST)
     4. Perform automatic installation of the operating system
     5. Save Licensed Internal Code




Selection
     3

Licensed Internal Code - Property of IBM 5769-999 Licensed
Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
rights reserved.  US Government Users Restricted Rights -
Use duplication or disclosure restricted by GSA ADP schedule
Contract with IBM Corp.d
```

Figure 7-33 IPL or Install the System display

3. Sign on to DST using an acceptable user ID and password as shown in Figure 7-34.

> **Tip:** If you are using a distribution or SAVSYS tape, the password is the same as it was when the tape was created. If you are using the distribution media, the default password for *qsecofr* is QSECOFR (uppercase).

```
              Dedicated Service Tools (DST) Sign On
                                                        System:   S101880D
Type choices, press Enter.

  DST user . . . . . . . . . . . .   qsecofr
  DST password . . . . . . . . . .











F3=Exit   F12=Cancel
Cannot access user profiles - defaults are in effect
```

Figure 7-34 Signing on to DST

4. When the Use Dedicated Service Tools (DST) menu (Figure 7-35) is displayed, discard the LPAR configuration found on the remote load source unit. To do so, select option 11 (Work with system partitions).

```
                      Use Dedicated Service Tools (DST)
                                                         System:  S101880D
 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system
      9. Work with save storage and restore storage
     10. Work with remote service support
     11. Work with system partitions


 Selection
      11
```

Figure 7-35 Using DST

5. This takes you to the Work with system partitions menu (Figure 7-36). We want to work with the LPAR configuration data. Select option 4 (Recover configuration data).

```
                       Work with System Partitions
                                                         System:  S101880D
 Attention:  Incorrect of this utility can cause damage
to data in this system.  See service documentation


  Number of partitions  . . . . . . . :  4
  Partition manager release . . . . . :  V5R1M0 L000

  Partition identifier  . . . . . . . :  2
  Partition name  . . . . . . . . . . :  SAMUEL
  Partition release . . . . . . . . . :  V5R1M0 L000

Select one of the following:

     1. Display partition information
     2. Work with partition status

     4. Recover configuration data

 Selection
      4

 F3=Exit   F12=Cancel
 Configuration data errors detected - see Product Activity Log.
```

Figure 7-36 Work with System Partitions display

6. In our example, where we are recovering a remote load source unit that was previously copied from another partition, LPAR configuration information is available on our remote load source unit. This is not valid for the partition we are recovering to, so we need to

discard it. Select option 3 to clear the LPAR configuration data from the non-configured remote load source unit.

```
                           Recover Configuration Data
                                                             System:  S101880D
Select one of the following:



     3. Clear non-configured disk unit configuration data

     5. Accept load source disk unit
     6. Copy configuration data to other side






Selection
     3
```

Figure 7-37 Clearing LPAR configuration data

7. You see with the Select Non-Configured Disk Unit for Configuration Data Clear display (Figure 7-38). This display shows the disk unit that contains the invalid LPAR configuration data.

   Select option 1 (Select) to identify the disk unit containing the invalid LPAR configuration data.

```
              Select Non-Configured Disk Unit for Configuration Data Clear
                                                            System:   S101880D
Type option, press Enter.
  1=Select

     I/O Resource                        ---Last Updated---  System
Opt  Description         Type-Model  Date      Time      Serial Number
 1   Disk Unit           2105-A81    03/23/01  17:57:53  5GG8M








F3=Exit    F10=Display serial/part numbers   F12=Cancel
```

Figure 7-38 Identifying the disk unit containing the invalid LPAR configuration data

8. The confirmation display (Figure 7-39) appears. Verify that this is the correct disk unit. Press F10 to display the serial numbers if you are not certain. When you are sure, press Enter to confirm.

```
          Confirm Non-Configured Disk Unit for Configuration Data Clear
                                                 System:   S101880D
Press Enter to confirm your choice to clear the logical
  partitioning configuration data residing on the disk
  units listed below.
Press F12 to return to change your choice.

I/O Resource                    ---Last Updated---  System
Description         Type-Model  Date      Time      Serial Number
Disk Unit           2105-A81    03/23/01  17:57:53  5GG8M







F10=Display serial/part numbers   F12=Cancel
```

Figure 7-39 Clearing invalid LPAR configuration details

9. You return to the Recover Configuration Data display. It shows the message:

   `Clear non-configured disk unit configuration data successful.`

   Press F12 until you reach the Use Dedicated Service Tools (DST).

You have now completed the tasks specific to an LPAR environment.

7.4.4 Continuing to the recover the load source

The following steps complete the tasks required to recover the remote load source unit:

1. On the Use Dedicated Service Tools (DST) menu (Figure 7-40), you may want to check the current disk configuration. To do so, select option `4` (Work with disk units).

```
                    Use Dedicated Service Tools (DST)
                                                       System:  S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Work with Licensed Internal Code
     4. Work with disk units
     5. Work with DST environment
     6. Select DST console mode
     7. Start a service tool
     8. Perform automatic installation of the operating system
     9. Work with save storage and restore storage
    10. Work with remote service support
    11. Work with system partitions



Selection
     4
```

Figure 7-40 Using DST

2. On the Work with Disk Units display (Figure 7-41), select option 1 (Work with disk configuration).

```
         Work with Disk Units

Select one of the following:

     1. Work with disk configuration
     2. Work with disk unit recovery




Selection
     1

F3=Exit        F12=Cancel
```

Figure 7-41 Work with Disk Units display

3. The Work with Disk Configuration menu (Figure 7-42) provides several options. We want to check the current disk configuration to ensure that we only have one disk unit available.

   Select option 1 (Display disk configuration).

```
         Work with Disk Configuration

Select one of the following:

     1. Display disk configuration
     2. Work with ASP threshold
     3. Work with ASP configuration
     4. Work with mirrored protection
     5. Work with device parity protection
     6. Work with disk compression




Selection
     1

F3=Exit        F12=Cancel
```

Figure 7-42 Work with Disk Configuration

4. From here, we want to check the status of the configured disks. On the Display Disk Configuration display (Figure 7-43), select option 1 (Display disk configuration status).

```
           Display Disk Configuration

Select one of the following:

     1. Display disk configuration status
     2. Display disk configuration capacity
     3. Display disk configuration protection
     4. Display non-configured units
     5. Display device parity status
     6. Display disk hardware status
     7. Display disk compression status




 Selection
     1

 F3=Exit        F12=Cancel
```

Figure 7-43 Display Disk Configuration

5. You should only see one configured disk unit on the status display (Figure 7-44). If you have more than one unit, go back to DST and remove all the other units apart from the LSU.

```
               Display Disk Configuration Status

            Serial                Resource
 ASP  Unit  Number      Type Model  Name        Status
   1                                            Unprotected
        1   00-CA71E    6713  030   DD001       Configured





 Press Enter to continue.

 F3=Exit       F5=Refresh        F9=Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 7-44 Display Disk Configuration Status

6. Press F12 (Cancel) until you reach the Work with Disk Units menu.
7. On the Work with Disk Units menu (Figure 7-45), select option 2 (Work with disk unit recovery).

```
          Work with Disk Units

Select one of the following:

     1. Work with disk configuration
     2. Work with disk unit recovery






Selection
     2

F3=Exit         F12=Cancel
```

Figure 7-45 Work with Disk Units

8. On the Work with Disk Unit Recovery display, scroll to the next page (Figure 7-46) and choose option 16 (Recover mirrored load source).

```
          Work with Disk Unit Recovery

Select one of the following:

    14. Correct device parity protection mismatch
    15. Recover unknown load source
    16. Recover mirrored load source
    17. Recover from start compression failure
    18. Migrate load source disk unit data




                                                                  Bottom

Selection
     16

F3=Exit   F11=Display disk configuration status   F12=Cancel
```

Figure 7-46 Beginning to recover mirrored load source

9. You see a confirmation display (Figure 7-47), which indicates the disk configuration that will be recovered along with the mirrored load source. Scroll through this display to check the disk configuration to be recovered.

```
                 Confirm Recover Mirrored Load Source

 To proceed, the system must copy all data from the remote load
 source to the new local load source.  Once you confirm on
 this screen, the control panel will display the percent of
 completion in the 'xx' places of the C6xx 4205 System
 Reference Code (SRC).  After the successful completion of this
 operation, the system will be IPLed to DST.

 Press Enter to confirm the recovery of the mirrored load source.

 The system will have the following configuration:

             Serial                 Resource
 ASP  Unit   Number      Type Model Name        Status
   1                                            Mirrored
        1    00-CA71E    6713  030  DD001       Active
        1    75-1044000  9337  59A  DD004       Active
        3    75-2044000  9337  59C  DD005       DPY/Active
        4    75-3044000  9337  5AC  DD006       DPY/Active
                                                                 More...

 F12=Cancel
```

Figure 7-47 First Confirm Recover Mirrored Load Source display

10.Verify the correct number of disk units. Press Enter to start the recovery.

```
          Confirm Recover Mirrored Load Source

 To proceed, the system must copy all data from the remote load
 source to the new local load source.  Once you confirm on
 this screen, the control panel will display the percent of
 completion in the 'xx' places of the C6xx 4205 System
 Reference Code(SRC).  After the successful completion of this
 operation, the system will be IPLed to DST.

 Press Enter to confirm the recovery of the mirrored load source.

 The system will have the following configuration:

             Serial                 Resource
 ASP  Unit   Number      Type Model Name        Status
        5    75-4044000  9337  5AC  DD007       DPY/Active
        6    75-5044000  9337  5AC  DD008       DPY/Active
        7    75-6044000  9337  5AC  DD002       DPY/Active
        8    75-7044000  9337  5AC  DD003       DPY/Active

                                                                  Bottom

 F12=Cancel
```

Figure 7-48 Last Confirm Recover Mirrored Load Source display

11.When the LSU is successfully recovered, you see a message on the Disk Configuration Information Report (Figure 7-49). The system comes up to DST from where you can continue to IPL the system and proceed as normal. You may want to check the other non-disk hardware configuration since it is likely to be different from the configuration on the source system.

12. At this point, verify that the system is set to IPL from side A. Recover Remote Load Source usually leaves the IPL as source A.

```
              Disk Configuration Information Report

 The following are informational messages about disk
 configuration changes started in the previous IPL.

 Information

 Recover mirrored load source completed successfully




 Press Enter to continue.
```

Figure 7-49 Remote load source recovered successfully

7.5 Examples of mirroring in conjunction with ESS

In this section we describe how to setup mirroring for the following scenarios:

- Mirroring of iSeries internal disks to ESS
- Mirroring of iSeries LUNs within the ESS
- Mirroring between two ESSs

7.5.1 Mirroring of iSeries internal disks to ESS

In this scenario we describe how to setup mirroring between internal disks and ESS which can be locally or remotely attached to iSeries. This implementation provides availability in case of a failure of disk, FC adapter, IOP, bus, ESS host adapter, connection, etc., it also enables concurrent maintenance of a disk, FC adapter, ESS host bay, etc. Besides, it provides disaster recovery solution: in case the iSeries fails, another iSeries can recover from mirrored half of disks on ESS. Refer to 7.4, "Recovering the load source unit" on page 179 for instructions how to recover from disk space on ESS. We can also use this scenario in combination with ESS Flashcopy for minimizing backup window: We take Flashcopy of ESS LUNs, recover iSeries disk space from Flashcopy by another iSeries, and take backup from this iSeries. For details about implementing Flashcopy for backups refer to Chapter 13. Taking Backups from Flashcopy.

In our scenario there are 4 internal disks: Load source has ca[pacity 70 GB, 3 other disks have capacity 35 GB each. This can be observer in DST by selecting 4. Work with disk units, 1. Work with disk configuration and 1. Display disk configuration status. The displayed disk status is shown in Figure 7-50. We can observe the configuration of internal disks also in SST. Internal disks are on iSeries bus 25.

Display Disk Configuration Status

ASP	Unit	Serial Number	Type	Model	Resource Name	Status
1						Unprotected
	1	21-00F26	4327	C50	DD070	Configured
	2	02-00135	6719	C50	DD031	Configured
	3	02-000A5	6719	C50	DD034	Configured
	4	02-00145	6719	C50	DD032	Configured

Figure 7-50 Mirroring internal disks to ESS - Configuration of internal disks

ESS is connected to iSeries via FC adapter 2787 on bus 26. First we have to look for the WWPN of this adapter. We achieve this in DST by selecting 7. Start a service tool, 4. Hardware service manager, 2. Logical hardware resources, 1. System bus resources. On the screen Logical Hardware Resources on System Bus we type the number of the bus over the value *ALL, as shown in Figure 7-51.

```
                   Logical Hardware Resources on System Bus

System bus(es) to work with . . . . . .   26_   *ALL, *SPD, *PCI, 1-511
Subset by . . . . . . . . . . . . . . .   *ALL  *ALL, *STG, *WS, *CMN,

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=I/O Debug
  7=Display system information
  8=Associated packaging resource(s)     9=Resources associated with IO

                                                                Resource
Opt Description                      Type-Model  Status          Name
 _  HSL I/O Bridge                    282F-      Operational     BC01
 _   Bus Expansion Adapter                -      Operational     BCC02
 _    System Bus                      282F-      Operational     LB01
 _     Multi-adapter Bridge           282F-      Operational     PCI01D
 _      Combined Function IOP  * <    284E-001   Operational     CMB01
 _      Combined Function IOP         2842-001   Operational     CMB03
 _  HSL I/O Bridge                    283B-      Operational     BC02

F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
F9=Failed resources      F10=Non-reporting resources
F11=Display serial/part numbers      F12=Cancel
```

Figure 7-51 Mirroring internal disks to ESS - Selecting a bus

On the screen which shows the selected bus, we type 9 at the Combined Function IOP, as shown in Figure 7-52.

```
                    Logical Hardware Resources on System Bus

System bus(es) to work with . . . . . .   26     *ALL, *SPD, *PC
Subset by . . . . . . . . . . . . . . .   *ALL   *ALL, *STG, *WS

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=I/O Debu
  7=Display system information
  8=Associated packaging resource(s)     9=Resources associated


Opt Description                        Type-Model  Status
 _  System Bus                          283B-      Operational
 _   Multi-adapter Bridge               283B-      Operational
 9    Combined Function IOP             2844-001   Operational





F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting r
F9=Failed resources      F10=Non-reporting resources
F11=Display serial/part numbers      F12=Cancel
```

Figure 7-52 Mirroring internal disks to ESS - IOP resources

On the screen Logical Hardware Resources associated with IOP we type at the relevant storage adapter, as shown in Figure 7-53.

```
                 Logical Hardware Resources Associated with IOP

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=I/O Del
  7=Verify           8=Associated packaging resource(s)


Opt Description                         Type-Model  Status
 _  Combined Function IOP                2844-001   Operational
 _   Communications IOA                  2849-001   Operational
 _    Communications Port                2849-001   Operational
 5   Storage IOA                         2787-001   Operational
```

Figure 7-53 Mirroring internal disks to ESS - Display IOA

On the screen Auxiliary Storage Hardware Resource Detail we observe the WWPN of this adapter, as shown in Figure 7-54.

```
                     Auxiliary Storage Hardware Resource Detail

 Description . . . . . . . . . . . . . :   Storage IOA
 Type-Model  . . . . . . . . . . . . . :   2787-001
 Status  . . . . . . . . . . . . . . . :   Operational
 Serial number . . . . . . . . . . . . :   1E-A308008
 Part number . . . . . . . . . . . . . :   0000000P4339
 Resource name . . . . . . . . . . . . :   DC01
 Port worldwide name . . . . . . . . . :   10000000C9306F82
 PCI bus . . . . . . . . . . . . . . . :
   System bus  . . . . . . . . . . . . :   26
   System board  . . . . . . . . . . . :   0
   System card . . . . . . . . . . . . :   16
 Storage . . . . . . . . . . . . . . . :
   I/O adapter . . . . . . . . . . . . :   4
   I/O bus . . . . . . . . . . . . . . :
   Controller  . . . . . . . . . . . . :
   Device  . . . . . . . . . . . . . . :



 F3=Exit      F5=Refresh       F6=Print
 F9=Change detail              F12=Cancel
```

Figure 7-54   Mirroring internal disks to ESS - WWPN of FC adapter

On ESS we define a host system name for this FC adapter, as shown in Figure 7-55. We can look for the WWPN of an iSeries FC adapter also in SST.

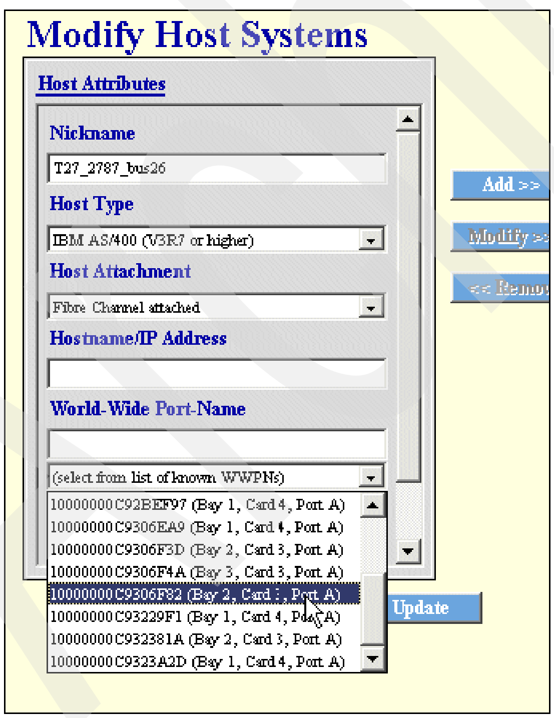

Figure 7-55   Mirroring internal disks to ESS - Define Host System name

The internal disks will be mirrored to ESS LUNs, therefore we have to define the same number and capacity of LUNs, as there are internal disks. the LUNs have to be defined as unprotected because if iSeries doesn't "see" them as unprotected it cannot mirror them. Defined LUNs are shown in Figure 7-56. Refer to Chapter 6, "Creating iSeries storage in

ESS" on page 127 i for detailed instructions about how to create a Host name and define LUNs on ESS.

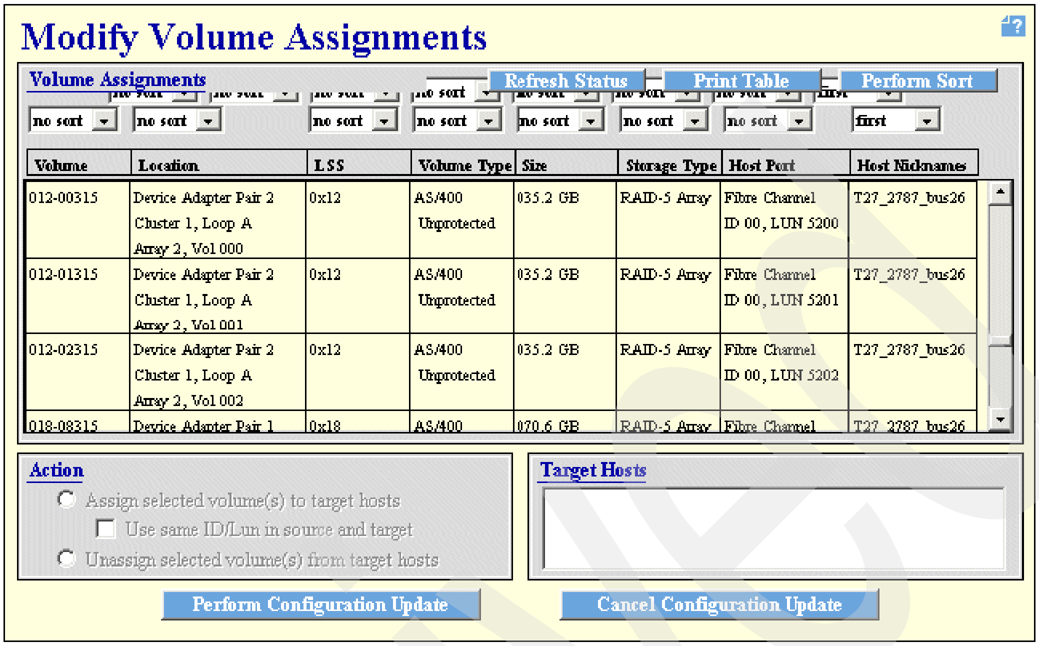

Figure 7-56 Mirroring internal disks to ESS - Defined LUNs

Mirroring is done on ASP level: in an iSeries we can have multiple ASPs, some of them mirrored, some of them RAID protected and some of them unprotected. In our example we want system ASP - ASP 1 to be mirrored. So, we have to add defined LUNs to ASP1 which already contains internal disks. To accomplish this we use DST function 4. Work with disk units, 1. Work with disk configuration, 3. Add units to ASPs; on the screen Specify ASPs to Add Units to we select the LUNs which we want to add to ASP by specifying the ASP number. Adding the 4 LUNs from our example is shown in Figure 7-57 and Figure 7-58.

We can also use SST to add LUNs to ASP.

```
                         Specify ASPs to Add Units to

 Specify the ASP to add each unit to.

 Specify  Serial                                       Resource
   ASP    Number        Type   Model   Capacity   Name
          02-00051      6719    050       35165   DD030
   ___    68-0C04EF9    6719    050       35165   DD082
   ___    02-000CF      6719    050       35165   DD028
   ___    21-00E03      4327    050       70564   DD057
   ___    68-0C0516A    6719    050       35165   DD081
   ___    02-001D1      6719    050       35165   DD029
   ___    68-0C05184    6719    050       35165   DD080
   ___    02-000B2      6719    050       35165   DD027
   ___    68-0C05052    6719    050       35165   DD079
   ___    68-0C0516D    6719    050       35165   DD083
   ___    68-0C04B43    6719    050       35165   DD078
   ___    02-000CC      6719    050       35165   DD033
   1__    75-1808315    2105    A84       70564   DD048
```

Figure 7-57 Mirroring internal disks to ESS - Selecting 70 GB LUN to add to ASP

```
                        Specify ASPs to Add Units to

Specify the ASP to add each unit to.

Specify  Serial                                  Resource
  ASP    Number       Type  Model  Capacity  Name
  ___    75-1809315   2105   A81       8589  DD049
  ___    75-180A315   2105   A81       8589  DD050
  ___    75-180B315   2105   A81       8589  DD051
  ___    75-1805315   2105   A81       8589  DD041
  ___    75-1806315   2105   A81       8589  DD043
  ___    75-1807315   2105   A81       8589  DD044
  1__    75-1200315   2105   A85      35165  DD052
  1__    75-1201315   2105   A85      35165  DD053
  1__    75-1202315   2105   A85      35165  DD054
```

Figure 7-58 Mirroring internal disks to ESS - Selecting 35 GB LUNs to add to ASP

When the LUNs are added to ASP they are initialized and formatted by iSeries, so it takes some time until they are added. During this time, the percentage of how much work is completed is shown on the screen (Figure 7-59).

```
                              Function Status

You selected to add units.




                                   5 % Complete
```

Figure 7-59 Mirroring internal disks to ESS - Adding LUNs to ASP

When mirroring is started iSeries automatically determines which will be the halves of mirrored disks. This decision is based on achieving maximal possible availability: it selects the two halves so that they are on separate HSLs, or Remote buses, or buses, or IOPs, or disks attached to the same IOP. The level of mirroring (bus level, IOP level, etc.) is determined by this selection. However, we can influence that mirroring will be done between disks on a set of buses and disks on another set of buses, so that we rename one set of buses to Remote buses. In our example we want internal disks on bus 25 to be mirrored to ESS LUNs on bus 26, so we rename bus 26 to be Remote bus.

In 7.3.1, “Setting up ESS mirroring” on page 173 we describe how to rename the bus to Remote. In this example we use a slightly different way, nevertheless both ways achieve the same result.

In order to rename bus 26 to Remote we select DST 7. Start a service tool, 4. Hardware service manager, 2. Logical hardware resources, 1. System bus resources. On the screen, Logical Hardware Resources on System Bus, we type the number of the bus over the value *ALL, as shown in Figure 7-51. On the screen, Logical Hardware Resources on System Bus, we type 2 by system bus, as shown in Figure 7-60.

```
                   Logical Hardware Resources on System

System bus(es) to work with . . . . . .   26    *ALL, *
Subset by . . . . . . . . . . . . . . .   *ALL  *ALL, *

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=
  7=Display system information
  8=Associated packaging resource(s)     9=Resources as


Opt Description                        Type-Model  Status
 2  System Bus                          283B-      Operatic
 _   Multi-adapter Bridge               283B-      Operatic
 _    Combined Function IOP             2844-001   Operatic






F3=Exit    F5=Refresh    F6=Print    F8=Include non-rep
F9=Failed resources      F10=Non-reporting resources
```

Figure 7-60 Mirroring internal disks to ESS - Change bus detail

On the screen Change Logical Hardware Resource Detail, we type a new name of the bus over previous one; the new one has to start with R. This is shown in Figure 7-61. Renaming a bus can be also done in SST.

```
                   Change Logical Hardware Resource Deta

 Description . . . . . . . . . . . . :   System Bus
 Type-Model  . . . . . . . . . . . . :   283B
 Status  . . . . . . . . . . . . . . :   Operational
 Serial number . . . . . . . . . . . :   38-0216059

 Current resource name . . . . . . . :   LB04


 Type changes, press Enter.

 New resource name . . . . . . . . .     rB04
```

Figure 7-61 Mirroring internal disks to ESS - Rename a bus to Remote

Now we are prepared to start mirroring. In order to start it we have to use DST, it cannot be accomplished in SST.

In DST we select 4. Work with disk units, 1. Work with disk configuration, 4. Work with mirrored protection. On the screen Work with Mirrored Protection we first select 4. Enable remote load source mirroring, and on the next screen confirm to enable it. We get a message that Load source is successfully enabled, as shown in Figure 7-62.

```
Selection
     2

F3=Exit         F12=Cancel
 Remote load source mirroring enabled successfully.
```

Figure 7-62 Mirroring internal disks to ESS - Enable remote load source mirroring

We press F12 to get back to the screen Work With Mirrored Protection and select 2. Start mirrored protection. On the screen Select ASP to Start Mirrored Protection we select the ASP for which we want to start mirroring. We are presented with the screen Confirm Start Mirrored Protection where we can check which disks will be in each mirrored half. In our example we can see that internal disks will be mirrored to ESS LUNs on Remote bus level. This is shown in Figure 7-63. For detailed instruction how to enable remote load source mirroring and start mirroring refer to 7.3.1, “Setting up ESS mirroring” on page 173 and 7.2, “Remote load source mirroring” on page 168.

```
                        Confirm Start Mirrored Protection

 Press Enter to confirm your choice to start mirrored
 protection.  During this process the system will be IPLed.
 You will return to the DST main menu after the IPL is
 complete.  The system will have the displayed protection.

 Press F12 to return to change your choice.

              Serial                   Resource
 ASP  Unit    Number      Type Model   Name         Protection
   1                                                Mirrored
         1    21-00F26    4327  050    DD070        Remote Bus
         1    75-1808315  2105  A84    DD048        Remote Bus
         2    02-00135    6719  050    DD031        Remote Bus
         2    75-1201315  2105  A85    DD053        Remote Bus
         3    02-000A6    6719  050    DD034        Remote Bus
         3    75-1200315  2105  A85    DD052        Remote Bus
         4    02-00145    6719  050    DD032        Remote Bus
         4    75-1202315  2105  A85    DD054        Remote Bus
```

Figure 7-63 Mirroring internal disks to ESS - Check the disks

After we press enter on the screen Confirm Start Mirrored protection the mirroring starts. After some time we get the message start mirroring completed successfully.

```
                   Disk Configuration Information Report

 The following are informational messages about disk
 configuration changes started in the previous IPL.

 Information

 Start mirroring completed successfully
```

Figure 7-64 Mirroring internal disks to ESS - Start mirroring completed

However, at this point the mirrored volumes are not yet synchronized. In DST we can observe that they have status “Resuming”, as shown in Figure 7-65.

```
                      Display Disk Configuration Status

              Serial                       Resource
ASP  Unit     Number       Type Model     Name          Status
  1                                                     Mirrored
        1     21-00F26     4327  050      DD070         Active
        1     75-1808315   2105  A84      DD048         Resuming
        2     02-00135     6719  050      DD031         Active
        2     75-1201315   2105  A85      DD053         Resuming
        3     02-000A6     6719  050      DD034         Active
        3     75-1200315   2105  A85      DD052         Resuming
        4     02-00145     6719  050      DD032         Active
        4     75-1202315   2105  A85      DD054         Resuming
```

Figure 7-65 Mirroring internal disks to ESS - Volumes after start mirroring

Now we IPL iSeries from DST. During IPL the mirrored volumes will synchronize as shown in Figure 7-66.

```
                 Licensed Internal Code IPL in Progress
                                                                  11/
IPL:
  Type . . . . . . . . . . . . . . . :   Attended
  Start date and time . . . . . . . :   11/11/03  16:15:12
  Previous system end . . . . . . . :   Normal

IPL step  . . . . . . . . . . . . . :   Storage Management Reco

  Elapsed time in minutes . . . . . :       41.5
  Reference code detail . . . . . . :   C6004205  3

Activity  . . . . . . . . . . . . . :   Synchronization of Mirr

  Current . . . . . . . . . . . . . :           95
  Total . . . . . . . . . . . . . . :          100
  Identifier  . . . . . . . . . . . :
```

Figure 7-66 Mirroring internal disks to ESS - Synchronizing of mirrored volumes

After IPL is performed we can observe in SST that the mirrored volumes are synchronized and have status Active. In SST we select 4. Work with disk units, 1. Work with disk configuration and 1. Display disk configuration status. This is shown in Figure 7-67.

```
                      Display Disk Configuration Status

              Serial                       Resource
ASP  Unit     Number       Type Model     Name          Status
  1                                                     Mirrored
        1     21-00F26     4327  050      DD070         Active
        1     75-1808315   2105  A84      DD048         Active
        2     02-00135     6719  050      DD031         Active
        2     75-1201315   2105  A85      DD053         Active
        3     02-000A6     6719  050      DD034         Active
        3     75-1200315   2105  A85      DD052         Active
        4     02-00145     6719  050      DD032         Active
        4     75-1202315   2105  A85      DD054         Active
```

Figure 7-67 Mirroring internal disks to ESS - Synchronized mirrored volumes

7.5.2 Mirroring between LUNs in the same ESS

In this example we use mirroring between iSeries LUNs within the same ESS. Load source is internal disk and is mirrored to ESS. The LUNs from mirrored half are assigned to iSeries adapters on separate buses. This implementation provides continuos work in case of failure of an iSeries FC adapter, ESS host adapter, link between iSeries and ESS, etc. It also enables concurrent maintenance of iSeries FC adapters, ESS host bays, etc. As with mirroring internal disks to ESS, this scenario also provides disaster recovery solution: In case iSeries fails another iSeries can recover from mirrored half of volumes.

In our example we used 4 LUNs for each mirrored half. One mirrored half of LUNs is attached to iSeries FC adapter 2787 on bus 26, the other half is attached to adapter 2766 on bus 25. This scenario is shown in Figure 7-68. The capacity of Load source is 70 GB, while the capacity of the 3 mirrored LUNs is 8 GB.

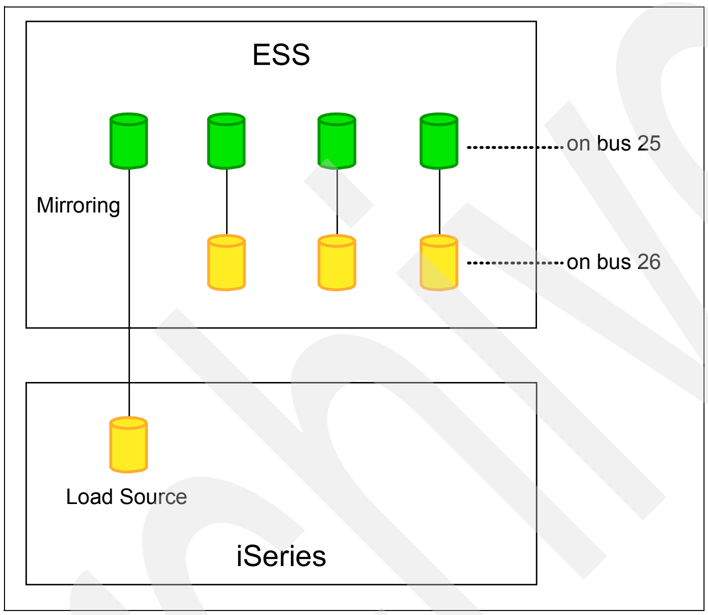

Figure 7-68 Mirroring within ESS - scenario

We first look for WWPNs of the two iSeries adapters. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to do this.

In ESS we define a host name for each adapter. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to do this.

In our example we define in ESS 3 LUNs of capacity 8 GB for Host name for adapter 2787. We define 3 LUNs of capacity 8 GB and 1 LUN of capacity 70 GB for Host name for adapter 2766. For instructions how to define LUNs refer to Chapter 6. Creating iSeries storage on ESS.

Figure 7-69 and Figure 7-70 show defined LUNs as can be observed in ESS StorWatch Specialist screen Modify volume assignments.

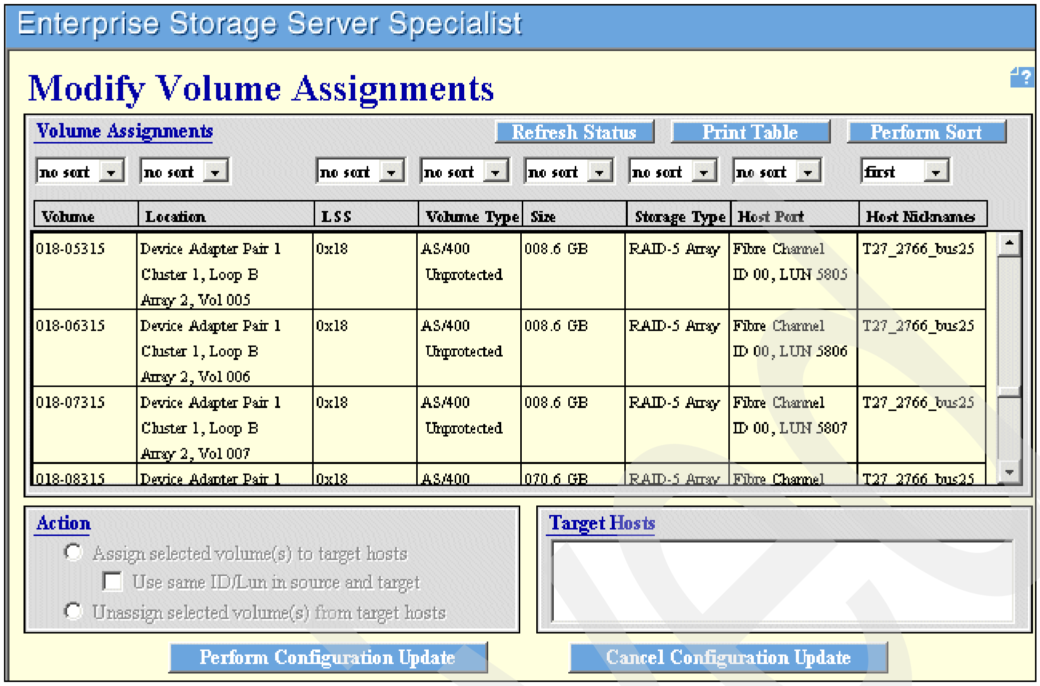

Figure 7-69 Mirroring within ESS - LUNs for one mirrored half

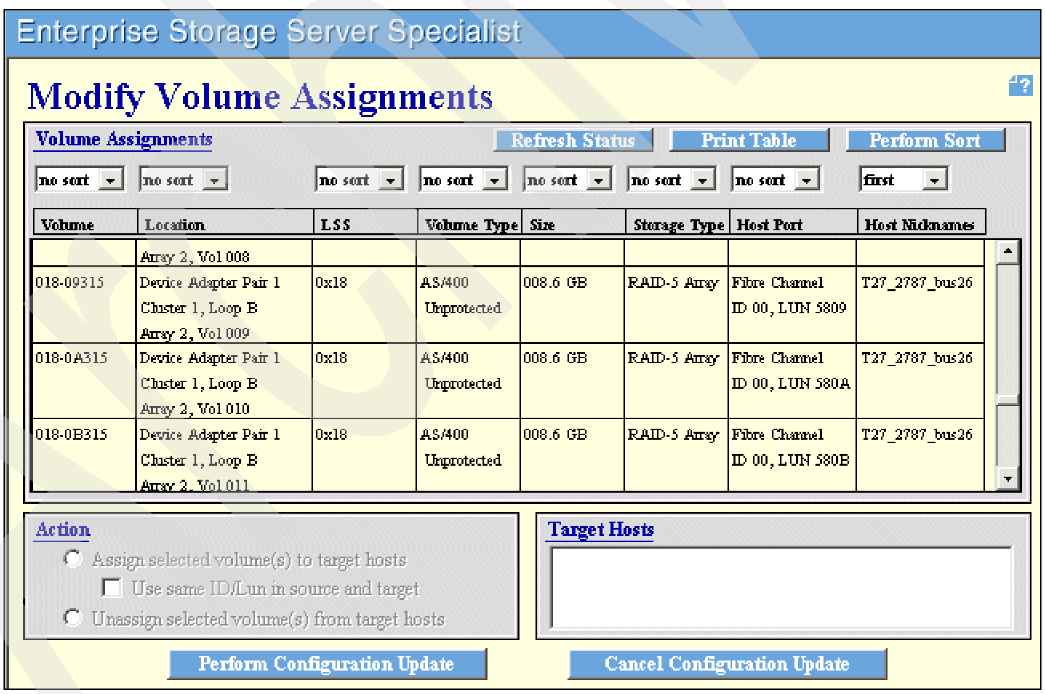

Figure 7-70 Mirroring within ESS - LUNs for the other mirrored half

In iSeries DST we add the defined LUNs to ASP as shown in Figure 7-71. For instructions how to add LUNs to ASP refer to Chapter 7. Adding ESS storage to the iSeries server.

```
                    Specify ASPs to Add Units to

  Specify the ASP to add each unit to.

  Specify  Serial                                  Resource
    ASP    Number        Type  Model  Capacity     Name
           02-00145      6719   050       35165    DD032
           02-00135      6719   050       35165    DD031
           75-1202772    2105   A81        8589    DD045
           75-1204772    2105   A81        8589    DD047
     1     75-1808315    2105   A84       70564    DD048
           75-1203772    2105   A81        8589    DD046
     1     75-1807315    2105   A81        8589    DD044
     1     75-1806315    2105   A81        8589    DD043
     1     75-1805315    2105   A81        8589    DD041
     1     75-1809315    2105   A81        8589    DD049
     1     75-180A315    2105   A81        8589    DD050
     1     75-180B315    2105   A81        8589    DD051
```

Figure 7-71   Mirroring within ESS - Add LUNs to ASP

After the volumes are added to ASP, we make sure that the volumes on bus 25 will be in one mirrored half, and the volumes on bus 26 and Load source will be in the other mirrored half. For this we rename the bus 25 to Remote bus, as shown in Figure 7-72. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to rename a bus to Remote bus.

```
                   Change Logical Hardware Resource Det

Description . . . . . . . . . . . . :   System Bus
Type-Model  . . . . . . . . . . . . :   283B
Status  . . . . . . . . . . . . . . :   Operational
Serial number . . . . . . . . . . . :   38-0216059

Current resource name . . . . . . . :   LB03


Type changes, press Enter.

New resource name . . . . . . . . .     rB03
```

Figure 7-72   Mirroring within ESS - Rename a bus to Remote bus

All actions which were performed on iSeries up to now could be done also in SST while iSeries is running. But now we enable Load source mirroring and start mirroring. This can only be performed in DST. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to enable Load source and start mirroring. In this example volumes will be mirrored as shown in Figure 7-73.

```
_                    Confirm Start Mirrored Protection

 Press Enter to confirm your choice to start mirrored
 protection.  During this process the system will be IPLed.
 You will return to the DST main menu after the IPL is
 complete.  The system will have the displayed protection.

 Press F12 to return to change your choice.

             Serial                  Resource
 ASP  Unit   Number      Type Model  Name         Protection
   1                                              Mirrored
         1   21-00F26    4327  050   DD070        Remote Bus
         1   75-1808315  2105  A84   DD048        Remote Bus
         3   75-1807315  2105  A81   DD044        Remote Bus
         3   75-180B315  2105  A81   DD051        Remote Bus
         4   75-1806315  2105  A81   DD043        Remote Bus
         4   75-180A315  2105  A81   DD050        Remote Bus
         5   75-1805315  2105  A81   DD041        Remote Bus
         5   75-1809315  2105  A81   DD049        Remote Bus
```

Figure 7-73 Mirroring within ESS - Check the volumes

After we check that volumes will be mirrored as desired we press enter on the screen Confirm Start Mirrored Protection and mirroring is starting. After some time we get the message Start Mirroring completed successfully. At this point the mirrored volumes are not yet synchronized, they will be synchronized during IPL.

7.5.3 Mirroring between two ESSs

In this example we use mirroring between iSeries LUNs on two ESSs, one locally attached to iSeries and the other attached on long distance. The volumes are mirrored between the two ESSs, internal Load source is mirrored to remote ESS. This scenario provides disaster recovery solution in case the site of iSeries and local ESS is damaged. In this case we use another iSeries on remote site to recover from mirrored volumes on remote ESS. Besides, this scenario provides the same availability and concurrent maintenance as are described in previous two examples.

In our example iSeries connects to local ESS with FC adapter 2787 on bus 26, and to remote ESS with FC adapter 2766 on bus 25. Load source has capacity 70 GB, all the LUNs have capacity 8 GB except the LUN to which Load source is mirrored: this one has the same capacity as Load source.

In DST we look for WWPNs of both FC adapters. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to do this.

In ESS we define a host name for each adapter. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to do this.

In our example we define on local ESS 4 LUNs of capacity 8 GB. On remote ESS we define one LUN of capacity 70 GB and 3 LUNs of capacity 8 GB. For instructions how to define LUNs refer to Chapter 6, “Creating iSeries storage in ESS” on page 127.

Figure 7-74 shows defined LUNs on local ESS as can be observed in ESS StorWatch Specialist screen Modify volume assignments.

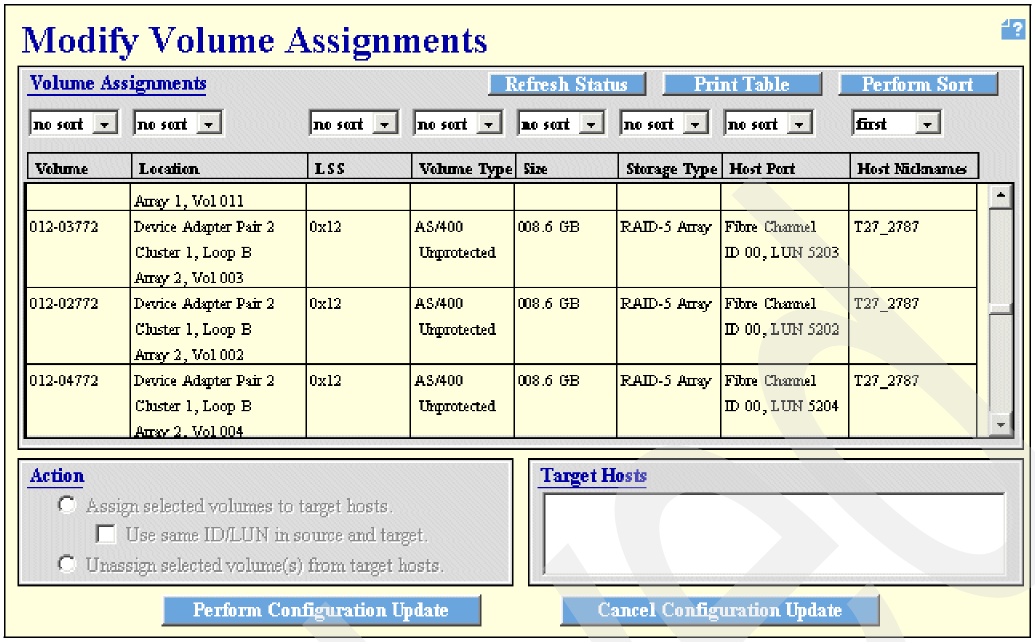

Figure 7-74 Mirroring between 2 ESSs - LUNs on local ESS

Figure 7-75 shows defined LUNs on local ESS as can be observed in ESS StorWatch Specialist screen Modify volume assignments.

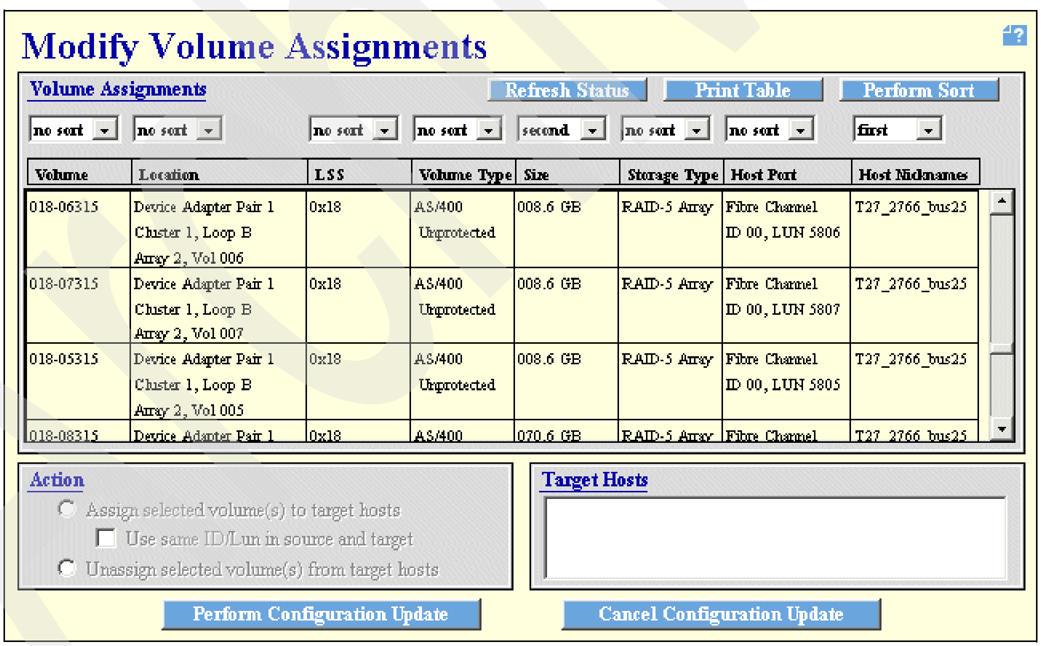

Figure 7-75 Mirroring between 2 ESSs - LUNs on remote ESS

After we define LUNs on both ESSs we add them to iSeries ASP. This can be done in DST or SST.

Adding LUNs from local ESS to ASP is shown in Figure 7-76.

```
                           Specify ASPs to Add Units to

 Specify the ASP to add each unit to.

 Specify  Serial                                  Resource
   ASP    Number        Type   Model   Capacity   Name
          02-000A6      6719    050       35165   DD034
          02-00135      6719    050       35165   DD031
   1      75-1202772    2105    A81        8589   DD045
    1     75-1203772    2105    A81        8589   DD046
    1     75-1204772    2105    A81        8589   DD047
```

Figure 7-76   Mirroring between 2 ESSs - Adding LUNs from local ESS

Figure 7-77 shows adding LUNs from remote ESS to ASP.

```
                           Specify ASPs to Add Units to

 Specify the ASP to add each unit to.

 Specify  Serial                                   Resource
   ASP    Number        Type   Model    Capacity   Name
          75-1210772    2105    A04        70564   DD002
          75-160F772    2105    A01         8589   DD004
          75-1213772    2105    A01         8589   DD014
          75-1420772    2105    A01         8589   DD036
          75-141F772    2105    A01         8589   DD007
          75-1218772    2105    A01         8589   DD005
          68-0C04EF9    6719    050        35165   DD082
          68-0C0516A    6719    050        35165   DD081
   1      75-1805315    2105    A81         8589   DD041
   1      75-1808315    2105    A84        70564   DD048
   1      75-1806315    2105    A81         8589   DD043
   1      75-1807315    2105    A81         8589   DD044
```

Figure 7-77   Mirroring between 2 ESSs - Adding LUNs from remote ESS

After the volumes are added to ASP, we make sure that the LUNs on local ESS will be in one mirrored half and the LUNs on remote ESS in the other mirrored half. To ensure this we rename the bus 25 to Remote bus, as shown in Figure 7-78. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions how to rename a bus to Remote bus.

```
                        Change Logical Hardware Resource Detail

 Description . . . . . . . . . . . . . :   System Bus
 Type-Model  . . . . . . . . . . . . . :   283B
 Status  . . . . . . . . . . . . . . . :   Operational
 Serial number . . . . . . . . . . . . :   38-0216059

 Current resource name . . . . . . . . :   LB03


 Type changes, press Enter.

 New resource name . . . . . . . . . .     rB03
```

Figure 7-78   Mirroring between 2 ESSs - Rename a bus to Remote bus

In DST we enable Remote load source mirroring and start mirroring. Refer to 7.5.1, “Mirroring of iSeries internal disks to ESS” on page 195 for instructions about how to enable Load source and start mirroring. The screen Confirm start Mirrored shows how the volumes will be mirrored between the two ESSs, as shown in Figure 7-79. Observe the models of ESS LUNs: models 8Ax show that the volumes are unprotected, the last digit of the model denotes the capacity.

```
                       Confirm Start Mirrored Protection

 Press Enter to confirm your choice to start mirrored
 protection.  During this process the system will be IPLed.
 You will return to the DST main menu after the IPL is
 complete.  The system will have the displayed protection.

 Press F12 to return to change your choice.

             Serial                     Resource
 ASP  Unit   Number      Type Model     Name         Protection
   1                                                 Mirrored
         1   21-00F26    4327  050      DD070        Remote Bus
         1   75-1808315  2105  A84      DD048        Remote Bus
         2   75-1805315  2105  A81      DD041        Remote Bus
         2   75-1202772  2105  A81      DD045        Remote Bus
         4   75-1806315  2105  A81      DD043        Remote Bus
         4   75-1203772  2105  A81      DD046        Remote Bus
         5   75-1807315  2105  A81      DD044        Remote Bus
         5   75-1204772  2105  A81      DD047        Remote Bus
```

Figure 7-79 Mirroring between 2 ESSs - Check and confirm mirroring

After we check that volumes will be mirrored as desired, we press Enter on the screen, Confirm Start Mirrored Protection, and mirroring is starting. After some time we get the message Start Mirroring completed successfully. At this point the mirrored volumes are not yet synchronized — they will be synchronized during IPL. After the mirrored volumes are synchronized during IPL they have status Active. This can be observed in SST by selecting 4. Work with disk units, 1. Work with disk configuration and 1. Display disk configuration status. The status of mirrored volumes is shown in Figure 7-80.

```
                         Display Disk Configuration Status

             Serial                     Resource
 ASP  Unit   Number      Type Model     Name         Status
   1                                                 Mirrored
         1   21-00F26    4327  050      DD070        Active
         1   75-1808315  2105  A84      DD048        Active
         2   75-1805315  2105  A81      DD041        Active
         2   75-1202772  2105  A81      DD045        Active
         4   75-1806315  2105  A81      DD043        Active
         4   75-1203772  2105  A81      DD046        Active
         5   75-1807315  2105  A81      DD044        Active
         5   75-1204772  2105  A81      DD047        Active
```

Figure 7-80 Mirroring between 2 ESSs - Mirrored volumes

8

Multipath for iSeries

In this chapter we describe the concept of Multipath for iSeries, which is available with OS/400 level V5R3. We also give guidelines for how to plan and implement Multipath for iSeries with new or existing installations.

8.1 Description of Multipath

Spreading disks over multiple I/O adapters can cause data availability issues for any system. For example, if one host adapter fails then access to the disks connected to this adapter would be lost. The concept of multipathing prevents such situations from occurring where the data is located on an Enterprise Storage Server (ESS). This support takes advantage of multiple paths between a host system and an ESS LUN or set of LUNs, as shown in Figure 8-5 on page 216. In a case where an adapter fails, the system automatically reroutes I/O operations to another available path. This support also provides the possibility of balancing I/O load among multiple paths, and in so doing prevents I/O bottlenecks.

Figure 8-1 shows why we need multipath support. There are a number of single points of failure between the iSeries bus and the first switch port. Beyond the switch port there can be redundancy in both the fabric and the ESS. Without multipath the internal bus cannot reroute I/O from a failed adapter.

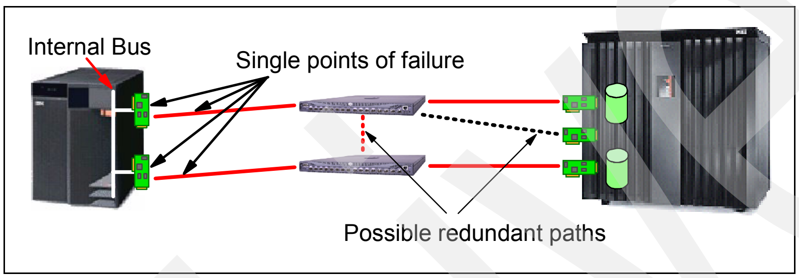

Figure 8-1 Possible failure points without Multipath

With multipath as in Figure 8-2 the internal bus has up 8 possible connection to transport IO. Typically we believe the maximum connections used with be 3.

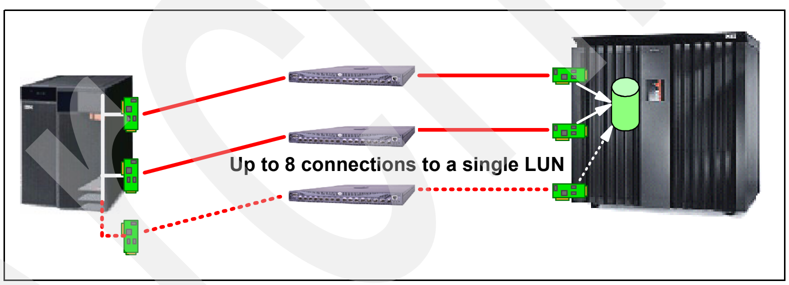

Figure 8-2 Multipath possibilities

Multipath concept is used by IBM software Subsystem Device Driver (SDD) which was developed for several host platforms like Windows, Unix and Linux.

iSeries offers two disk protection possibilities, Mirroring and RAID-5. OS/400 mirroring maintains the duplicate copy of data on separate disks which are attached to separate adapters. It can be used among internal disks, between internal disks and ESS LUNs, between the LUNs off the same ESS, or between two ESSs. Among others, OS/400 mirroring enables continuous access to data in case an adapter fails, and has been used up to now for the same purpose as SDD on other platforms. However, OS/400 mirroring provides also other availability and Disaster recovery possibilities, not only the continuous work in case of adapter failure.

OS/400 V5R3 has a new built-in capability for multipathing. If there is more than one path to the same set of ESS LUNs it will automatically establish multipathing for them. There can be up to 8 paths for the same set of LUN.

In ESS we achieve multiple paths to the same set of LUNs by assigning them to multiple iSeries FC adapters. After we do this, iSeries recognizes the multiple paths to the same set of LUNs and starts to use them: It spreads the I/O operations to these LUNs across all available paths using “Round Robin” rule to determine through which path will an I/O operation go. If one path will become unavailable because of failure of iSeries FC adapter, iSeries bus, ESS Host adapter, etc., the I/O load for the same LUNs will be automatically transferred to other available paths.

Multipath can be observed on iSeries in the following ways:

- In iSeries Navigator for V5R3 select Hardware, **disks**
- In SST you will notice that device name for a LUN with multiple paths starts with DPMxx, while device name for a LUN with only one path starts with DDxx.

8.2 Planning

This topic provides guidelines to help you plan how to configure your iSeries for multipath connectivity to ESS. The guidelines are intended to work for most customer, but are certainly not applicable to all situations. You must work with an IBM @server iSeries Representative and SAN Solution Representative.

Software requirements

Multipath is supported with iSeries from OS version V5R3 and does not need any additional software to enable multipath support.

ESS requires code of 2.2.0.X or above to support multipath with iSeries.

Hardware requirements

The hardware requirements of iSeries and ESS do not change much with the addition of Multipath support. Additional hardware is required under the following conditions

- One adapter on the iSeries
- One FC adapter on the ESS
- One switch connected between the ESS and iSeries

The adapters used for connecting the iSeries to ESS are 2766 and 2787. The planning, placement rules and performance for 2766 and 2787 are explained in chapter.

The Feature Code of the adapter used in ESS for connecting to iSeries is FC#3025 Short Wave® Fibre Channel adapter.

The best configuration would be to connect multiple iSeries adapters to multiple ESS adapter using redundant switches. The various scenarios of connecting ESS to iSeries are explained in chapter.

The number of FC adapters required on the iSeries is dependant on the existing LUNs per FC card and the 32 LUN per adapter limit. For example in Figure 8-3, there are four 2766 FC adapters on the iSeries and there are only 10 LUNs per adapter. This low number of LUNs per adapter could be for performance or just and small implementation.

In this case you could:

- Simply start multipath on each pair of 2766s as the total LUNs is 20 (of 32 possible) with 2 connections. This would should not have an performance impact and may show a performance improvement. Under adapter failure the performance would be degraded as all I/O would flow through 1 adapter.
- Add two additional 2766s, making 3 connections. Effectively 20/3 LUNs per adapter, this could show improved performance. Under adapter failure, the performance would be equal to the original configuration 20/2. So with multipath support this could still show a performance improvement.
- You should add an additional adapter to each original LUN. This would be the same as 5 LUNs per adapter. Again there could be a performance improvement. This is really an unwarranted configuration,

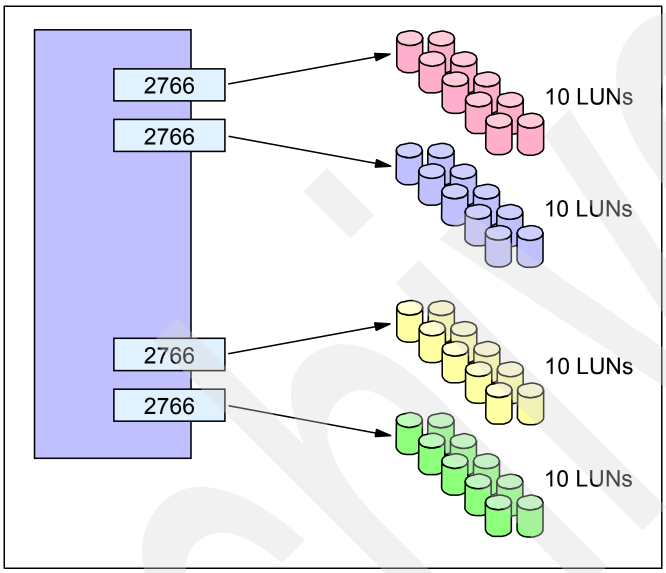

Figure 8-3 Multipath Example 1

In our next example there are 20 LUNs per adapter. This is a very typical configuration for the LUN to adapter ratio. We will look at the options we used in the previous example:

Simply start multipath between each pair of 2766s. This would result in an invalid configuration as the total possible LUNs under an adapter failure condition would be 40, which is greater than the maximum of 32.

Add one 2766 per adapter pair resulting in 40/3 LUNs per adapter. Under single adapter failure conditions there would be again be 40/2 adapter ratio which is invalid.

if you add an adapter per existing adapter under single adapter failure conditions you would have 40/3 LUNs per which is valid and would still offer better performance than the existing configuration.

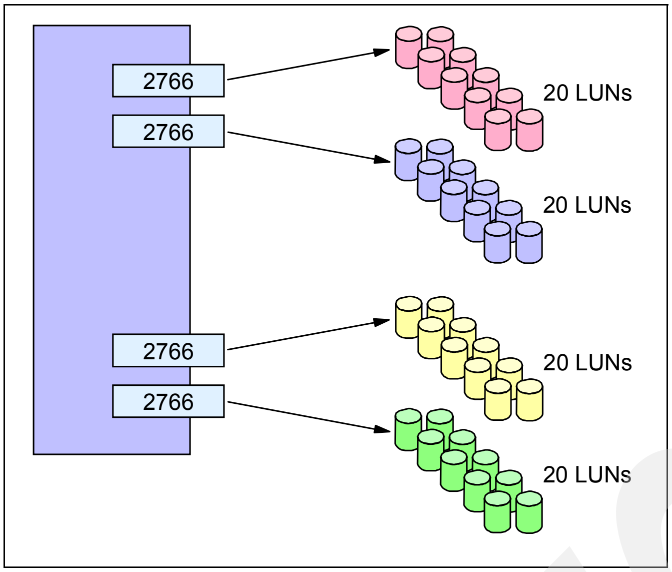

Figure 8-4 Multipath planning 2

Note: Up to 8 paths are supported for the same set of LUNs.

8.3 Implementation and examples

In this section we describe how to implement multipath in a new environment of iSeries and ESS and in an existing environment of iSeries connected to ESS. We also describe how to implement in simple and more complex environments.

8.3.1 Implementing multipath in a new environment

When you are planning multipath for a new ESS with iSeries, you should consider the following:

- Decide which ASPs and IASPs will be protected with multipath.
- Plan the number of paths for each ASP or IASP.
- You may decide to use some FC adapters in several multipath protections. This can be observed in Figure 8-5: ASP1 is protected with two paths (FC adapter 1 and 2), ASP2 is protected by three paths (FC adapter 2, 3 and 4). FC adapter 2 is used in both protections. This decision can be based on capability of FC adapter, for instance, you decide to share more capable FC adapter 2787.

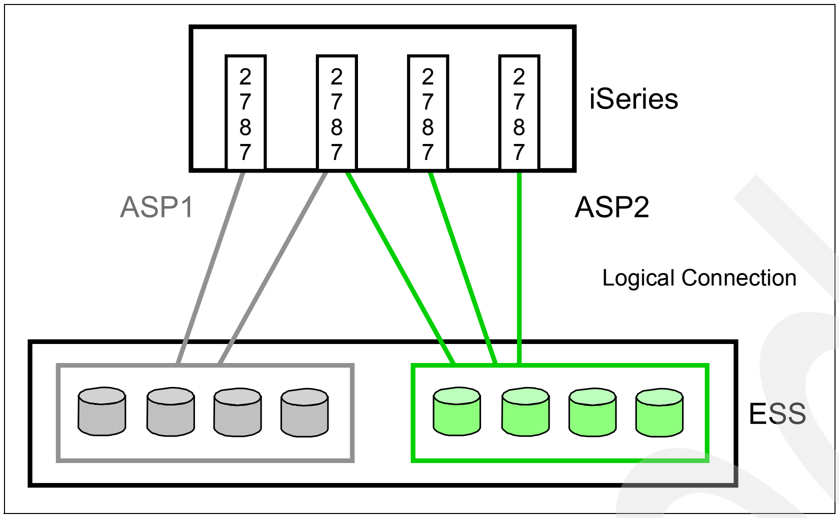

Figure 8-5 Considering FC adapters for multipath

- Define LUNs which you want to protect and assign them to all FC adapters in a multipath environment.

In the following example we define 2 LUNs and assign them to 3 FC adapters. We add these 2 LUNs to an ASP which will be protected by multipath to the 3 FC adapters.

In ESS we define host names for the 2 iSeries FC adapters. We define 2 LUNs for the host name of one FC adapter. Refer to Chapter 6. Creating iSeries storage for ESS for instructions about how to define LUNs.

We assign the 2 LUNs to the other host name of FC adapter. This is shown in Figure 8-6.

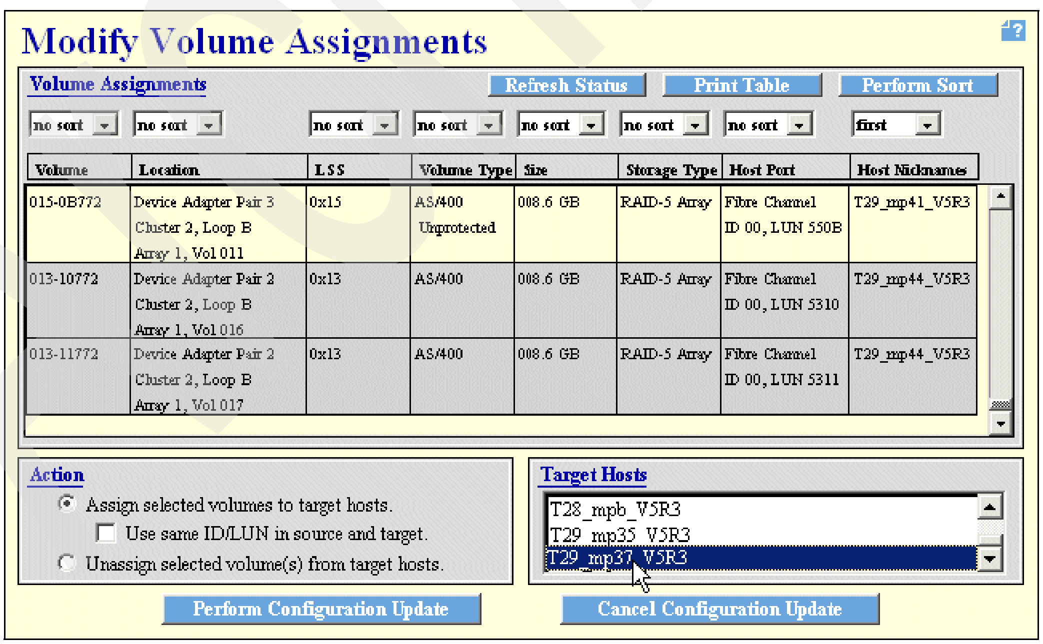

Figure 8-6 Multipath of one ASP - Assign LUNs to multiple FC adapters

Multipath for the 2 LUNs is now established and is active. We observe this in iSeries navigator the following way: On the left panel we expand Configuration and services, we expand

Hardware, we expand Disk Units, enter SST user ID and password, and click the Non-configured disk units. On right-hand panel we observe that deice name of two LUNs now starts with DMPxx as shown in Figure 8-7.

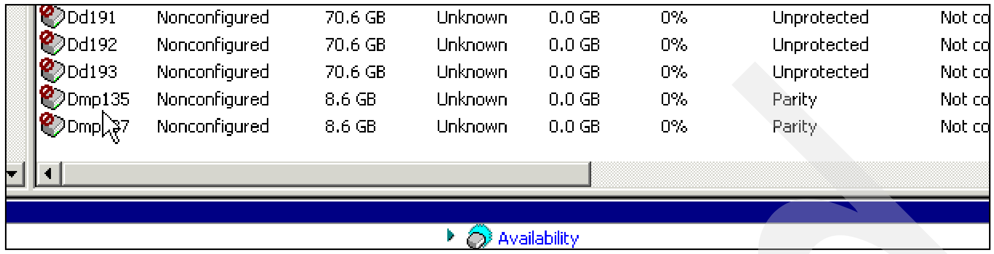

Figure 8-7 Multipath of one ASP - device names

We right-click one LUN and select Retrieve disk log from pull-down menu. In the Properties window for this LUN, we click the Connections tab and observe the number of I/O adapters for this LUN, which are 2. This is shown in Figure 8-8.

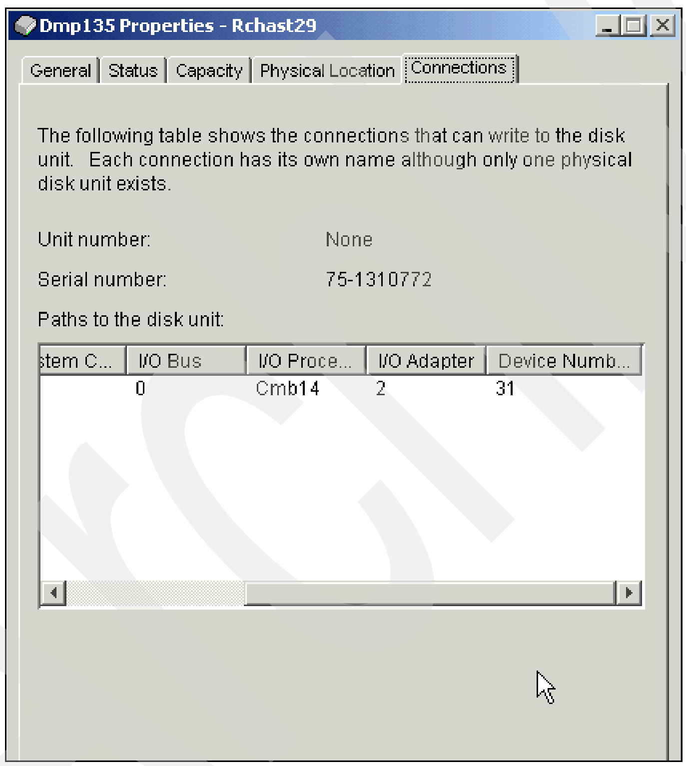

Figure 8-8 Multipath of one ASP - 2 I/O adapters for one disk

We add the two LUNs in a new ASP 3 as shown in Figure 8-9.

```
                         Specify ASPs to Add Units to

 Specify the ASP to add each unit to.

 Specify  Serial                                      Resource
   ASP    Number       Type   Model   Capacity   Name
   ___    21-662C5     4326    050       35165   DD124
   ___    21-54782     4326    050       35165   DD136
   ___    21-65CAA     4326    050       35165   DD084
   3__    75-1310772   2105    A01        8589   DMP135
   3__    75-1311772   2105    A01        8589   DMP137
```

Figure 8-9 Multipath of one ASP - add LUNs to ASP

8.3.2 Implementing multipath in an existing environment

For implementing multipath in an existing environment consider the following:

- You may use existing FC adapters and assign the LUNs from different ASPs to multiple FC adapters. This can be done while iSeries is running. but you should carefully plan in advance which FC adapters to use for each ASP.
- You may consider purchasing additional iSeries FC adapters and additional ESS host adapters and install them for multipath.

8.3.3 Implementing multipath for an IASP

In this example we describe how to implement multipath with an independent auxiliary storage pool (IASP).

To create an IASP we use iSeries navigator. On the left panel, expand Configuration and services, expand Hardware, Disk Units, enter SST user ID and password, right-click the Disk pools, and select New disk pool from the pull-down menu. In the New disk pool wizard, we specify the type and name of a new IASP and check the box Protect data in this disk pool, as shown in Figure 8-10.

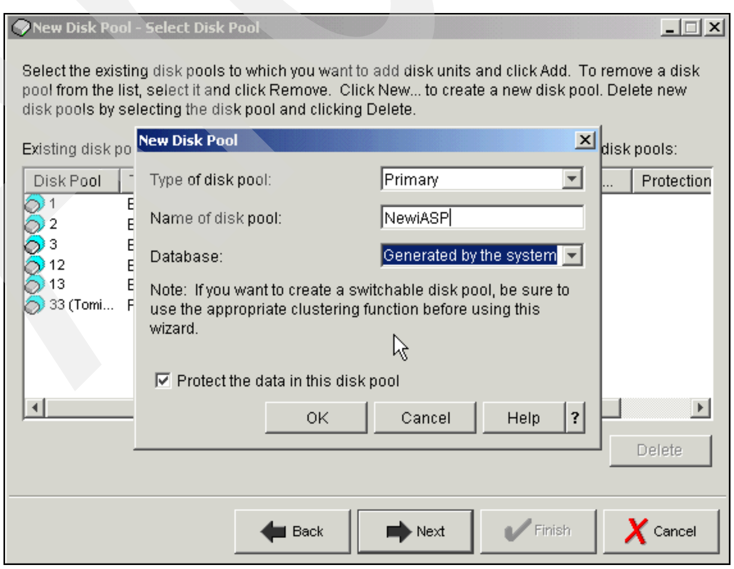

Figure 8-10 Multipath of an IASP - create IASP

In the New disk pool window, select disk pool and click Next, as shown in Figure 8-11.

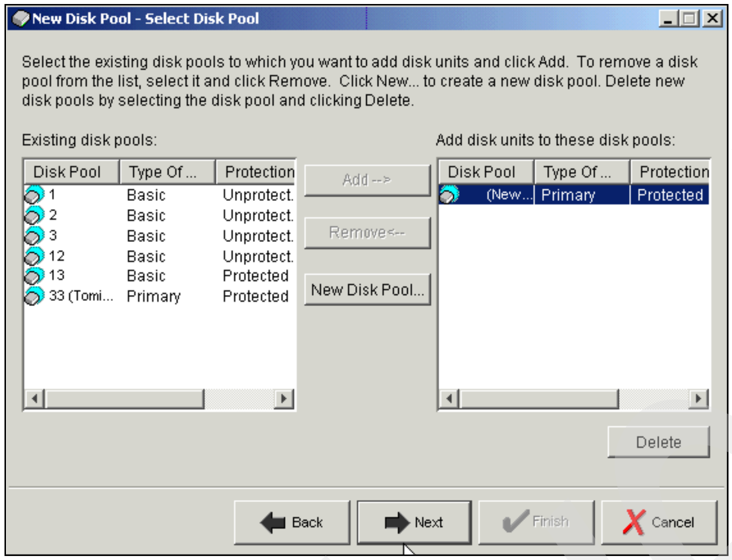

Figure 8-11 Multipath of an IASP - create IASP - continue

On the New disk pool window - add to disk pool xxx, select the type of protection, and click Next.

On the New disk pool window - add disks, select disks to be added to IASP and click Add. In this example we select two disks which are protected by multipath as can be seen from device names DMPxx. This is shown in Figure 8-12.

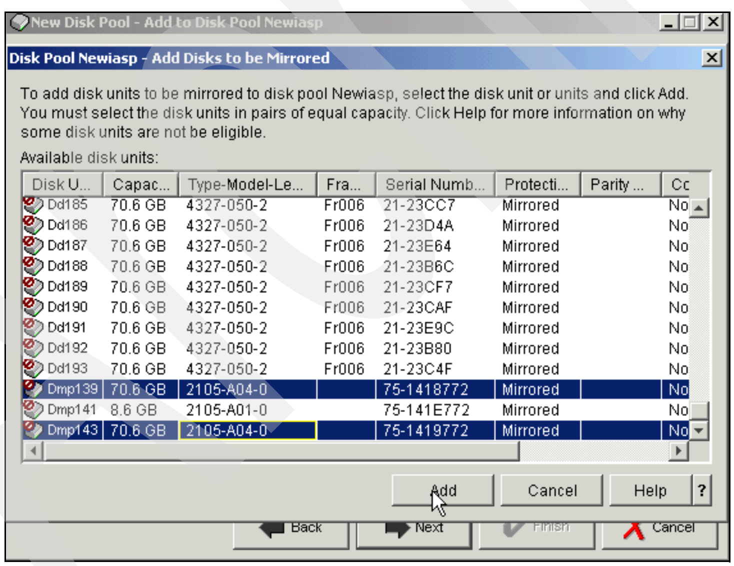

Figure 8-12 Multipath of an IASP - add multipath disks to IASP

On the New disk pool window - Summary, check that the correct LUNs will be added and click Finish, as shown in Figure 8-13.

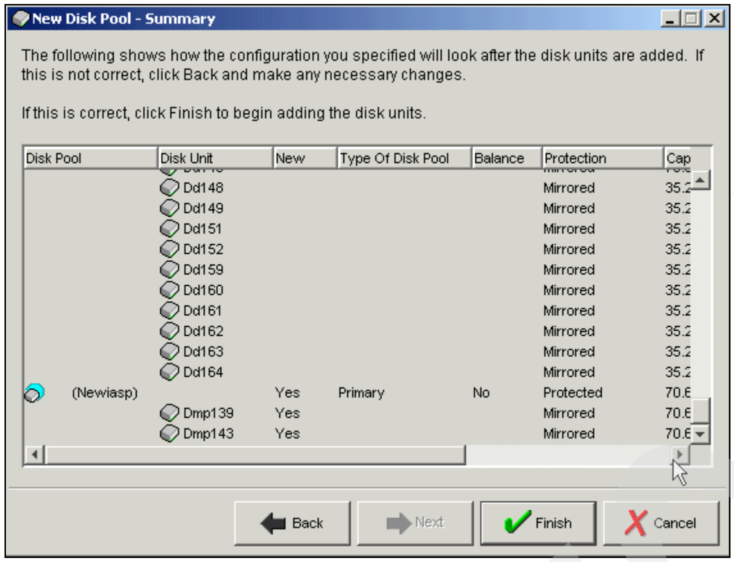

Figure 8-13 Multipath of an IASP - add multipath disks - check and finish

9

Migrating external disk from SCSI to FC

The Enterprise Storage Server (ESS) is supported on a range of iSeries and AS/400 models, operating system releases, SCSI adapters and Fibre Channel adapters. With the introduction of OS/400 Version 5 Release 1 and beyond, support for Fibre Channel is introduced on the 8xx and 270 range. The new Fibre Channel adapters feature 2787 and 2766 for disk and feature 2665 for Fibre Channel tape support. Previous releases and hardware do not support Fibre Channel. This chapter discusses the steps required to migrate from SCSI to Fibre Channel, the reasons to migrate, and it introduces the world of SAN to iSeries.

9.1 Supported hardware and software

For SCSI support, iSeries and AS/400 hosts, which can support an SPD feature 6501 SCSI attachment card running OS/400 Version 3 Release 2 for CISC systems and Version 3 Release 7 or higher for RISC systems, can attach to ESS. AS/400 hosts running Version 3 Release 1 and Version 3 Release 6 are also supported, but are limited to emulating 4.19 GB 9337 drives only.

When using an SCSI attachment, the LUNs emulate 9337 drives. Fibre Channel attached LUNs report to the iSeries as type 2105.

For Fibre Channel support the new adapters (described on page 221) for tape and DASD are required. The new Fibre Channel adapters are supported in models i5, 8xx and 270 only. Version 5 Release 1 or higher of OS/400 is required for Fibre Channel support.

The iSeries supports ESS models F800, F750, E10, E20, F10 and F20. A microcode update is required on the ESS to support iSeries and ESS copy services functions.

9.2 Why to migrate

There are a number of reasons to migrate to Fibre Channel. Some of these are listed here. This is not intended as an exhaustive list.

- Distance

  The Small Computer Systems Interface (SCSI) is limited to a maximum distance of 25 metres. Fibre Channel has a a supported distance of 500M for direct attached fibre devices and up to 10KM, via the use of hubs between source system and target storage. This allows for far greater flexibility in areas such as machine room design, sharing of tape devices between systems and the ability to back up to a remote tape drive. The diagram below illustrates some of the supported scenarios.

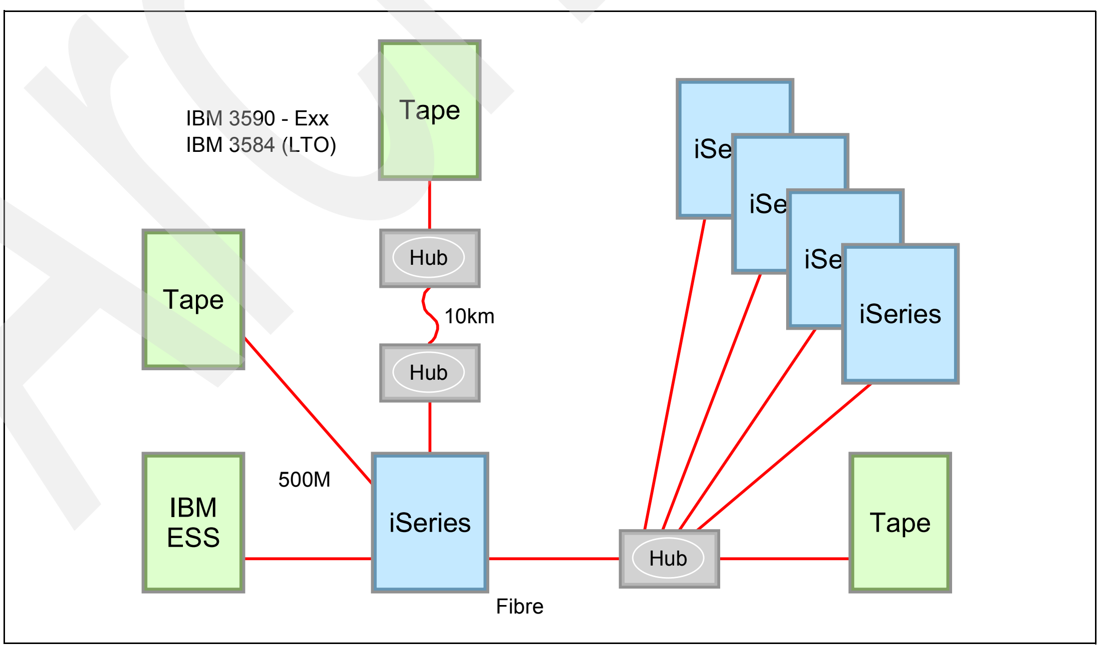

Figure 9-1 iSeries SAN

- Participation in SAN infrastructure

  A SAN is a specialized, high speed network connecting servers and storage devices. A SAN allows “any to any” connection across a network of interconnected elements such as routers, gateways, hubs and switches. Migrating to Fibre Channel allows the iSeries to participate in a SAN and realizes the flexibility and benefits this brings.

- Performance

  The new Fibre Channel interface is capable of running up to 100 MB/s full duplex. For DASD fewer Fibre Channel host adapters will be required on the ESS to achieve the same throughput as multiple SCSI cards. As a rule of thumb, between three and four SCSI 6501 DASD cards could be replaced by a single Fibre Channel adapter (use three to one if the current SCSI adapters are very busy). High end tapes can perform closer to their full potential and tests show that on average a 3590 is around 19% faster using Fibre Channel when compared to SCSI. The chart below illustrates the tape performances improvement for a mix of workloads. Individual workloads may vary. More details can be found in the Performance Capabilities manual on the Internet at:

  http://www.as400.ibm.com/developer/performance/index.html

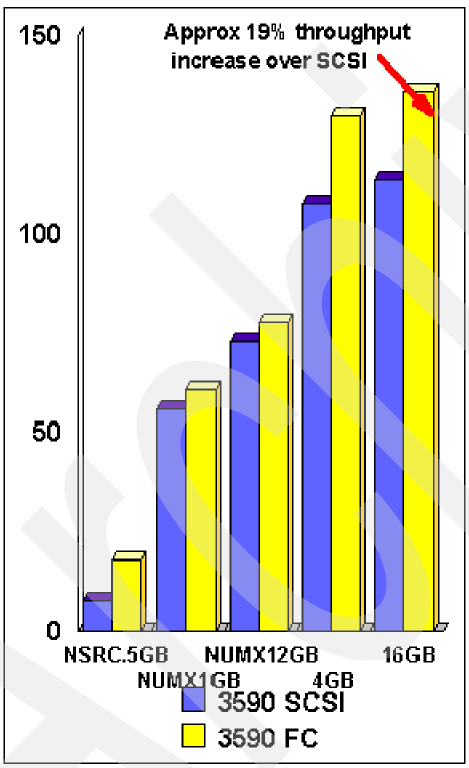

Figure 9-2 FC versus SCSI tape performance

- Number of Addressable LUNs increased

  LUNs connected via a 6501 SCSI card emulate external 9337 DASD type. A 9337 can support up to eight drives per draw and the 6501 card can support one draw for each of its two ports. Therefore the 6501 card could support a maximum of sixteen LUNs, eight on each port. With Fibre Channel this is increased to 32 LUNs. The LUNs can be of mixed sizes for both SCSI and Fibre Channel connections.

9.3 Migration considerations

The following considerations relate to migrations that include external disk and tape.

9.3.1 ESS and SAN considerations

- When performing migration from SCSI to native Fibre Channel attachment, some ESS adapters must be added or removed from the ESS I/O bays. It is necessary to put the I/O bay into service mode whenever this adapter add/removal is performed. This action will cause the other adapters within that I/O bay to be temporarily unusable, disrupting any I/O access through that bay to or from any hosts. When additional paths to those hosts have been implemented for load balancing and/or failover through any of the other I/O bays, they will be utilized. (Note: iSeries does not support failover except via host mirroring; see 9.3.3, “iSeries considerations” on page 224). When no additional paths to those same hosts exist, then I/O access from that host will be disrupted for the duration of the hardware action.

  You should minimize the hardware action to a single I/O bay at a time and minimize the elapsed time for this hardware change to reduce the impact on host systems. Proper planning should be performed to determine the impact to all host systems.

  The migration from SCSI to fibre will require new Fibre Channel cards in the iSeries and new Fibre Channel host adapters in the ESS. If required, the removal of SCSI adapters in the host and ESS needs to be performed. These actions are the responsibility of IBM Customer Engineering for ESS and the customer’s responsibility for iSeries, as well as other steps necessary to perform the migration.
- Fabric support.

  Currently the iSeries supports switched fabric and hubs in Quick loop mode. Direct Fibre Channel attachment between the iSeries and ESS is also supported.
- SAN data Gateway not supported.

  The SAN Data Gateway provides support for industry-standard Fibre Channel attachment to Intel- based servers running Windows NT and UNIX-based servers from IBM and Sun Microsystems. It allows you to extend distances between the server and SCSI-attached storage systems up to 10 kilometers. This is not supported on iSeries at this release.
- The 2103 hub is not supported by iSeries.
- Tape drives cannot be configured in a point-to-point mode, and are required to be set up in a Fibre Channel arbitrated loop (FC-AL) topology over shortwave transceivers.
- Fibre Channel cable supply and fitting is a customer responsibility.

9.3.2 StorWatch Specialist considerations

ESS StorWatch Expert will require a PTF to include Fibre Channel support. The PTF number is IP21996 and can be downloaded from the following Web site:

`http://SSDDOM01.storage.ibm.com/techsupp/swtechsup.nsf/support/essfix`

9.3.3 iSeries considerations

- A maximum of one 2787/2766 Fibre Channel DASD adapter and two 2765 tape adapters are allowed per input/output processor (IOP), feature 2743.

- A maximum of 32 LUNs are supported per Fibre Channel DASD adapter.

  A maximum of two 2765/2766/2787s (any combination) are allowed per Multi-Adapter Bridge Boundary. For a full discussion on PCI placements rules see “PCI Card Placement Rules” in the *iSeries System Handbook*, GA19-5486.
- A maximum of two 2765/2766’s (any combination) allowed per model 270 system unit. A maximum of four 2765/2766’s (any combination) allowed per model 820 system unit and 5075 expansion tower. A maximum of five 2765/2766’s (any combination) allowed per model 830 and 840 system units and 5074 expansion tower.
- A 64 bit PCI slot is preferred, though the adapters are supported in 32 bit slots.
- The adapters are supplied with an LC type cable connector. The following kits are available when connecting SC type cables to 2765 or 2766:
  a. Feature 0371 is a LC-SC adapter kit using 50um (50 Micron) Fibre containing two metres of cable. This kit contains an LC-ST cable and ST-SC adapter.
  b. Feature 0372 is a LC-SC adapter kit using 62.5um (62.5 Micron) Fibre containing two metres of cable. This kit contains an LC-ST cable and ST-SC adapter.
- Fibre Channel attached tape drives cannot be specified as Alternate IPL devices. However, they can be selected from a D-IPL using OS/400 boot manager software.
- Fibre Channel attached disk drives cannot be selected as the load/source device. You must select a suitable internal drive and controller for this function.

Note: You can select an ESS LUN to act as the remote load/source mirror.

- With one FC adapter such as the 2765 you can target 16 tape devices, or with the 2787/2766 you can target 1-32 Enterprise Storage Server devices. This could be one tape device or library for the 2765 adapter and one ESS per 2787/2766 adapter.
- Heterogeneous loops are not supported on the iSeries with V5R1 or later. However, multiple iSeries servers can access the same storage device in a so called homogeneous FC loop.
- 4 GB LUNs cannot be migrated from SCSI to Fibre Channel, since the minimum LUN size supported on Fibre Channel is 8.58 GB. This will have a particular impact on Version 3 Release 6 systems and a Version 3 Release 1 system, which can only support 4 GB LUN sizes. These users will either need to perform a scratch installation or can use ASP storage management functions in SST or DST to achieve migration.
- OS/400 supports native mirroring. By defining the LUNs as “unprotected” and starting mirroring via OS/400, this could be used to provide failover support. It is also possible to have a redundant path between the ESS and a hub, though this does not provide a load/balancing solution.
- Independent Auxiliary Storage Pools (IASPs) were introduced with the introduction of V5R1 of OS/400. IASPs are supported on external drives.

9.4 Migration process from SCSI to Fibre Channel

Migration from SCSI to Fibre Channel is very straight forward. We have several scenarios below for migrating from SCSI to Fibre Channel. In most cases the LUNs can be simply re-assigned from the SCSI host adapter to the Fibre Channel adapters with the data intact. In this case you should follow the steps in Scenario 1 on page 226.

Scenario 2 on page 233 describes a migration using PPRC combined with a load/source mirror and load/source recovery.

Scenario 3 on page 235 assumes a box swap.

Scenario 4 on page 241 assumes an iSeries upgrade, combined with a migration to SCSI.

Scenario 5 on page 242 uses standard ASP management tools to migrate the data to the new Fibre Channel LUNs.

Be sure to check the iSeries, ESS, and SAN service support Web pages for the latest information at:

- `http://www.as400service.ibm.com/`
- `http://www.storage.ibm.com/ess`
- `http://www.storage.ibm.com/ibmsan/index`

As is normal with any major system change, we strongly recommend that *at least* one full system back up is performed prior to the changes.

9.4.1 Scenario 1 - SCSI to Fibre Channel LUN re-assignment

In this example we upgrade an existing iSeries system from SCSI to support Fibre Channel. The LUNs will be re-assigned to Fibre Channel without the need for a full reload. If insufficient host adapter bays exist in the ESS for the SCSI and Fibre Channel adapters to coexist, you will have to remove the SCSI adapters to install the Fibre Channel adapters.

We expect most customers who want to migrate from SCSI attached ESS LUNs to Fibre Channel will use this procedure. This procedure could also be used for environments with a mix of internal and external DASD. In this case, the internal, including the load/source, will not change.

> **Note:** When re-assigning LUNs, all LUNs must be re-assigned, otherwise the system will be unusable. Single Level Storage will place data on all available drives, so all available LUNs must be re-assigned to ensure single level storage is intact.

Advantages

- Simple upgrade process on the iSeries
- No need for a full save/restore
- Simple migration of LUNs on ESS
- Fairly quick migration, typically 2-3 hours
- No need for duplicate storage on ESS

Upgrade Process

- Upgrade the iSeries source system to OS/400 V5R1 or greater.
- Install new Fibre Channel adapters in the iSeries.
- Upgrade the ESS microcode to a level supporting Fibre Channel.
- Install the new host Fibre Channel adapters in the ESS. If installing multiple Fibre Channel adapters, we recommend they are installed in different I/O bays supported by different RISC processors.
- Power down the iSeries before reconfiguring the ESS for Fibre Channel.

- From the ESS Specialist perform the following tasks.

1. At the Specialist main menu select Storage Allocation. Since all data between the ESS and the ESS Specialist is encrypted you will have to proceed through the Site Certificate screens. For additional security you will be prompted for user name and password. See *Web Interface Users Guide for the ESS Specialist and ESS Copy Services*, SC26-7346 for more details.

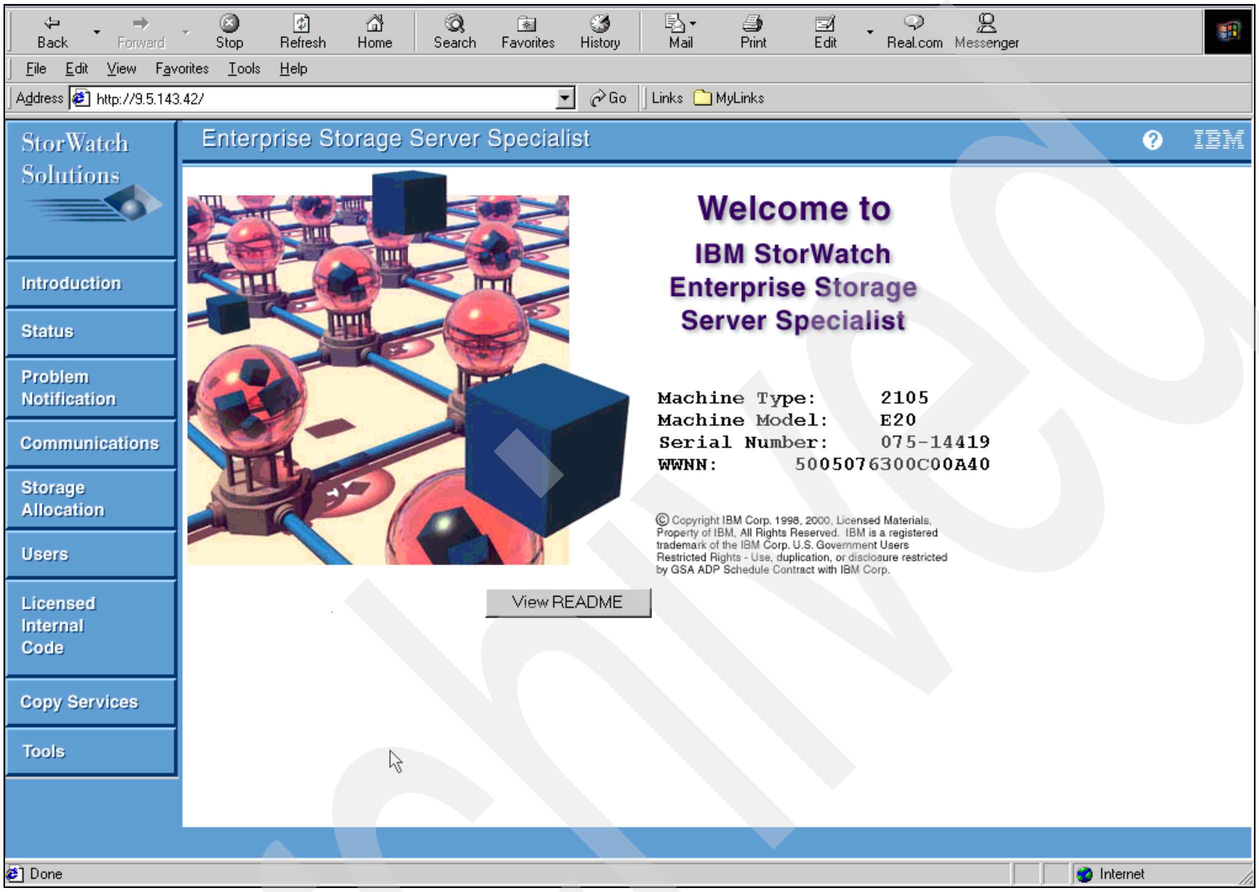

Figure 9-3 StorWatch Specialist Welcome

Note: The default supplied user name is *storwatch*. The default password is *specialist.*

2. At the StorWatch Allocation screen select Open Systems Storage.

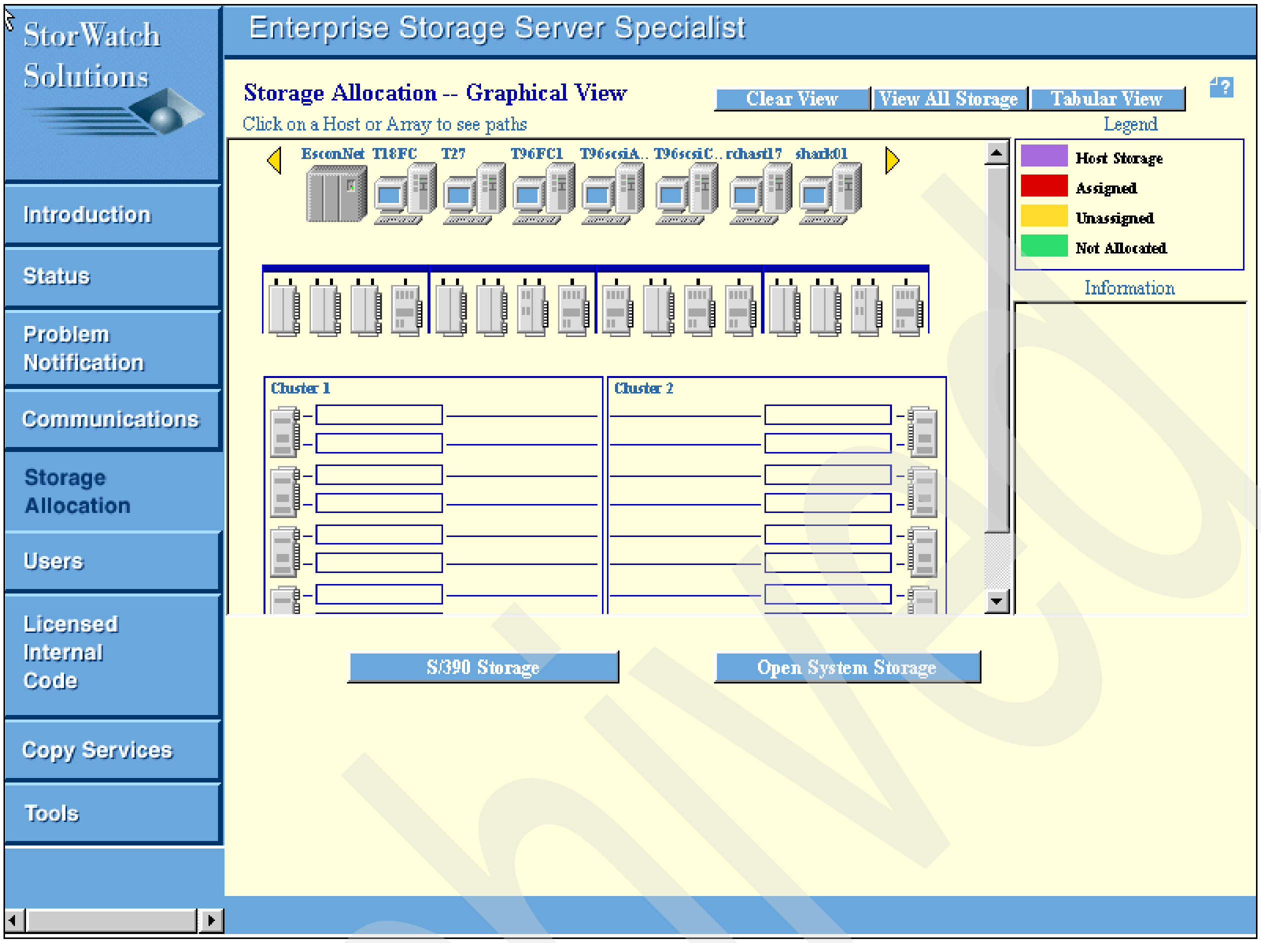

Figure 9-4 Storage Allocation

3. At the Open Systems Storage Screen, select Modify Host Systems.

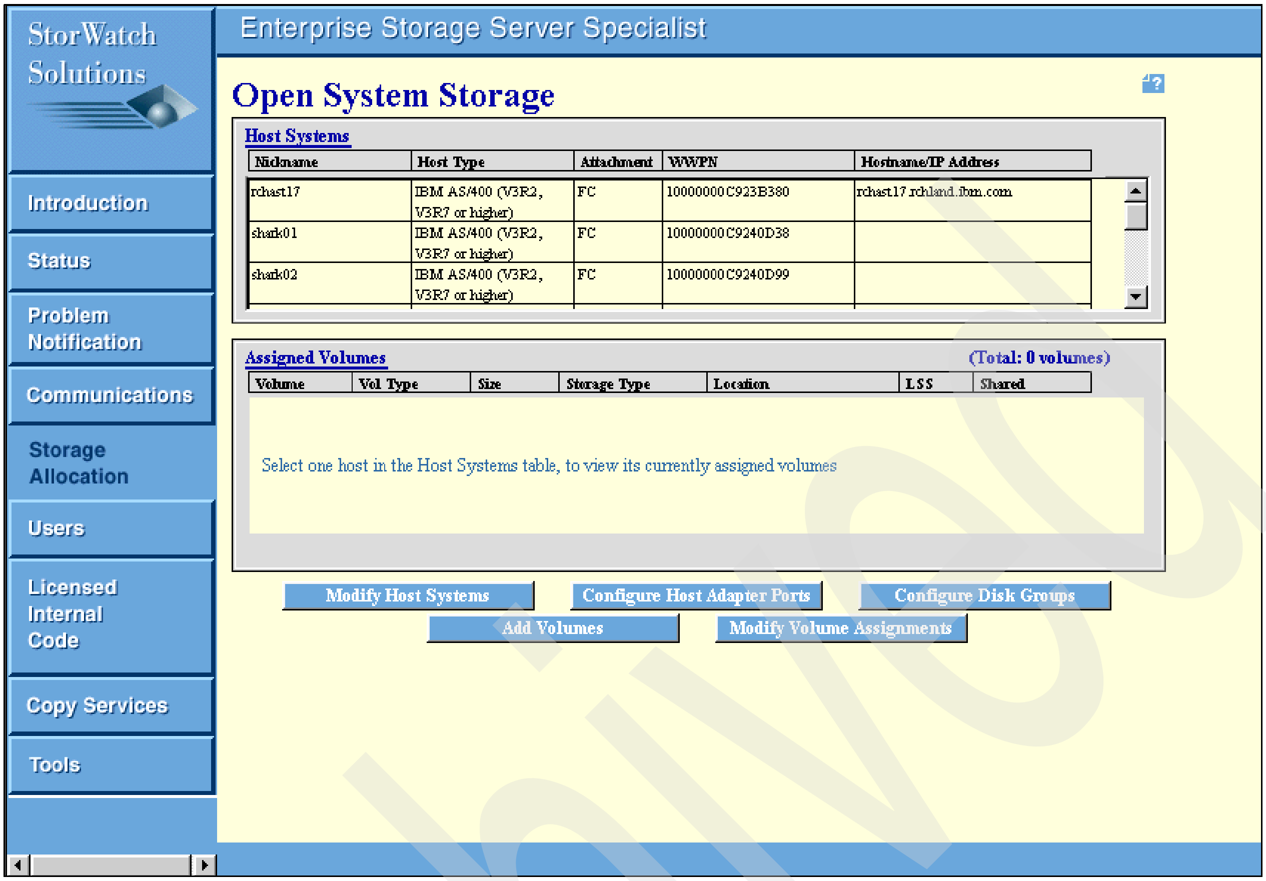

Figure 9-5 Open System Storage

a. In the host name field, choose a host name.

> **Note:** Each Fibre Channel connection on the same AS/400 will need a unique host name.

b. In the Host Type field, choose “AS/400 V3R2, V3R7 or higher” from the drop down list.

c. In the Host Attachment field, choose “Fibre Channel Attached” from the drop down list.

d. The hostname/IP address field is optional.

e. Select the World-Wide-Port-Number of the new Fibre Channel adapter in the iSeries you want configure. This field identifies the WWPN of the attachment card inside the iSeries.

f. Select the Fibre Channel port in the ESS. This field identifies the Fibre Channel bay and adapter position inside the ESS.

> **Note:** For alternate path support (for example via a hub or switch) select two or more ESS Fibre Channel ports by holding down the shift key. You can also select all installed adapters.

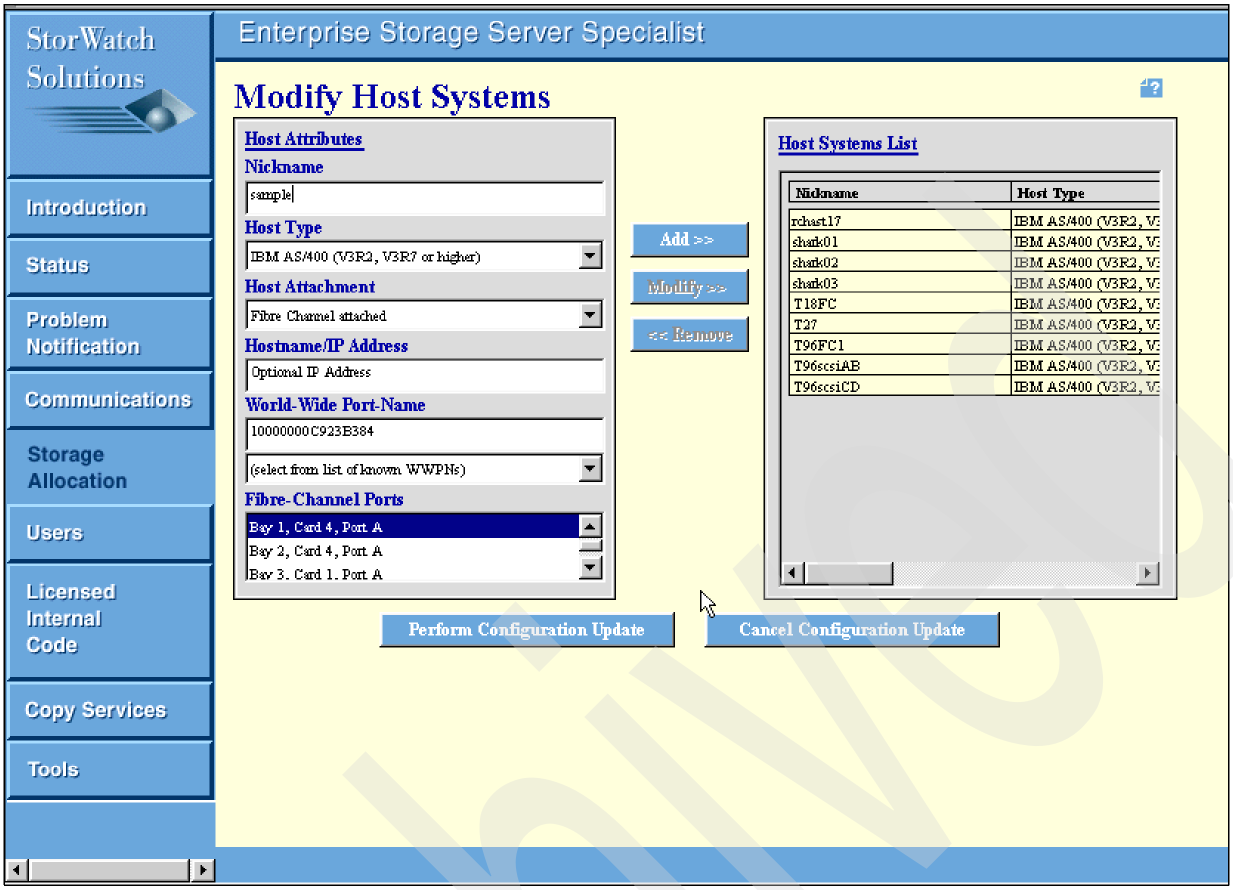

Figure 9-6 Add Port to Host

g. When you have finished completing the fields, select ADD.

h. Repeat the above tasks for additional Fibre Channel connections.

i. Finally, select Perform Configuration Update to process the configuration changes.

4. You will now be at the Open System Storage display. Select “modify Volume assignments”. In the next steps we will unasign the SCSI attached volumes and then re-assign them as Fibre Channel attached volumes.

   a. To more easily identify the iSeries SCSI volumes you want to unassign, sort the table by host system name by selecting “first” from the sort drop down list field above host nickname. Then select “perform sort”.

   b. Click on the LUNs you want to unassign. To select multiple LUNs hold down the shift key.

   c. Select “unassign selected volumes from target host”.

   d. Select “Perform configuration update”.

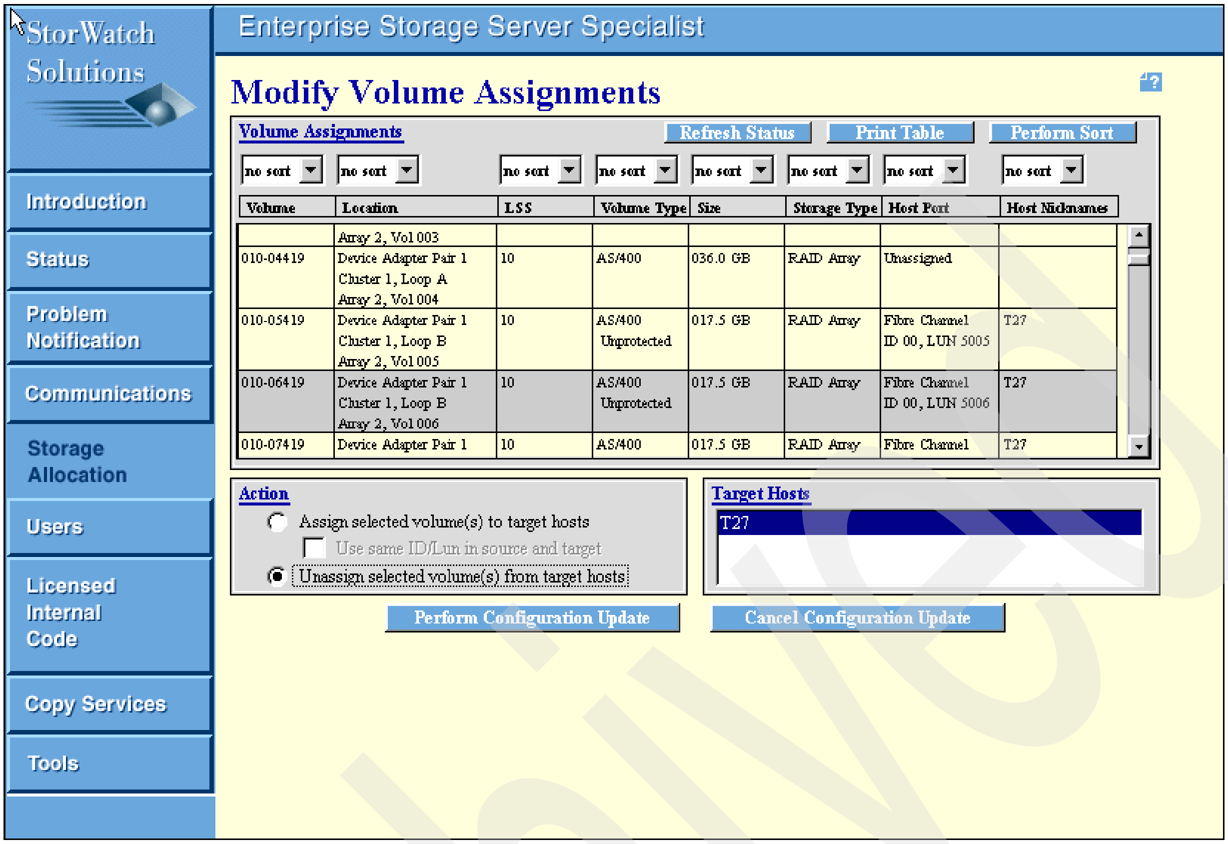

Figure 9-7 Open System Storage display

5. We will now be adding the LUNs to the Fibre Channel connection.
   a. To more easily identify the iSeries volumes you want to assign sort the table by host system name by selecting “first” from the sort drop down list field above host nickname. Then select “perform sort”.
   b. Click on *all* the LUNs from the source system. To select multiple LUNs hold down the shift key.

Note: You must select all drives from the source system. Due to single level storage, the data will be on all drives. Failure to select all drives will result in an unusable system.

   c. Select “assign selected volumes to target host”.
   d. In the Target Hosts box select the “iSeries host”.
   e. Select Perform Configuration Update to complete this task.
6. You may optionally remove the SCSI host adapter definitions.
   a. From the Open Systems Storage display select Configure Host Adapter Ports.
   b. At the top of the screen highlight the SCSI port to be removed from the ESS configuration.

Note: SCSI adapters have two ports and both may need to be removed, however, the two available SCSI ports on the ESS adapter could be used by different host systems. Be careful to select the correct ports.

c. In the “Map of SCSI id bus” box highlight your iSeries or AS/400 host.

d. Press the remove button and repeat the above step for each SCSI host adapter.

e. Select Perform configuration Update to complete the changes.

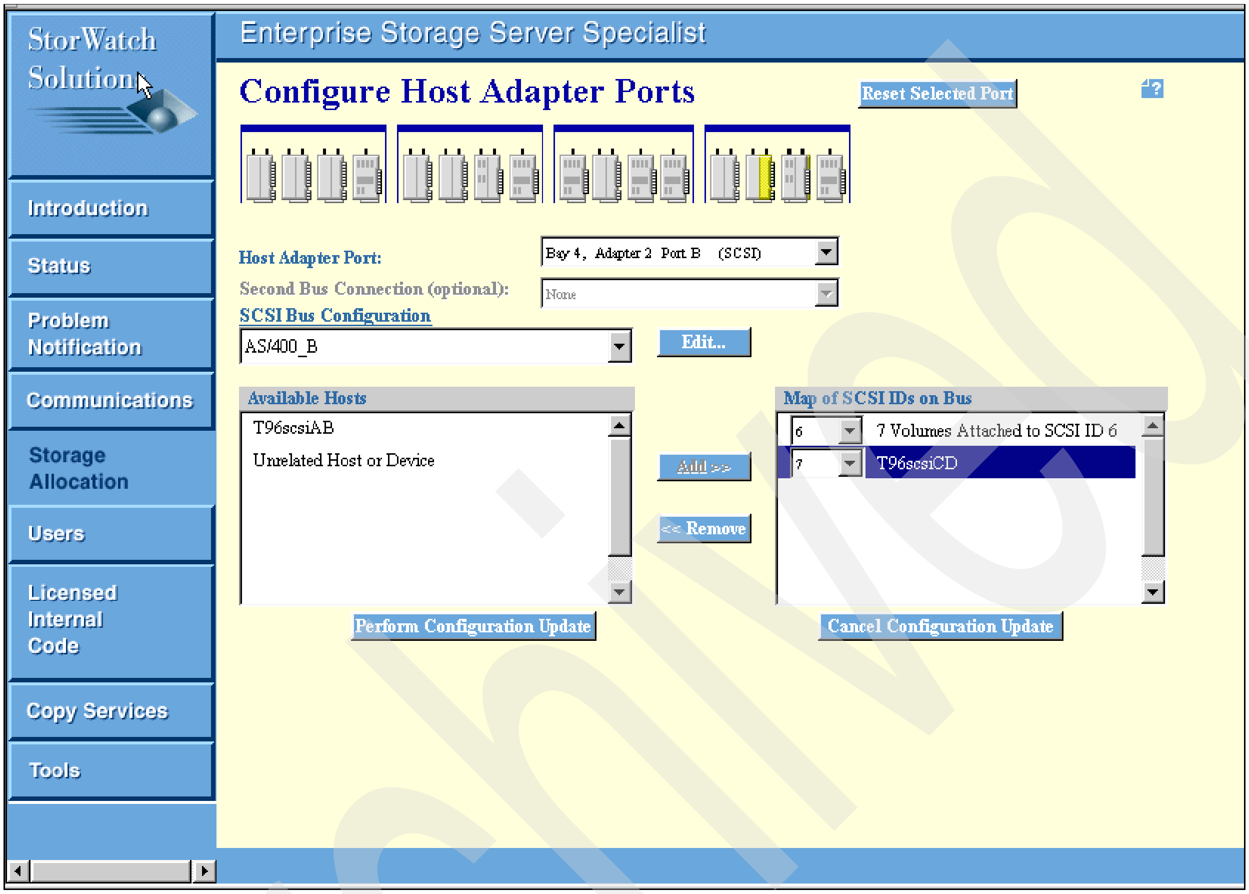

Figure 9-8 Configure Host Adapters

If you choose to do so, the original/source SCSI adapters may be removed from the ESS.

9.4.2 Scenario 2 - Box swap using PPRC

In this scenario a new iSeries system is ordered to replace an existing iSeries or AS/400. The new box is installed along side the existing iSeries or AS/400 system. This involves setting up PPRC from StorWatch Copy Services. You will then break the PPRC connection and perform a Remote Load Source Recovery on the new iSeries to recover the data, system and all the LUNs without the need for a re-load, and activate the new load/source internal drive. Remote load/source setup and recovery is fully discussed in Chapter 7, “Mirroring to an ESS” on page 165. This will require a second set of LUNs to act as the PPRC target.

Advantages

- You can easily create two identical systems and repeat the process many times without impact to the live system. This enables you to prove the concept and reduce the risk to the business.
- The second system can be used for other functions such as testing, user training, software testing or development.

- This process avoids the need for a full system save and restore.
- The downtime is minimized. Typically the recover load/source takes between 1-2 hours; much faster than a full scratch install of the iSeries.

Considerations

- Resource mappings may have to be changed. For example, the Token Ring line may use resource CMN11 on the old system, but requires resource CMN03 on the new system. You will have to manually fix this, by changing the appropriate configuration object. This would be required on any migration between two iSeries.
- All storage must be defined in the ESS including the load/source mirror, to ensure the entire system can be migrated.
- PPRC must be purchased, if not already used.
- Other devices, such as tape drives may assume new identities, since they have new resources. You should consider application dependencies on these devices.
- If both source and target systems are to be used at the same time, then system network configurations will have to be changed on one system (for example, IP addresses and system name).
- You will need double the original storage requirement on the ESS; storage must be available to accommodate this.
- The source load/source must be equal or smaller than the target load/source.
- Remote load/source recovery is relatively complex.

Upgrade process

The process is summarized here:

- Remote load/source mirroring must be set up on the source system prior to the migration. The remote mirror must be a LUN defined in the ESS. This is because single level spreads data across all available drives, including the load/source. See chapter Chapter 7, “Mirroring to an ESS” on page 165 for instructions on load/source mirroring.
- All DASD must be defined in the ESS including the load/source mirror to ensure the entire system can be migrated.
- Ensure the new iSeries system has OS/400 V5R1 or greater.
- Install new Fibre Channel adapters in the iSeries.
- Upgrade the ESS microcode to a level supporting Fibre Channel.
- Install the new host Fibre Channel adapters in the ESS. If installing multiple Fibre Channel adapters, they should be installed in different I/O bays supported by different RISC processors.
- Set up load/source mirroring on the source system as described in Chapter 7, “Mirroring to an ESS” on page 165.
- Define the new LUNs to the new target host as described in 12.2, “Peer-to-Peer Remote Copy” on page 329.
- Setup PPRC between the source and new target LUNs as described in 12.2, “Peer-to-Peer Remote Copy” on page 329. All LUNs must be in the PPRC configuration. All data must reside on the ESS.
- Power down the source system. This is to ensure all data is written to disk. This ensures data integrity.

- Terminate the PPRC connection as defined in 12.2, “Peer-to-Peer Remote Copy” on page 329.
- Power on the source system if it is still required.
- Perform the load/source recovery on the target system as described in 11.2.1, “Recovering the Remote Load Source” on page 310.
- Fix resource issues and system configuration, if required.
- Perform, as a minimum, a save of SAVSYS on the on the new system to ensure all changes are saved.
- You now have two systems with identical data (apart from resourcing and network changes).

9.4.3 Scenario 3 - Box swap

In this scenario a new iSeries is purchased to replace an existing system. For example, the new system is a 640 with SCSI attached LUNs, the new system is an 830, with Fibre Channel support.

We have two options:

1. Perform a full save and scratch install to recover the data on the new system. This would involve assigning LUNs of sufficient capacity on the new system (see Chapter 6, “Creating iSeries storage in ESS” on page 127) and then following the normal save restore process. Since this is a normal save/restore procedure we will not attempt to describe this here. Refer to *The Backup and Recovery Version 5*, SC41-5304 for details or follow your normal disaster recovery plans if these are documented and tested.
2. Re-assign the existing LUNs from SCSI to Fibre Channel and then perform a load source migration on the target system. Remote load/source mirroring must be enabled to ensure the entire system is migrated. All data must reside in the ESS. All LUNs must be migrated. We describe this situation here.

Advantages

- Avoids the need for a full system save and restore.
- The second system can be used for other things, such as testing, user training, and software testing.

Considerations

- Resource mappings may have to be changed. For example, the Token Ring line may use resource CMN11 on the old system, but requires resource CMN03 on the new system. You will have to manually fix this, by changing the appropriate configuration object. This would be required on any migration between two iSeries.
- Other devices, such as tape drives may assume new identities, since they have new resources. You should consider application dependencies on these devices.
- The source load/source must be equal or smaller than the target load/source.
- All data must reside in ESS LUNs, including a mirrored copy of the load/source.

Upgrade Process

- Ensure the new iSeries system has OS/400 V5R1 or greater.
- Install new Fibre Channel adapters in the iSeries.
- Upgrade the ESS microcode to a level supporting Fibre Channel.

- Install the new host Fibre Channel adapters in the ESS. If installing multiple Fibre Channel adapters, they should be installed in different I/O bays supported by different RISC processors.
- Set up load/source mirroring on the source system as described in Chapter 7, “Mirroring to an ESS” on page 165.
- Define the new LUNs to the new target host as described in Chapter 12, “Peer-to-Peer Remote Copy” on page 327.
- Ensure all data resides on ESS LUNs. The internal load/source drive will have been mirrored to a LUN is the ESS.
- Power down the source system. This is to ensure all data is written to disk. This ensures data integrity.
- From the ESS Specialist perform the following tasks.

1. At the Specialist main menu, select Storage Allocation. Since all data between the ESS and the ESS Specialist is encrypted you will have to proceed through the Site Certificate screens. For additional security you will be prompted for user name and password. See *Web Interface Users Guide for the ESS Specialist and ESS Copy Services*, SC26-7346 for more details.

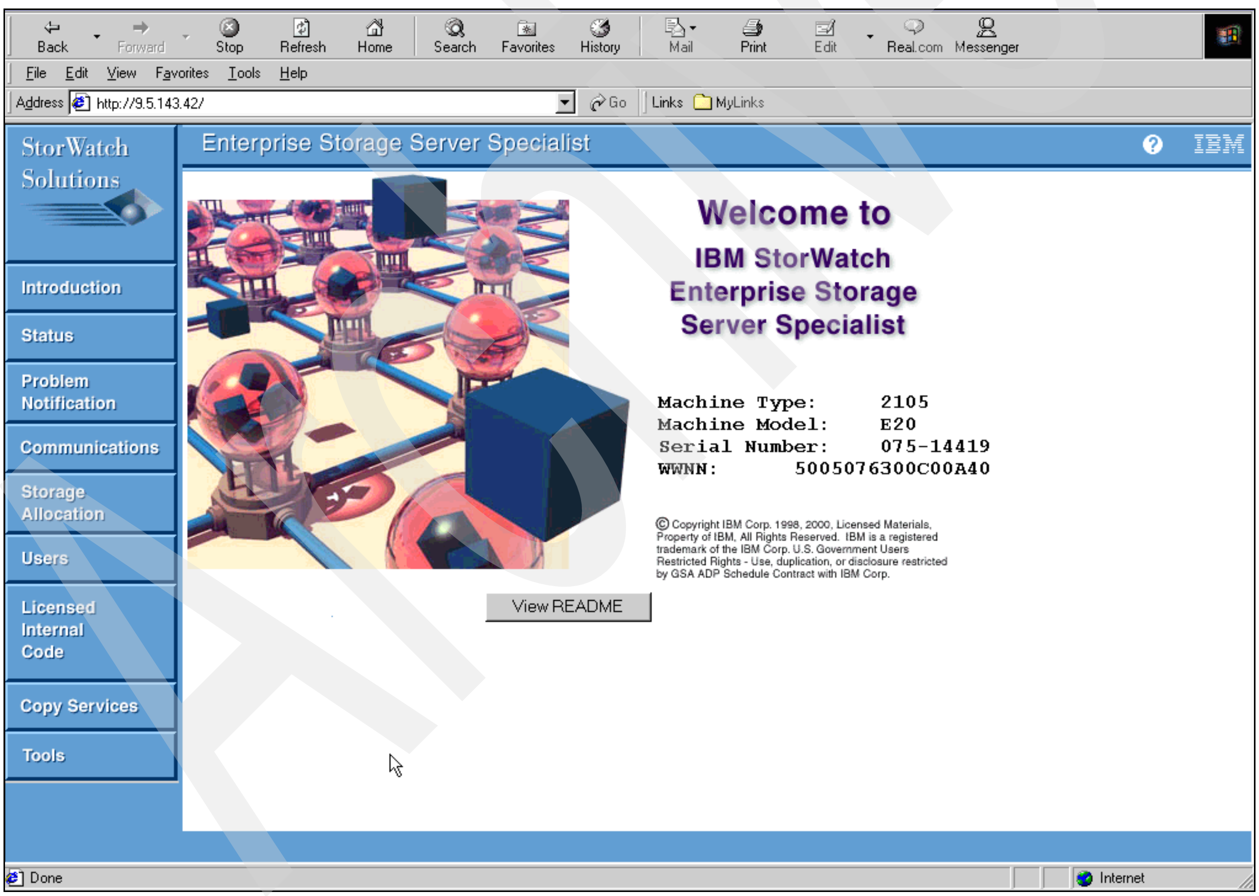

Figure 9-9 StorWatch Specialist Welcome

> **Note:** The default supplied user name is *storwatch*. The default password is *specialist.*

2. At the StorWatch Allocation screen, select Open Systems Storage.

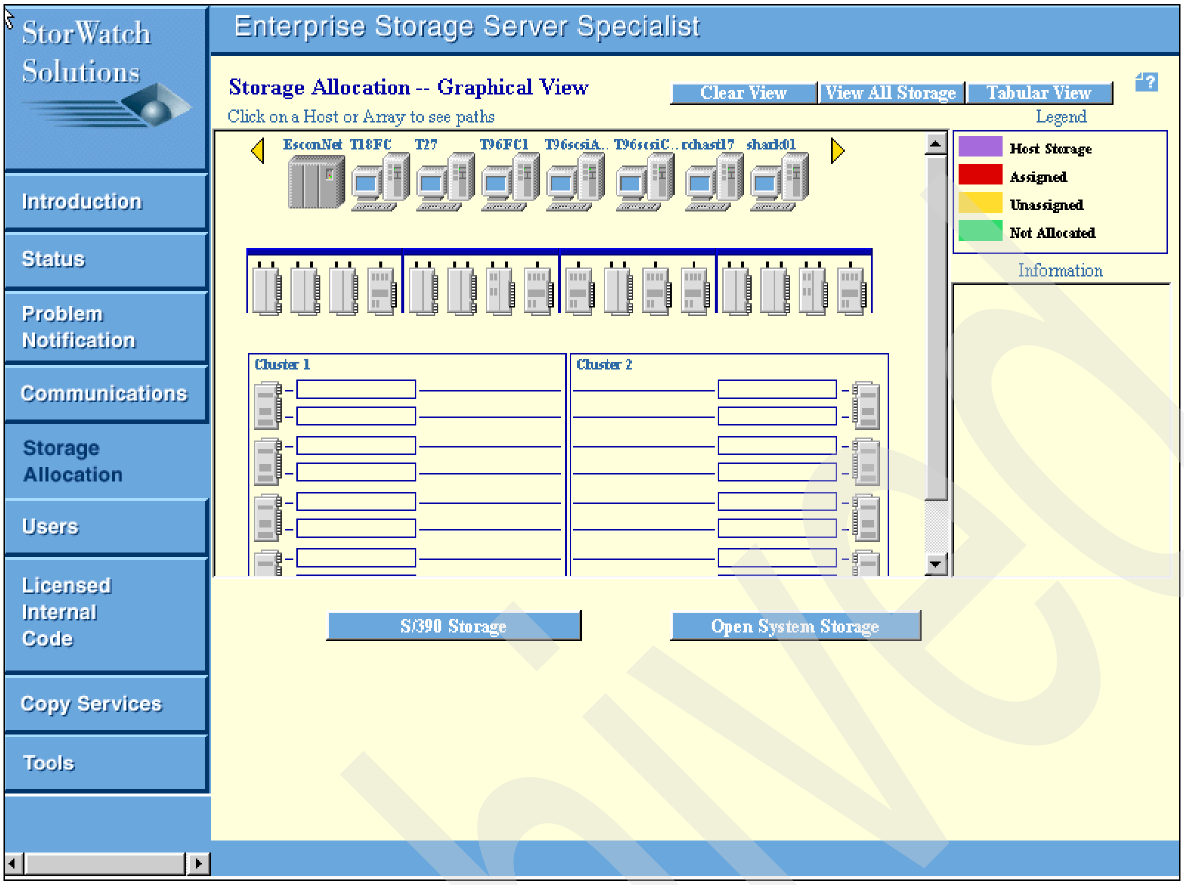

Figure 9-10 Storage Allocation

3. At the Open Systems Storage Screen, select Modify Host Systems.

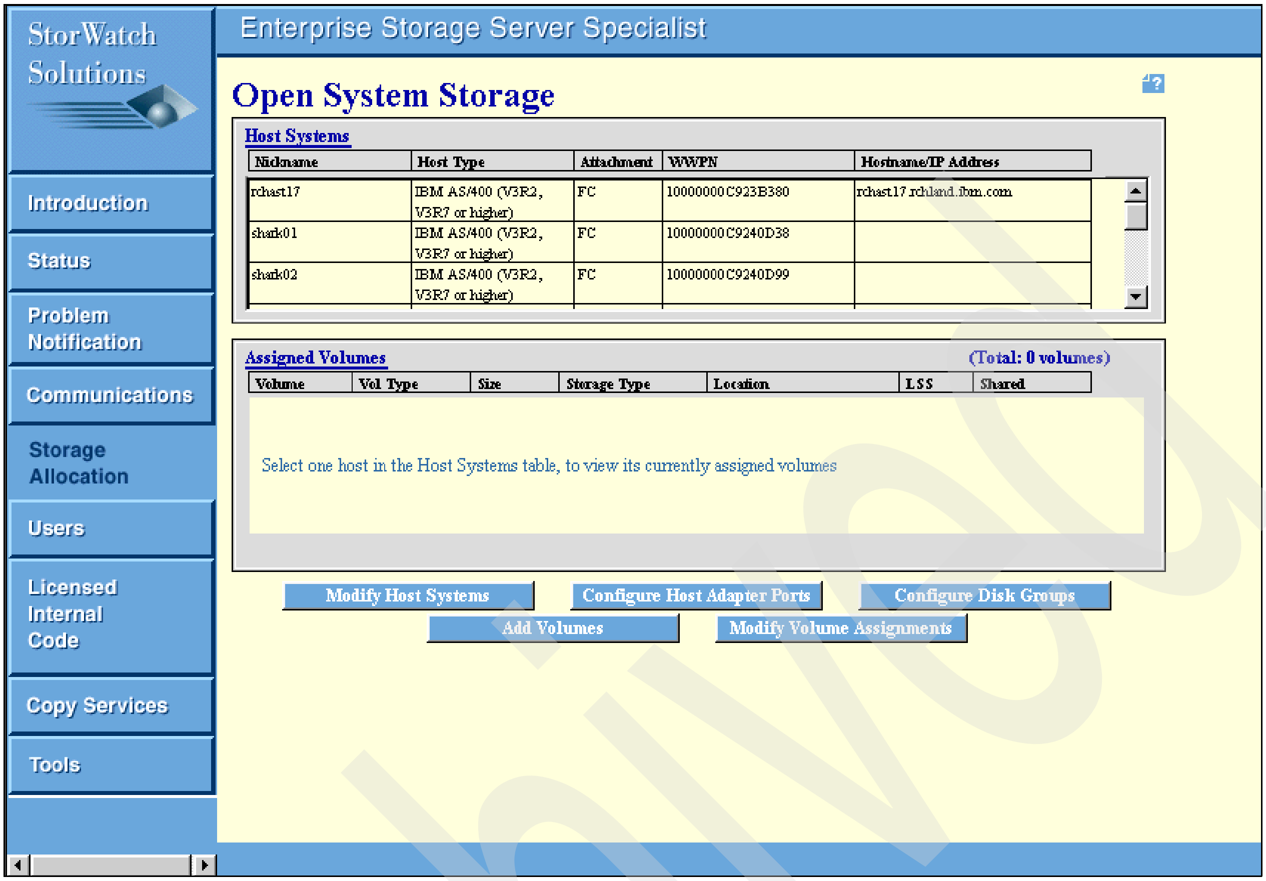

Figure 9-11 Open System Storage

a. In the host name field, choose a host name.

> **Note:** Each Fibre Channel connection on the same AS/400 will need a unique Host name.

b. In the Host Type field choose “AS/400 V3R2, V3R7 or higher” from the drop down list.

c. In the Host Attachment field choose “fibre Channel Attached” from the drop down list.

d. The hostname/IP address field is optional.

e. Select the World-Wide-Port-Number of the new Fibre Channel adapter in the iSeries you want configure. This field identifies the WWPN of the attachment card inside the iSeries.

f. Select the Fibre Channel port in the ESS. This field identifies the Fibre Channel bay and adapter position inside the ESS.

> **Note:** For alternate path support (for example, via a hub or switch) select two or more ESS Fibre Channel ports by holding down the shift key. You can also select all installed adapters.

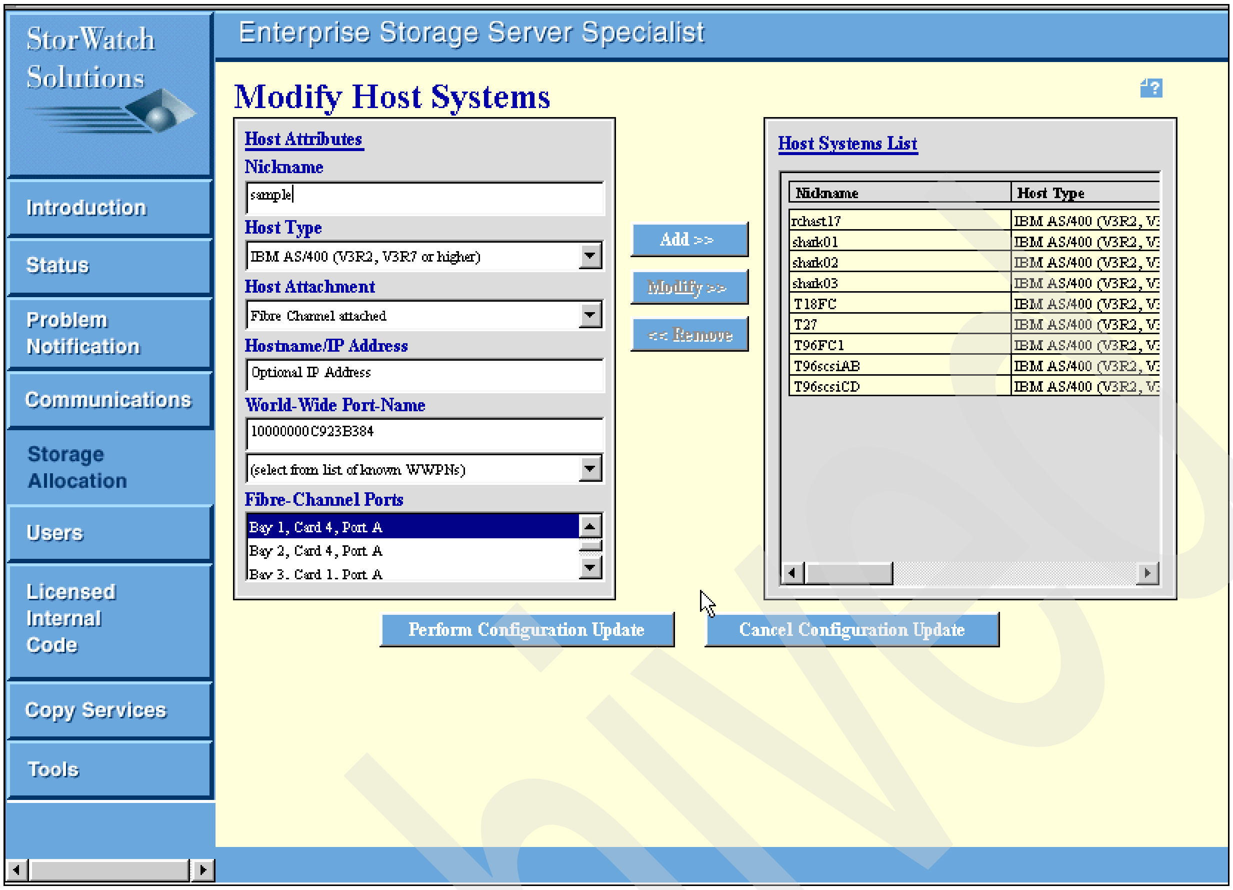

Figure 9-12 Add Port to Host

- g. When you have finished completing the fields, select ADD.
- h. Repeat the above tasks for additional Fibre Channel connections.
- i. Finally, select Perform Configuration Update to process the configuration changes.

4. You will now be at the Open System Storage display. Select “modify Volume assignments”. In the next steps we unassign the SCSI attached volumes and then re-assign them as Fibre Channel attached volumes.
   - j. To more easily identify the iSeries SCSI volumes you want to unassign, sort the table by host system name by selecting “first” from the sort drop down list field above host nickname. Then select “perform sort”.
   - k. Click on the LUNs you want to unassign. To select multiple LUNs hold down the shift key.
   - l. Select “unassign selected volumes from target host”.

   Select “Perform configuration update”.

> **Note:** You must select all drives from the source system. Due to single level storage, the data will be on all drives, failure to select all drives will result in an unusable system.

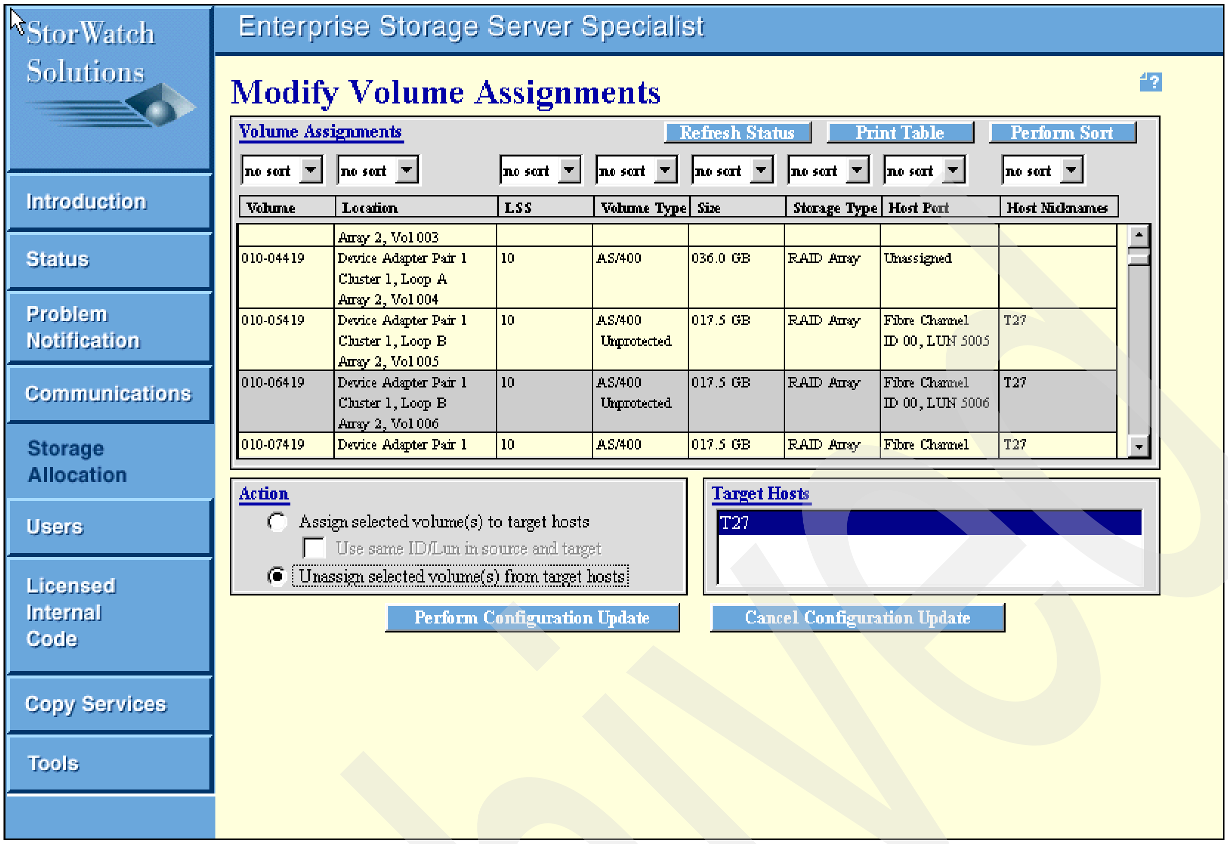

Figure 9-13

5. We now add the LUNs to the Fibre Channel connection.
   a. To more easily identify the iSeries volumes you want to assign, sort the table by host system name by selecting "first" from the sort drop down list field above host nickname. Then select "perform sort".
   b. Click on *all* the LUNs from the source system. To select multiple LUNs hold down the shift key.

> **Note:** You must select all drives from the source system. Due to single level storage, the data will be on all drives. Failure to select all drives will result in an unusable system.

   c. Select "assign selected volumes to target host".
   d. In the Target Hosts box, select the "iSeries host".
   e. Select Perform Configuration Update to complete this task.
6. Now we need to remove the SCSI host adapter definitions.
   a. From the Open Systems Storage display, select Configure Host Adapter Ports.
   b. At the top of the screen highlight the SCSI port to be removed from the ESS configuration.

> **Note:** SCSI adapters have two ports both may need to be removed, however, the two available SCSI ports could be used by different host systems. Be careful to select the correct ports.

c. In the “Map of SCSI id bus” box, highlight your iSeries or AS/400 host.

d. Press the remove button and repeat the above step for each SCSI host adapter.

e. Select Perform configuration Update to complete the changes.

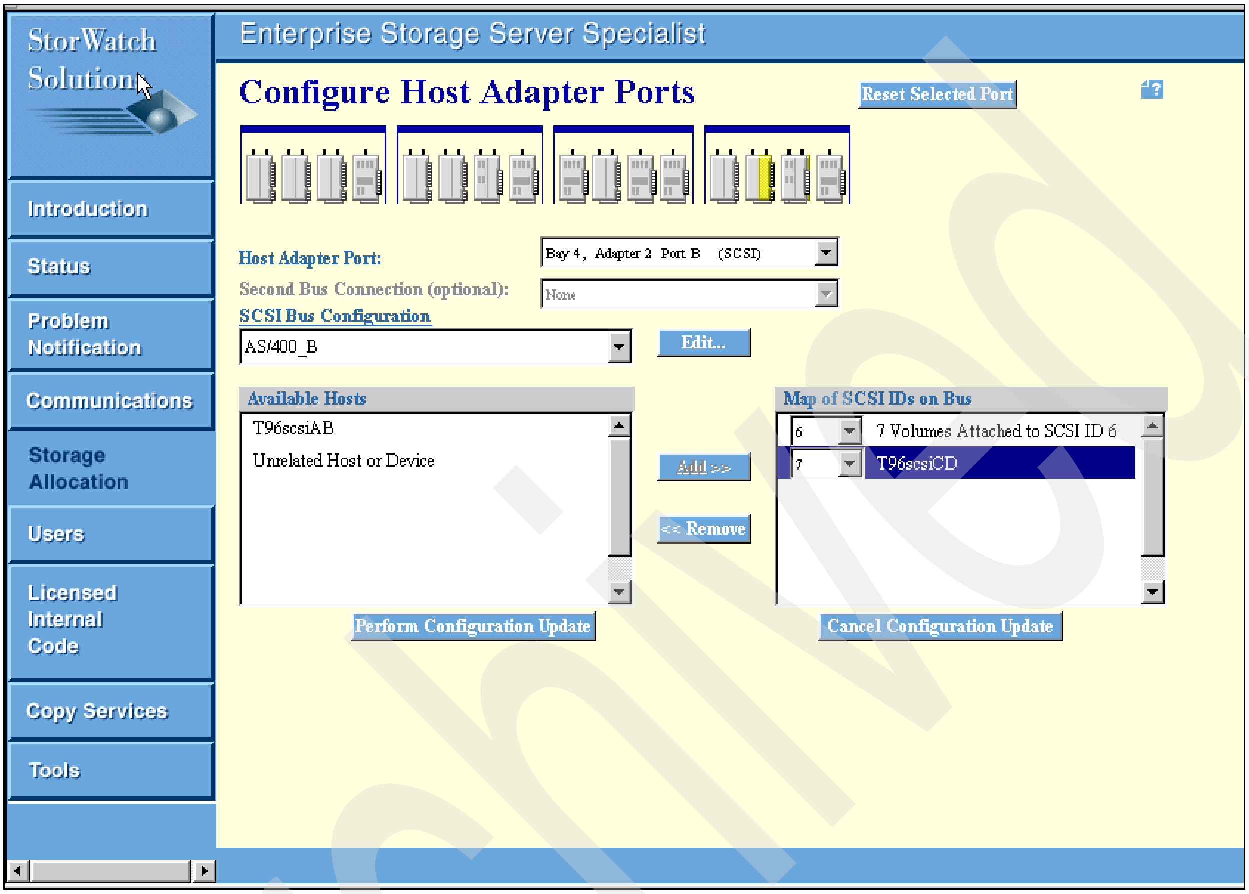

Figure 9-14 Configure Host Adapters

- If you choose to do so, the original/source SCSI adapters may be removed from the ESS.
- Perform the load/source recovery as described in 11.2.1, “Recovering the Remote Load Source” on page 310. This is because we must migrate the mirrored copy of the load/source to a supported internal drive.

Note: The iSeries only supports an internal load/source. This step must be completed or you will be unable to IPL your system.

- Fix resource issues and system configuration, if required.
- Perform, as a minimum, a save of SAVSYS on the on the new system to ensure all changes are saved.

9.4.4 Scenario 4 - iSeries upgrade combined with migration from SCSI to Fibre Channel

In this scenario an existing iSeries will be upgraded while at the same time Fibre Channel adapters are added to accommodate a change from SCSI to Fibre Channel. The process to achieve this is identical to that described in Scenario 1. You simply need to re-assign the LUNs from the SCSI connection to the Fibre Channel connection and perform a load/source migrate.

We suggest that the upgrade be phased to avoid making too many changes at once.

9.4.5 Scenario 5 - Migration from SCSI or internal drives to Fibre Channel using ASP storage management

In this scenario we use native iSeries tools to migrate the data from SCSI attached LUNs to the newly attached Fibre Channel attached LUNs. This example assumes a single iSeries systems which has Fibre Channel adapters added along side existing SCSI adapters. If you have all storage, including the remote load/source in the ESS you will more likely use Scenario 1. If you have a mix of internal and external storage, or a very small system then this maybe more appropriate.

Advantages

- A full save/restore is not required.
- Data can be migrated piece meal to suit the production/live requirements.
- Avoids the need for the relatively complex remote load source recovery and remote load/source setup.
- Can be used to migrate existing internal DASD and/or external SCSI attached DASD.

Considerations

- This is likely to take longer than a re-assigning LUNs depending on the number of DASD arms involved and amount of data involved.
- Additional storage will be required in the ESS if the SCSI storage is already installed in the ESS.
- Spare host adapter bays may be required.

Upgrade process

- Upgrade the iSeries system to OS/400 V5R1 or greater.
- Install new Fibre Channel adapters in the iSeries.
- Upgrade the ESS microcode to a level supporting Fibre Channel.
- Install the new host Fibre Channel adapters in the ESS. If installing multiple Fibre Channel adapters, we recommend they should be installed in different I/O bays supported by different RISC processors.
- Refer to Chapter 6, “Creating iSeries storage in ESS” on page 127 for ESS configuration details.
- You will now have all new LUNs in an nonconfigured state. When you (or your service representative) physically attach a new disk unit to your system, the status of the new disk unit is nonconfigured. Nonconfigured status means that a disk unit has not yet been assigned to an ASP on the system. You can assign disk units to an existing ASP or to a new ASP. You create a new ASP simply by assigning disk units to it.

To assign nonconfigured disks to an ASP, do the following:

- If you are not already using DST, perform a manual IPL to start DST.
- From the Use Dedicated Service Tools (DST) menu, do the following:

  – Select option 4 - Work with disk units
  – Select option 1 - Work with Disk Configuration
  – Select option 3 - Work with ASP Configuration
- **Or**, from the System Service Tools (SST) menu, do the following:
  – Select option 3 - Work with disk units
  – Select option 2 - Work with disk configuration
- Select the option to add units to ASPs. The specify ASPs to add unit to display appears. It lists all the disk units that have a status of nonconfigured.

> **Note:** All the disk units on the system may not report in right away. Verify that the number of disk units and LUNs displayed matches the number of disk units physically attached to the system. If they do not match, wait a few minutes and press F5 (Refresh) until all of the disk units report in.

- Type an ASP number next to each disk unit that you want to configure and press Enter.

  The Confirm Add Units display shows what the entire system configuration will be when you add the units. If you have more than one ASP on your system, verify this configuration against your planned configuration.
- You can press F9 (Resulting capacity) to see how the change will affect your disk utilization. Press F12 (Cancel) to return to the Confirm Add Units display.
- If you are satisfied with the configuration, press Enter to add the disk units to the ASP. If you want to make changes, press F12 to return.

> **Note:** Adding units can take from several minutes to several hours. The system updates the display periodically with the current status.

- You can press F16 to return to the Use Dedicated Service Tools (DST) menu if you have other tasks to perform. However, you cannot perform any disk configuration tasks or end DST until the system has finished adding disk units. Otherwise, wait until the adding of units is complete and press F12 until you reach the Work with Disk Configuration display.

Now you have all the existing LUNs and/or internal drives configured along side the new Fibre Channel attached LUNs. Next we remove all the unwanted SCSI LUNs and/or internal drives. Consider the following:

- The remove process must be performed from DST.
- The process of removing the units will cause the data to be moved to the remaining drives in the ASP. Sufficient space must exist on the remaining drives for the data to reside.
- You cannot remove the load source.
- You cannot remove disk units from a user ASP that is overflowed.
- When mirrored protection is active for the ASP that contains the disk units, you must remove both units of a mirrored pair.
- If you are going to physically remove a disk unit from a 6502, 6512, 6751, or 6532 IOP, you must either exclude the disk unit or stop device parity protection first.
- If you are going to physically remove a disk unit from a “real” 9337 Disk Array Subsystem, you must either exclude the disk unit or stop device parity protection.

To remove DASD units, perform the following:

- If you are not already using DST, perform a manual IPL to start DST.
- From the Use Dedicated Service Tools (DST) menu, do the following:
  – Select option 4 - Work with disk units
  – Select option 1 - Work with Disk Configuration
  – Select option 3 - Work with ASP Configuration
  – Select option 7 - Remove units from configuration
- You are shown the Remove Units from Configuration display.
- Type a 4 (Remove unit from configuration) in the OPT column for each unit you want to remove and press Enter. If the remove operation would leave you with insufficient space you will receive an error message. Otherwise, the Confirm Remove Disk Units display, may be shown if the storage management directories are not usable.
- Determine whether you want to cancel the procedure or continue. If you want to continue, press Enter. To proceed the system must perform internal processing that may take several minutes during which the system may appear inactive.
- When the remove is complete the display is returned to Work with ASP Configuration display. Press F12 until you reach the DST main menu.

You have now completed the data migration from the old drives or LUNs to the new Fibre Channel attached LUNs. If you want to set up remote/load source mirror from the internal load/source to the ESS, follow the instructions in Chapter 7, “Mirroring to an ESS” on page 165.

Part 4

FC attached tape

Until now iSeries tape has been limited to SCSI attach. With OS/400 V5R1 you can attach suitable tapes via FC to an iSeries. This is a great improvement to both a distance connection and the sharing of tapes.

10

Using Fibre Channel attached tape devices

With the advent of Fibre Channel support on the iSeries, the possibilities for tape attachment to iSeries have increased significantly. This chapter covers the new facilities available for attaching Fibre Channel and how to implement Fibre Channel attached tape devices.

10.1 Overview

Prior to V5R1, external tape was only supported on a SCSI connection. This had a number of limitations, the most significant being a distance limitation of 25 meters on all models of system and the maximum of two hosts systems attached to each 3590 drive and one host per 3580 drive. On AS/400 models 7xx and earlier, there was also a significant performance limitation due to the BUS speed and IOP/IOA capabilities, the result being that high speed tape drives such as 3580 and 3590 could not perform at speeds approaching their maximum rates.

With the introduction of the iSeries models with High Speed Link (HSL) and new I/O adapters, these performance limitations were effectively removed but the 25 meter distance limitation still existed. Now, with a Fibre Channel adapter available on iSeries at V5R1, the distance for attaching tape drives has been increased to 500 meters when directly attached or to 10Km when using a pair of cascaded hubs.

For many larger iSeries customers, this will be a significant enhancement and may in itself be a reason for moving to iSeries from the older models of AS/400.

10.2 Hardware supported

There are two types of hardware required to support Fibre Channel tape - the host I/O adapter and the tape device. Optionally, hubs or switches can also be used to extend the distance or connectivity.

10.2.1 Host I/O adapter

The I/O adapter for Fibre Channel tape support is feature #2765, which attaches to the feature #2783 I/O processor. Although feature #2783 can support up to four I/O adapters, only two of them can be feature #2765 for bandwidth considerations. Feature 2765 requires a minimum of OS/400 V5R1 and a minimum of 270 or 8xx hardware

AtV5R1, there are additional restrictions on the use of feature #2765.

- Feature #2765 can support a single *target* per adapter. On 3590 and 3584, the drive is a target. On 3583, a integrated SAN Data Gateway module is a target
- The iSeries must only see the drives across a single path

When using a hub or switch, zoning can be used if applicable to satisfy these restrictions. Additional information regarding these restrictions can be found in the section entitled 10.5, “Planning for Fibre Channel tape” on page 254.

You should check the iSeries System Builder for a complete set of I/O card placement rules.

10.2.2 Tape devices

Currently, there are three drive types that support Fibre Channel ports. These are the IBM Magstar® 3590-Exx drives and the IBM Linear Tape Open (LTO) drives in 3583 and 3584 libraries. New models of these drives/libraries are available with Fibre Channel ports.

Existing 3590 Exx drives can be upgraded from SCSI to Fibre Channel, thus offering 2 Fibre Channel ports rather than 2 SCSI ports. Bxx drives must be upgraded to Exx drives in order to gain fibre support. As part of the B to E upgrade, customers can choose the SCSI or fibre connection. Hence it is in the customer’s best interests financially to upgrade directly from Bxx to Exx fibre, rather than doing a two-step upgrade from 3590-Bxx to 3590-Exx SCSI and then to 3590-Exx fibre, if possible.

The 3583 Ultrium Scalable Tape Library offers fibre attachment via an integrated SAN Data Gateway, feature #8005, that offers 2 fibre ports. The SAN Data Gateway then connects the SCSI drives to the fibre ports. Whereas the SCSI version of this library was limited to a single host and single drive when attached to the AS/400 or iSeries, the fibre version of this library allows 2 hosts directly attached, or multiple hosts when attached via a hub or switch. With appropriate configuration, it is possible for these hosts to access all 6 drives in the 3583. Customers who purchased the original SCSI 3583 library can upgrade to the fibre library by purchasing the SAN Data Gateway feature. Since the original 3583 required HVD drives on the iSeries, and the fibre 3583 requires LVD drives, customers can optionally convert their drives from HVD to LVD, but only at the time that they purchase the SAN Data Gateway module.

For customers with 3584 libraries, new fibre drives can be purchased and added to the library. There is no conversion path for existing SCSI LTO drives in a 3584.

The IBM 3490, 3570, 3575, 8mm and QIC drives do not support Fibre Channel. For those customers who want to use the increased flexibility of Fibre Channel with the external devices, it may be possible to use third party “gateway” devices to convert the fiber cable to SCSI. However, it is not within the scope of this book to cover these devices. You should consult the suppliers of these third party devices for further information. The IBM SAN Data Gateway is not supported on the iSeries, except for the integrated SAN Data Gateway in the 3583.

10.2.3 Hubs and switches

Hubs and Switches can be used to increase the distance and connectivity of Fibre Channel tape drives. See Chapter 3, “Storage Area Networks (SAN)” on page 35 for switches supported at time of writing. See the iSeries Storage Web page for the latest information on supported hubs and switches:

http://www-1.ibm.com/servers/eserver/iseries/storage/sansupport.html

The IBM 3534 SAN Fibre Channel Managed Hub offers 8 ports that can connect hosts and devices in an FC-AL configuration. It runs QuickLoop by default.

The IBM 2109 SAN Fibre Channel Switch has two models offering 8 or 16 ports that can connect hosts and devices. At V5R1, the iSeries requires that this switch have the Quickloop software PRPQ #8S0521 licensed so the switch can be used in FC-AL. Customers may choose to purchase a switch today with Quickloop, rather than a hub (which is less expensive), to position themselves for fabric-mode in the future once it is supported on the iSeries.

At V5R2 and V5R3 the supported SAN devices increased dramatically along with switched fabric support. See the previously mentioned Web page for the latest support.

At V5R1 on the iSeries, two hubs or switches can be cascaded to offer a distance of 10km between, thus allowing 11 km total between host and device. An individual Quickloop can include at most two hubs or switches. However, in complex, multi-platform, enterprise wide SANs it may be possible for some ports on these devices to be used to connect them into a larger SAN with more cascaded hubs and switches. However, this has not been tested on the iSeries.

10.3 Benefits of Fibre Channel attached tape

Fibre Channel is the newest generation of tape attachment technology. It brings several key benefits, including device sharing among multiple hosts, extended distances, and the potential for increased performance.

10.3.1 Sharing tape devices among multiple hosts

SCSI attached tape drives have a set number of ports that allow attachment to host CPUs. For example, the SCSI LTO drives have a single port that allows attachment to one CPU. The 3590 has two ports that allow attachment to two CPUs. In years gone by, customers have sometimes had to buy more tape drives than they may otherwise have needed in order to satisfy their tape attachment requirements. For example, a customer with 6 small systems had to buy 6 SCSI LTO drives, even though each drive may only have been used for a few hours each night.

With the introduction of Fibre Channel support on the iSeries, it is now possible to attach many more hosts to a single drive via hubs and switches, and in the case of the 3583, via an integrated SAN Data Gateway.

Consider the example in Figure 10-1. Here we have three iSeries hosts connected via a 2109 switch to one 3590 tape drive. At V5R1, the 2109 must be running the Quickloop PRPQ to behave like a hub. Alternatively, a 3534 hub could be used. With Fibre Channel, all three hosts can have access to the drive, although not all at the same time. Standard “vary on/vary off” functions control access to the drive. If the 3590 is used in sequential mode, then the “vary on/vary off” can be done manually, or via a media management solution such as IBM’s Backup Recovery and Media Services/400 (BRMS) that manages the device allocation automatically. If the drive is in random mode, then the library can be varied on to all hosts concurrently, with OS/400 managing the allocation of the drive resource among hosts so only one host accesses the drive at a time.

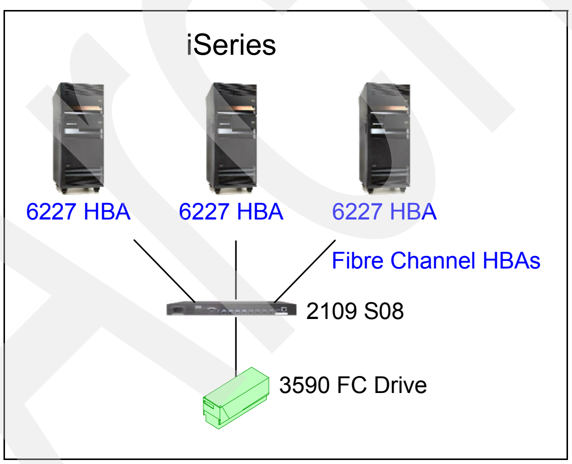

Figure 10-1 Sharing a tape drive among multiple host systems

Figure 10-2 extends this configuration to include a 3494 Automated Tape Library which can load tapes into any of the drives contained within the library. Such a solution will give the greatest flexibility. In the example shown, both 3590 drives could be available to any of the iSeries hosts.

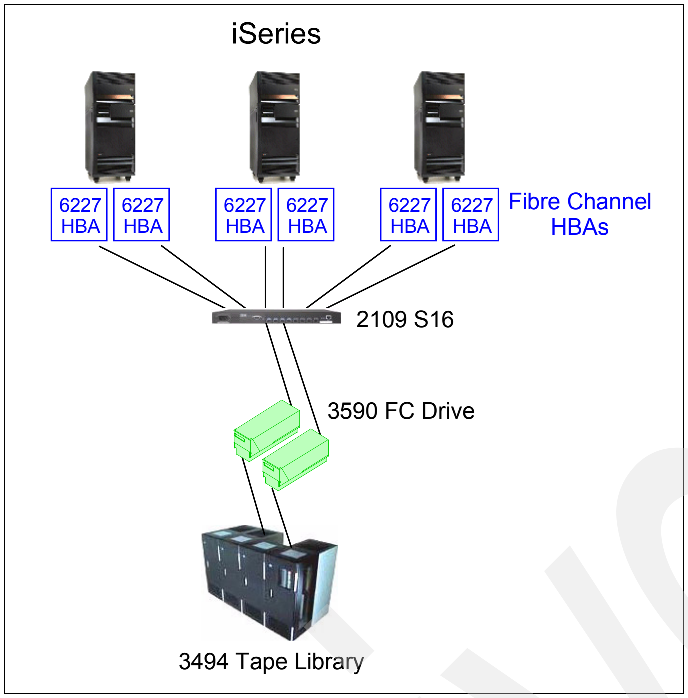

Figure 10-2 Extended flexibility with an IBM Automated Tape Library

At V5R1, each #2765 Fibre Channel adapter only supports one target, so you need to set up different zones on the switch/hub for each 3590. Media management software such as BRMS should be used to manage the library functions such as selecting volumes, and ejecting volumes for off-site rotation.

A common question is “How will the drives show up on my iSeries?”. To answer this question, we need to understand that Fibre Channel is merely a new connection technology. Each iSeries that can see the drive will auto-configure device and media library descriptions just as they did when attached by SCSI. In a future release, OS/400 will support fabric mode whereby extensive SAN configurations can be set up throughout the enterprise. At this time, the iSeries could theoretically see every drive in your enterprise and will auto-configure descriptions for any device that is supported on the iSeries. In practice, most organizations will use zoning to restrict the drives that the iSeries actually sees.

It is also important to note that although SAN technology makes it possible for multiple systems to connect to a given drive, appropriate software and operating system functions will be required to control that sharing, both in terms of use of the drive and security for the tape cartridges. In a homogeneous iSeries environment, OS/400 and BRMS can control that sharing. Once heterogeneous environments become prominent, other controls will be required. For example, in terms of drive sharing, it is important that all applications use the “reserve /release” functions so other systems can’t steal the drive away while a given system is using it. If the default parameters are used in OS/400, then OS/400 does the reserve/release function automatically when drives are varied on and varied off, and this function then carries over to BRMS. Tivoli Storage Manager (TSM), a backup product that is popular in the Unix and Windows world, also uses reserve/release. For other software packages, you should confirm with the vendor that reserve/release is used, otherwise the package does not protect itself in a shared environment.

Another common question is “Now that I have access to lots of drives, will my save automatically move to another drive if the drive it is running on fails?”. The answer to this question is “no”. When OS/400 detects a drive failure, it ends the save job. Human intervention is then needed to decide the best way to continue the operation. Although some organizations may want the save to pick up and continue on another drive, other organizations may not, since that other drive may be required for a more critical save elsewhere. Some organizations may choose to restart the backup from the beginning, while others may want to pick up where the save left off, or possibly even skip the save that evening due to time constraints. Human intervention or extensive automated management software is needed to handle situations like this. SAN should be treated as an enabler for this type of systems management, rather than the systems manager itself.

10.3.2 Extended distance

One of the major drawbacks of SCSI tape drives was the attachment limit of 25 meters. Although this could be extended with the use of 3rd party SCSI-Fibre Channel converters, such a solution proved too expensive or impractical for many customers who wanted a simple solution for locating their tape drives further away from the host system - often only in a different room.

Fibre Channel provides this capability. The following figures will illustrate some of the capabilities of Fibre Channel tape:

Figure 10-3 shows a direct connection. This simple solution uses short wave multi-mode fiber. It allows an LTO or Magstar 3590 tape drive to be located up to 500 meters from the single iSeries with 50 micron fiber or 175 meters with 62.5 micron fiber.

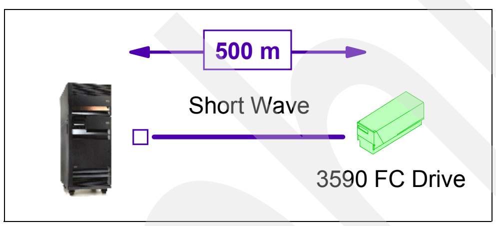

Figure 10-3 Fibre Channel direct connection

Figure 10-4shows the tape drive connected via a hub. The fiber on each side of the hub can be up to 500 meters, although in practice, the hub may be located closer to either the host or the tape drive.

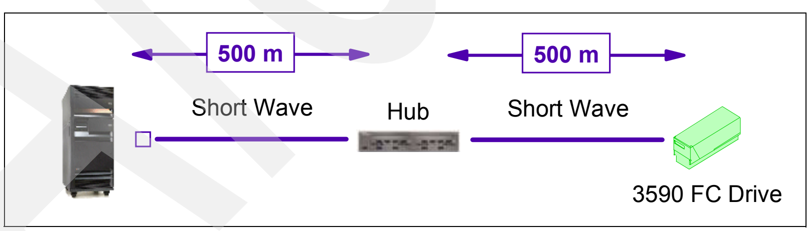

Figure 10-4 Fibre Channel hub connection over shortwave fibre

Figure 10-5 shows two cascaded hubs connecting the host and the drive over a longer distance. Shortwave fibre is used between the host and the hub, followed by long-wave fibre between the hubs, and shortwave fibre from the second hub to the tape drive. With such a topology, the tape drive can be located over 10 km from the host, allowing backups to be done “remote” from the central site, effectively providing an immediate off-site copy of the backup data. This type of solution is most likely to be used with tape libraries where minimal or even no operator intervention is required to mount tapes.

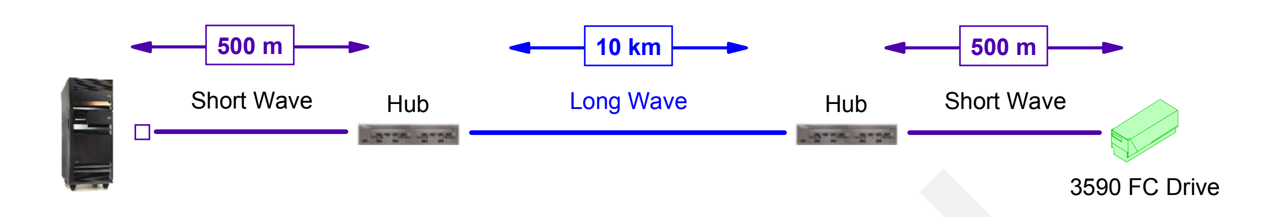

Figure 10-5 Fibre Channel cascaded hub connection over long wave fibre

10.3.3 Performance

The Fibre Channel technology provides the potential for significantly faster data rates than earlier tape connection technologies. However, the speed of the tape drives and CPUs available today is considerably less than what the Fibre Channel can carry, and as a result, Fibre Channel tape drives are only slightly faster than their most recent SCSI counterparts, when attached to the latest iSeries hardware.

Generally, customers will achieve the greatest performance gain when moving from older tape drives (e.g. 3590-B) to the 3590E and LTO technologies (SCSI or fibre), and by moving from older CPU hardware to the iSeries (i.e. 270 and 8xx and beyond). Often when the customer moves from the older drives to the newer drives, it makes sense to move to fibre technology at the same time, but this will depend on the customer's ability to upgrade to V5R1 which is a pre-requisite for fibre. Depending on the customer's upgrade time lines and specific performance requirements, he will need to evaluate the most appropriate timing for upgrade to Fibre Channel technology.

For customers who want to research this further, more detailed performance information is included in 10.5.2, "SAN tape performance considerations" on page 259.

10.4 Positioning

Fibre Channel tape technology will be attractive to customers who are looking for greater connectivity and distance between their CPUs and tape drives. To a lesser extent at V5R1, fibre attached tape may be attractive to customers who are looking for increased performance.

We will now look at some specific scenarios:

For customers who are already on V5R1 and iSeries hardware (i.e. 270 or 8xx models) and require new tape hardware, it would make sense to purchase the fibre-capable drives. This will provide the most recent technology available with all its associated benefits, and save the cost of upgrading in the future.

For customers who are already on V5R1 and iSeries hardware (i.e. 270 or 8xx models) and have existing SCSI 3590-E or LTO tape hardware, upgrading to fibre tape technology may or may not be warranted in the near term. If increased connectivity or distance is required, then upgrading makes sense. If increased performance is the requirement, then a more detailed examination of the performance benchmarks will be in order, since the performance delta between SCSI-attached tape on feature #2749 and fibre is small. The more significant performance delta is achieved when upgrading from 3590B to 3590E hardware, and from the older SCSI attachment cards to the feature #2749 on the 270 and 8xx CPUs. That being said, in some customer environments, even a small performance gain is helpful, and in this case, fibre would be warranted. For detailed performance data, 10.5.2, "SAN tape performance considerations” on page 259 or read the save/restore chapter of the iSeries Performance Capabilities Reference which is available at the following Web site:

http://www-1.ibm.com/servers/eserver/iseries/perfmgmt/resource.htm

For customers who are on earlier versions of OS/400 and/or older CPUs and are looking at purchasing new tape hardware, a business decision needs to be made. They need to weigh the benefits of purchasing Fibre Channel hardware in the near term to enjoy the Fibre Channel benefits and avoid the additional costs of having to upgrade SCSI to fibre at a later date, against the costs of doing a CPU and operating system upgrade in the near term. In many cases, customers find that the significant tape performance improvements between 7xx and 270/8xx CPUs is a compelling reason to upgrade.

For customers who are on earlier versions of OS/400 and/or older CPUs and/or older tape technologies and who want faster tape performance or more function, serious consideration should be given to upgrading to V5R1 and iSeries hardware and possibly fibre technology tape.

10.5 Planning for Fibre Channel tape

Fibre Channel offers considerably more flexibility in tape configuration than has been available in the past. However, there are a number of rules that must be adhered to at V5R1, and also some performance considerations to bear in mind. As a result, careful planning is required to ensure you have designed a suitable Fibre Channel tape configuration.

In this section, we will outline the rules and performance considerations, then show a series of configurations for each type of fibre drive, some valid and some not, in order to illustrate the various rules and considerations.

A checklist for the rules/considerations are:

Configuration rules:

- Fibre Channel Arbitrated Loop (FC-AL) is required
- All Hosts on the FC-AL must be iSeries (no other platforms)
- Maximum of one target per FC-AL
- Maximum of one target per iSeries fibre adapter
- Maximum of one path visible to the iSeries for any given device
- Both the iSeries fibre adapter and the tape fibre port must be attached to shortwave fibre (with longwave between hubs if needed)
- Maximum of 2 cascaded hubs or switches (V5R1)
- Design configuration to minimize impact of Loop Initialization Primitive (LIPs)

Performance rules:

- Choice of Drive
  - Choose 3590 for average file sizes and start/stop operations. Choose LTO for large file streaming operations
- Bandwidth
  - Do not exceed the 100 MB/s bandwidth of the fibre connection
  - On the 3583, do not exceed the 80 MB/s bandwidth of the SCSI buses

10.5.1 Configuration rules

Fibre Channel Arbitrated Loop

In Fibre Channel, there are 3 different topologies as follow:

– Point to Point - a dedicated connection between the two nodes that provides the full bandwidth (100 MB/s) to the connection

– Fibre Channel Arbitrated Loop - a token-ring like topology where multiple devices share the ring and its bandwidth (100 MB/s), with only one conversation underway at a time

– Switched Fabric - an any-to-any topology implemented via a switch such that each connection has the full bandwidth (100 MB/s)

At V5R1, OS/400 only supports the FC-AL topology for tape, although a point-to-point connection can be implemented that runs FC-AL.

From V5R2 and later releases OS/400 and i5/OS support switched fabric.

Homogeneous Fibre Channel Arbitrated Loop

At V5R1, iSeries hosts must be in a *homogeneous* FC-AL, without any other server platforms mixed in. This is shown in Figure 10-6. If other platforms needed to share the same hub or switch, they could be zoned into their own homogeneous FC-AL's, possibly sharing the same tape device.

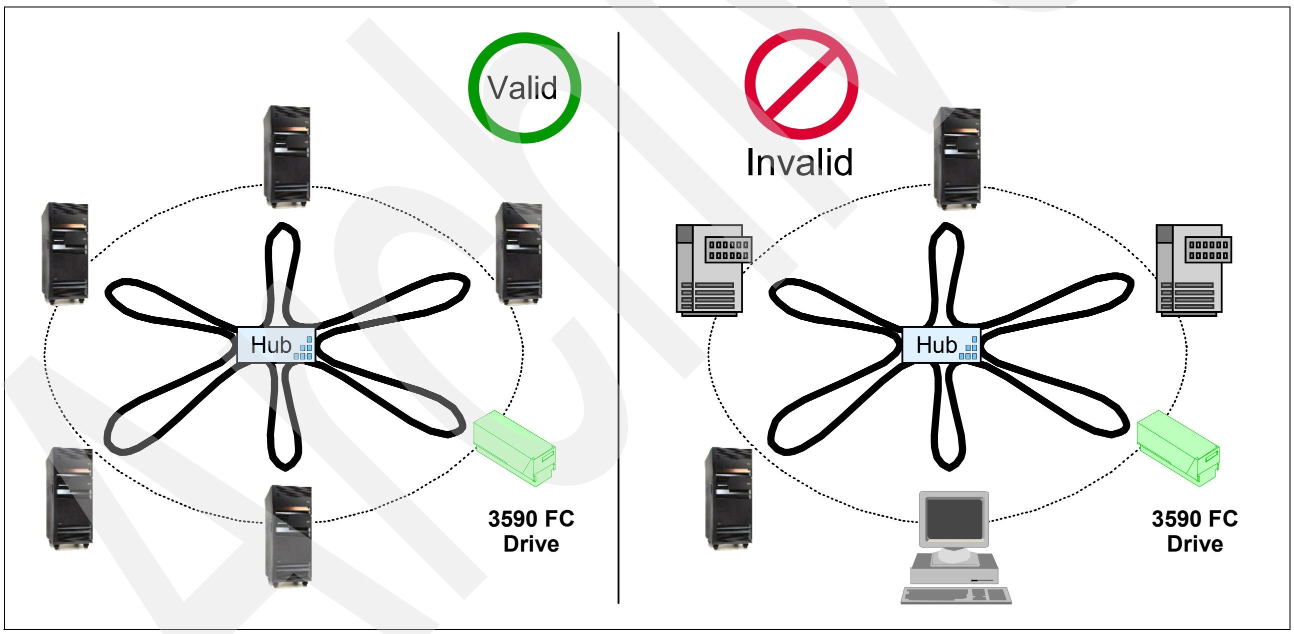

Figure 10-6 Homogeneous Hosts in FC-AL

One Target per Fibre Channel arbitrated loop

In Fibre Channel, we have initiators and targets on the FC-AL. Initiators are the iSeries hosts. For 3590 and 3584, the targets are the tape drives. For the 3583, each fibre port on the integrated SAN Data Gateway is a target.

At V5R1, there is a maximum of one target on each FC-AL. The reason for this is that the iSeries device addressing scheme does not let it differentiate more than one target on the FC-AL.

As an example, as shown in Figure 10-7, it would be acceptable to put five iSeries hosts and one 3590 drive on an FC-AL, but it would not be acceptable to add a second drive to the same FC-AL since OS/400 would not be able to distinguish between the two drives.

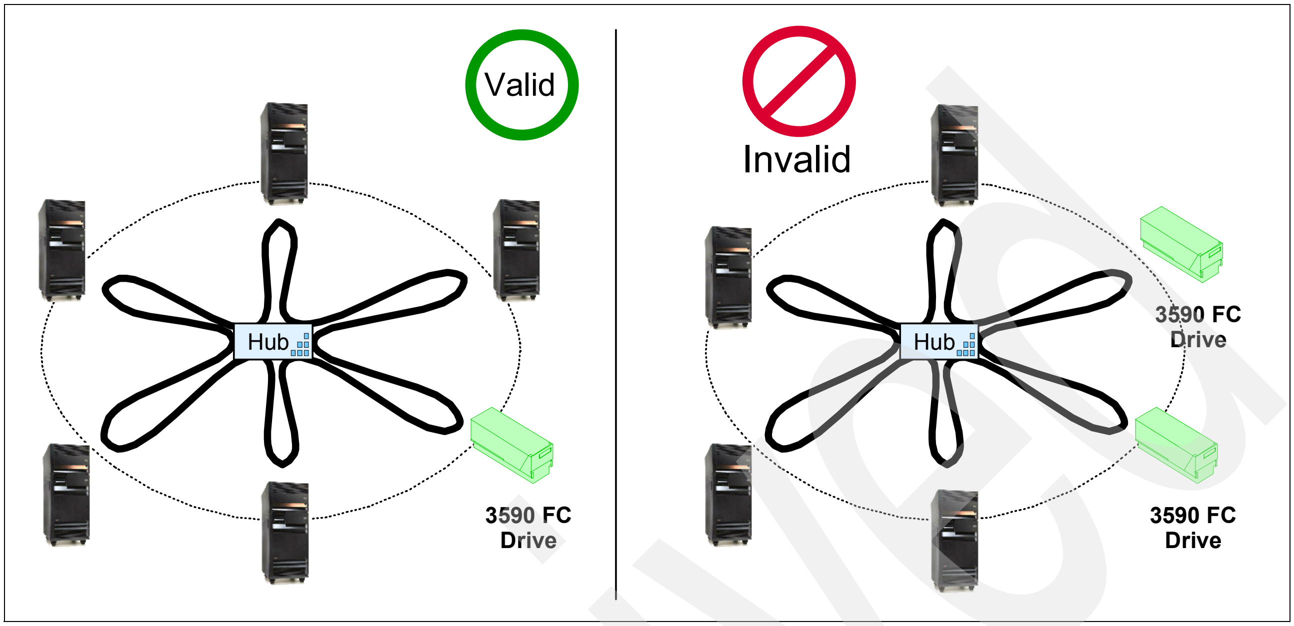

Figure 10-7 One target per Fibre Channel arbitrated loop

Here is an example of where it could be very serious to have more than one target in the FC-AL. Suppose, that you put both a disk adapter and a tape adapter in the same FC-AL. When you IPLed the system, you could lose contact with your disks if the iSeries picked up the tape adapter as the single target rather than the disk adapter

> **Important:** Do not mix tape and disk adapters in one FC-AL. Failure to do this will cause significant operational problems. For example, when IPLing an iSeries system when there are both tape and DASD devices in the same zone can cause the iSeries to lose contact with the LUNs in the Enterprise Storage System.

Maximum of one target per iSeries adapter

At V5R1, each iSeries adapter can only attach to one Fibre Channel target. The reason for this is that the iSeries device addressing scheme does not let it differentiate more than one target on the FC-AL

For example, as shown in Figure 10-8, if there were two 3590 tape drives, and 3 hosts, each with one FC #2765 adapter, it would not be acceptable to make two overlapping FC-AL's, each with all the hosts and one of the tape drives, since each adapter would have more than one target.

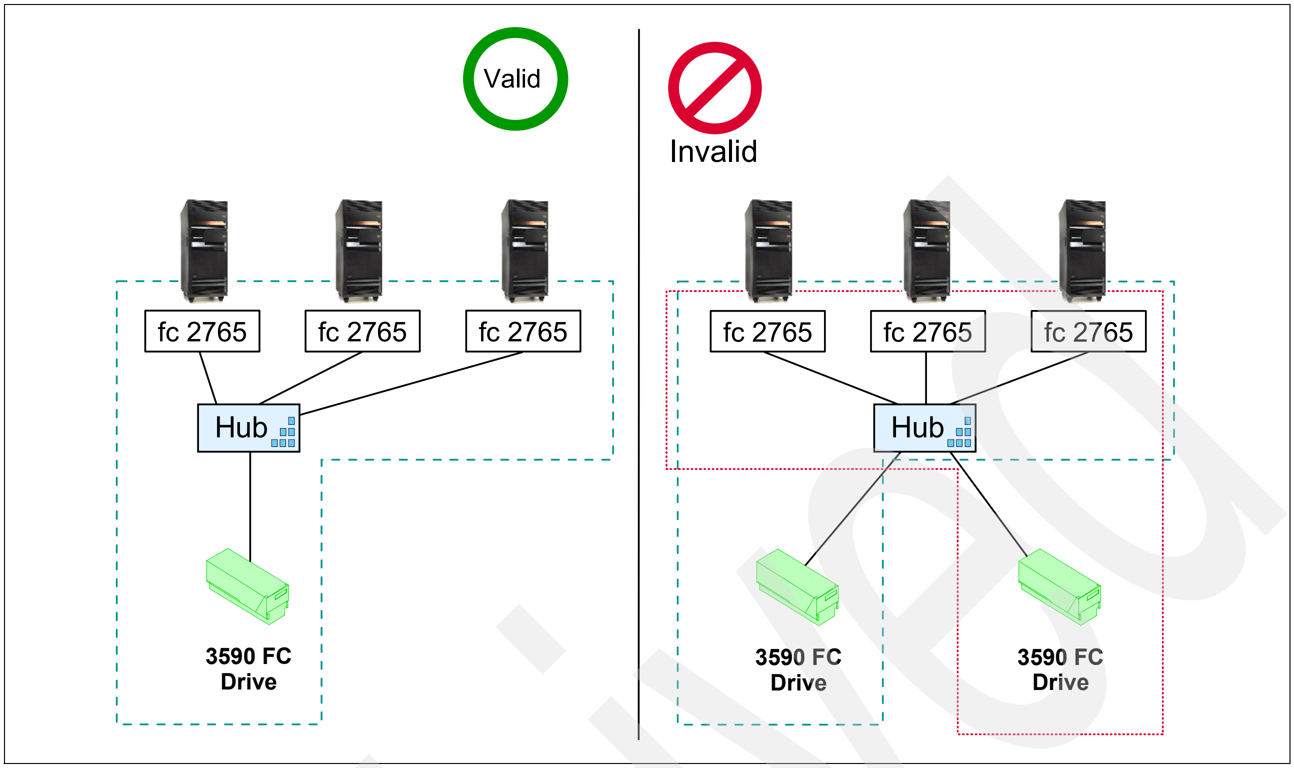

Figure 10-8 Maximum of one target per iSeries adapter

Maximum of one path visible to the iSeries for any given device

At V5R1, OS/400 does not support multiple paths to the same device. When the system powers on, it walks the bus to determine the devices that are available based on the device serial number. Suppose there is more than one path to the same device: OS/400 will see the device on the first path and assign resource names. However, shortly thereafter, it will see the same device on a second path and assume that the device has been moved, so the first set of resources will be marked as not-operational. As a result, valid configurations at V5R1 should only have one path from an iSeries to any device. This is shown in figure Figure 10-9.

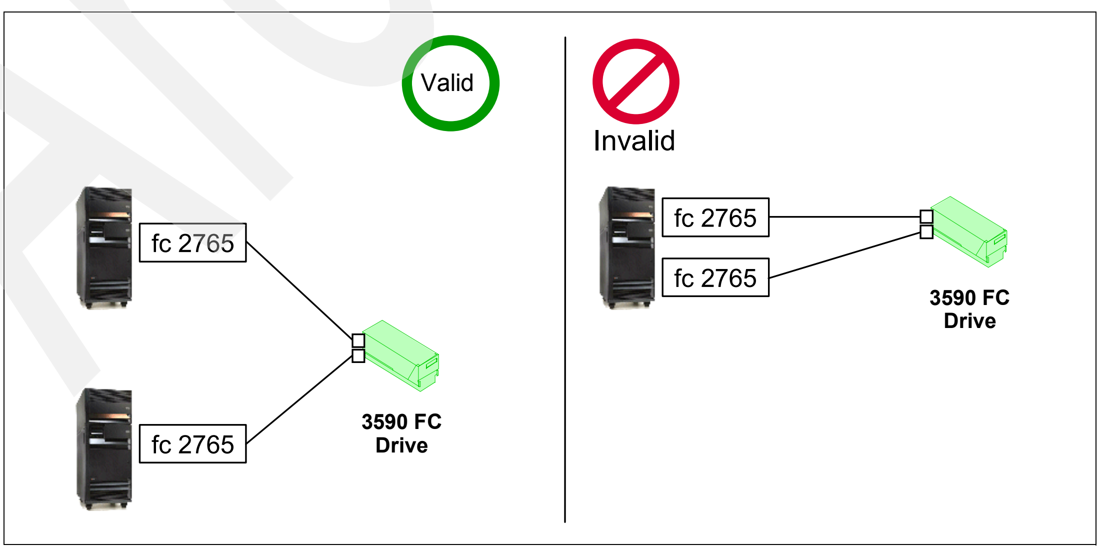

Figure 10-9 Maximum of one path visible to the iSeries from any device

Shortwave fibre to target/initiator

The fibre ports on the iSeries host adapter cards and the tape drives can only accept shortwave fibre cables. This is usual in the fibre tape marketplace. As a result, if distance is required, then hubs must be used to link the shortwave fibre cable to the long-wave fibre cable. This is shown in Figure 10-10.

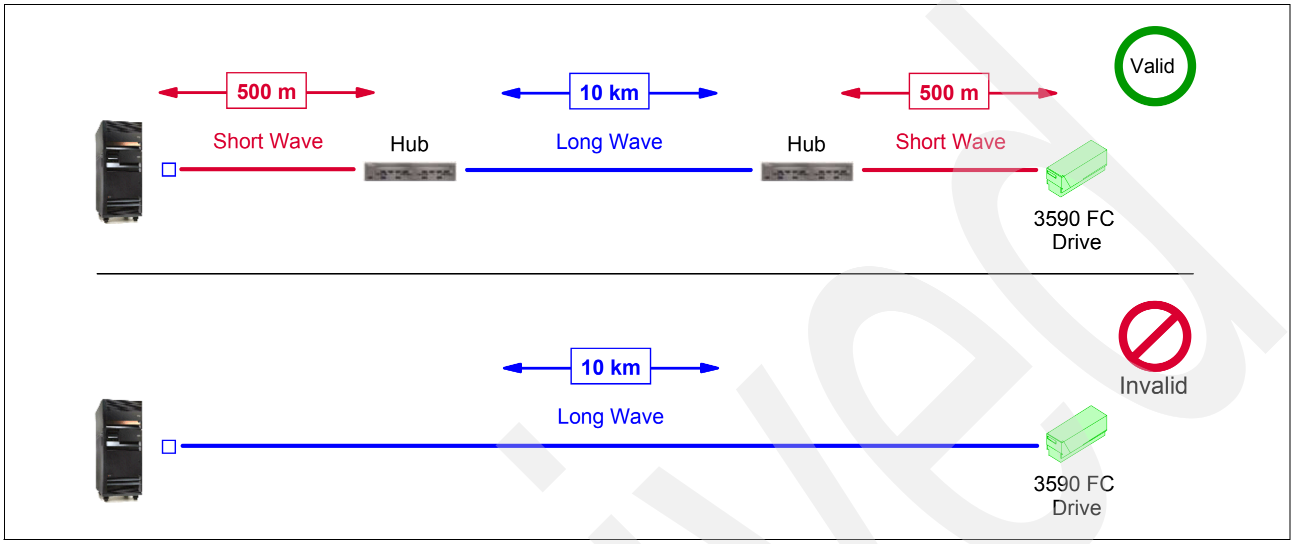

Figure 10-10 Shortwave fibre to target/initiator

Maximum of 2 cascaded hubs or switches

AtV5R1, iSeries SAN tape attachment is via FC-AL. The 3534 managed hub runs Quickloop by default to offer FC-AL, and the 2109 switch needs to have the Quickloop PRPQ so the iSeries ports can behave like FC-AL instead of fabric ports. Quickloop allows at most 2 hubs or switches to be joined in a partner group, and as a result, V5R1 SAN configuration should be kept to at most 2 hubs/switches or one hub and one switch in each loop.

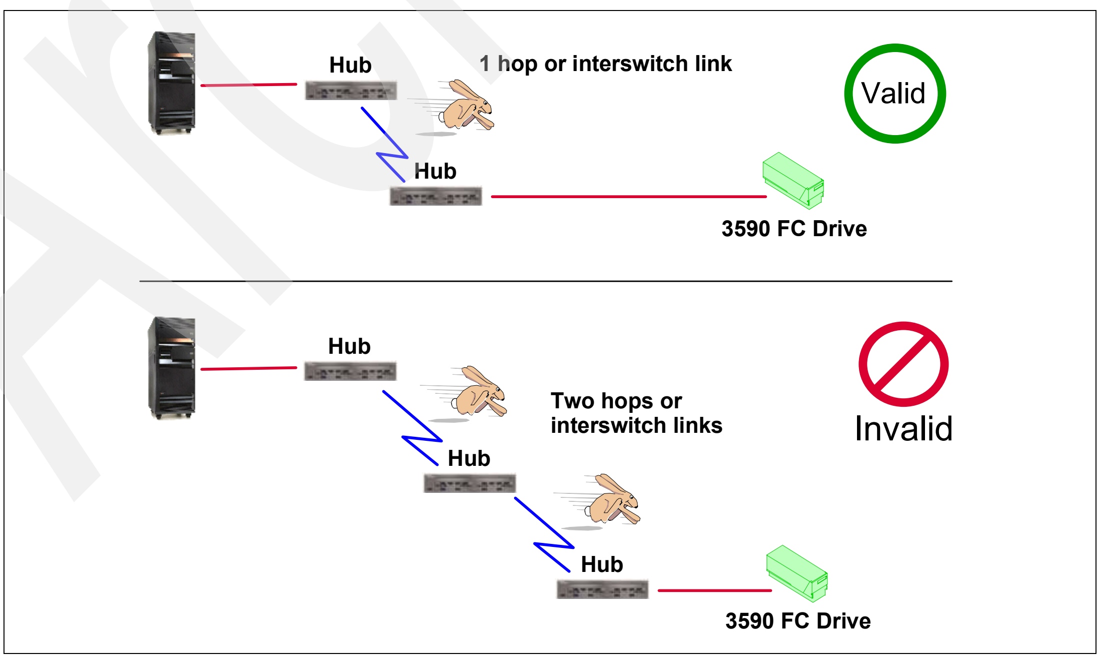

Figure 10-11 Maximum of two cascaded hubs

Minimize impact of Loop Initialization Primitive (LIPs)

When a new device is added to a loop, the loop needs to re-initialize and ensure that all loop members have a loop-address. This would occur when a net new device is added to the loop, or an existing device is hot-plugged from one switch port to another, or an existing device powers on, etc. The entire LIP may only take a few milliseconds, but it is possible that devices may change their loop addresses as part of the process.

On other platforms, if there is a tape operation going on when a LIP occurs, it is possible for the save to fail and need to be restarted. As a result, the size and arrangement of the FC-AL loops will be designed to minimize the impact of a LIP on these platforms.

On the iSeries, error recovery is built into the operating system, so in most cases, a tape operation that is underway at the time of the LIP will pause briefly, then continue. However, in mission critical environments, customers may want to take the LIP into account, and design a configuration that will minimize the impact. For example, the number of hosts in each FC-AL loop may be kept to a minimum using zoning so there are fewer other systems to potentially IPL, and hosts with similar IPL schedules may be selected to share the same loop. Overlapping loops might be avoided since a LIP will propagate from one loop to another via the nodes or devices that are in both loops.

10.5.2 SAN tape performance considerations

Although Fibre Channel attachment for tape opens up new opportunities and is the best performing option for iSeries save/restore operations, there are still some performance considerations to be taken into account.

Firstly, although Fibre Channel provides the potential for very high save/restore speeds, today's tape drives and CPUs are not fast enough to take advantage of this entire bandwidth. As a result, fibre tape drives perform only slightly faster than their most recent SCSI counterparts when running on the latest iSeries hardware. By understanding which upgrades have the greatest impact on save/restore speeds, you can choose the most appropriate upgrade path and timing to accomplish your goals.

Secondly, the 3590 and LTO tape drives have different performance characteristics. Although both drives are able to reach comparable speeds, they excel in different areas. It is important to choose the drive that fits with your specific save/restore requirements.

Thirdly, although Fibre Channel provides high bandwidth connections, it is still possible to exceed the bandwidth available, when multiple systems and devices are sharing fibre media. By careful design of your Fibre Channel configuration you can maximize the tape drive performance.

This section will explain the considerations in each of these areas.

Choice of SCSI or fibre

Several recent enhancements in iSeries hardware have offered significant improvements in tape performance. Although Fibre Channel tape drives provide the optimum performance, there are SCSI configurations that provide performance that is almost as good. By understanding the various upgrades and their contribution to tape performance, customers can select the most appropriate tape upgrade strategy to meet their needs.

In this section, we will review the various components of tape performance, and explain the performance benefit that each provides.

Figure 10-12, "Tape Drive Performance Path" on page 260 shows the path used when data moves between the iSeries and a tape drive. For a save, the data moves from right to left across the top of the chart, starting in the CPU, then moving out through the IOP/IOA to the bus, and on to the drive. Below the pictures are performance figures indicating the speeds of the various components that could be used for a certain link in the path. Many people hear that a tape drive can run at a certain raw speed, and assume that is the speed that they will see on their system. However, by looking at the chart, you can see that there are many factors involved in the process of saving data on an AS/400 or iSeries. The chart will help you to identify the slowest link in the tape path, and thus determine the upgrade that would be required to improve your backup speed. Notice that the right-hand section of the chart is divided into the older processors at the top, and the newer processors at the bottom. The reason for this split is that the 270 and 8xx processors introduced a new bus technology that has made a tremendous difference in save/restore rates.

Drive | Bus | IOP/IOA | CPU

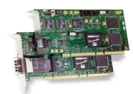

LTO
SCSI: 31 MB/s
Fibre: 46 MB/s

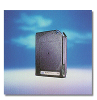

3590
B-SCSI: 14 MB/s
B-Ultra: 27 MB/s
E-SCSI: 34 MB/s
E-fibre: 43 MB/s

* Drive speeds shown assume 3:1 compression

7xx and below

25 MB/s

fc 6501: 17 MB/s
fc 6534: 17 MB/s
fc 2729: 13 MB/s

Save Rates	Source File	User Mix	Large File	X-Large File
LTO SCSI	4	12	17	17
LTO Fibre	n/a	n/a	n/a	n/a
3590B SCSI	2	7	14	14
3590B ultra fibre	2	9	17	17
3590E SCSI	2	10	17	17
3590E fibre	n/a	n/a	n/a	n/a

8xx and above

700 MB/s

fc 2749: 38 MB/s
fc 2765: 100 MB/s

Save Rates	Source File	User Mix	Large File	X-Large File
LTO SCSI	4	19	26	30
LTO Fibre	5	22	34	40
3590B SCSI	3	9	14	14
3590B ultra fibre	3	20	24	28
3590E SCSI	4	21	31	32
3590E fibre	5	22	37	39

* Save speeds shown at right were achieved in benchmarks and take into account the processor power and also the other links in the drive/bus/iop chain. Save rates are for source file, user mix and large file data respectively. For details, see the iSeries Performance Capabilities Reference at http://publib.boulder.ibm.com:80/cgi-bin/bookmgr/books/as4ppcpf

Figure 10-12 Tape Drive Performance Path

Starting at the right of the chart, we see that the processor gathers up the files to be saved. Depending on the file size, this can be very processor intensive, and can greatly impact the speed of the save. As a result, tape performance benchmarks are generated for multiple different workload mixtures as follow:

- Source File: these are very small files that take a lot of processor resource to save. As a result, save speeds are slow since the processor is not able to send the data quickly enough to keep the tape drive busy. Examples of source files are the source programs on a development system.
- User Mix: this is a mixture of small, medium and large-sized files intended to represent a typical environment.
- Large File: these are large files such as database files. The large file benchmark uses 2 GB files and 4 GB files depending on the benchmark.

- Extra Large File: As tape drive speeds have increased, the 2 GB and 4 GB large file workloads were no longer sufficient to drive them to their maximum speeds. As a result, an extra-large file workload with files as big at 16 GB was introduced. It is unlikely that a real customer environment would have files this large.

Notice that on the 7xx processors, the save speeds are the same for large file and extra large file data. The reason for this is that the tape path is bottle necked on the I/O adapter, and hence sending larger (more efficient to save) objects, does not help to speed up the backup.

Note that in the examples shown, very large processors were used. Customers with small or medium-sized processors may find that their saves are restricted by the processor speed more so than shown in the chart. Also note that the figures shown are for saves. Depending on the tape drive technology used, restores can vary from comparable speeds to considerably slower speeds. Consult the iSeries Performance Capabilities Reference Guide to evaluate whether the restore performance of your chosen tape technology meets your business requirements. It can be found at the following URL:

http://publib.boulder.ibm.com:80/cgi-bin/bookmgr/books/as4ppcpf

From the CPU, the data is sent to the IOP. Notice that the IOPs are the slowest link in the tape chain, in all the large-file benchmarks, except the 3590B examples. Hence the reason that the faster #2749 SCSI IOP on the 8xx processors was such a help for tape performance. Notice that the fibre adapter is so fast that it is no longer the bottleneck in the tape performance path. This is the reason that upgrades to fibre adapters do not provide as big a performance boost as people may presume when they hear that the adapter runs at 100 MB/s: other factors are bottle necking the save, not the fibre adapter. However, the fibre infrastructure positions people well for future tape drives with higher performance, and larger CPUs that can send user-mix data faster.

After the IOP, the data crosses the system bus on its way to the drive. Notice the dramatic difference in speed between the buses on the 7xx processors and below (25 MB/s), compared with the 270 and 8xx processors that have the HSL bus technology (700 MB/s or more). On a 7xx processor or below, if you put two SCSI adapters on the same bus, you would saturate the entire bus bandwidth. As a result, customers who required the fastest tape performance possible would but a single tape IOP on each bus, and move their DASD to separate buses. In contrast, the HSL bus on the 270 and 8xx processors is so fast, that it is possible to put many tape adapters on the same bus before it is saturated. In a benchmark, the lab ran ten 3590E drives simultaneously and still had bandwidth to spare, a feat that was unheard of on the earlier bus technologies.

After crossing the bus, the data runs along the cable and finally reaches the tape drive. Notice that almost all the tape drives are considerably faster than the I/O adapters that are available on 7xx processors, whereas the I/O adapters on the 8xx processors are comparable to or faster than the tape drives. This is the reason that customers get such a tape performance boost when they upgrade to 8xx technology, because they are no longer constrained by the bus/adapter combination. Notice that in order to get the benefit, customers need to migrate to one of the new adapters: leaving the tape drives on older adapters on migrated buses, leaves them in the same tape performance situation as they were in on their older CPUs.

Now that we have reviewed the tape performance path, we will put it into context for customers who are planning the timing of their Fibre Channel migration.

Start by calculating the approximate tape performance that you are seeing today. To do this, calculate the number of GB of actual data on your system. As a rough approximation, you can use the statistics on the WRKSYSSTS display, remembering to subtract off any unused disk space. Then time a full save of your system. Divide the total GB by the number of hours for the save, then divide by 3.6 to get the approximate MB/s. Compare your approximate save speed to the appropriate table at the right-hand side of the chart, and determine whether your data mix tends towards source file, user mix, or large file, based on the speeds you are seeing.

Now look to the tape drive, bus and adapter speeds to determine whether any of them is the limiting factor in your save speed. If not, then your processor is likely gating your speed, and as a result, a processor upgrade would help you to increase your save speed. If one of the drive, bus, adapter is the limiting factor, then look at your possible upgrade options to see how you could remove this bottleneck, and what kind of a performance improvement you could then expect. This will help you to contrast the benefits you will see by upgrading to Fibre Channel, compared with the benefits you will see by adjusting your SCSI configuration.

As an example, suppose a customer had a 3590E tape drive running on a 6501 IOP on a 7xx processor and wants to increase the speed of his backups. Suppose he calculates his approximate backup speed of his current backup as 17 MB/second. Looking at the chart for the 7xx processors, we see that his save speed is at the right-hand side of the chart, and hence the customer likely has a large-file workload, since the extra-large file workload is not normally found in real life environments. Looking at the drive/bus/adapter speeds, we see that the drive is capable of far more than 17 MB/s, but it appears to be bottle necked on the #6501 adapter. In order to speed up the backup, the customer needs to get a faster adapter. To do this, he needs to upgrade to an 8xx processor, so he can choose between the #2749 SCSI adapter and the #2765 fibre adapter. Next we look at the backup speed chart for the 8xx processors, to estimate the backup speed he would get with each of these adapters with his large-file data and a 3590E SCSI or 3590E fibre drive. We learn the following:

- 3590E SCSI: 31 MB/s (1.82 times faster than his 7xx backup)
- 3590E Fibre: 37 MB/s (2.17 times faster than his 7xx backup)

The 6 MB/s speed difference between these two drives amounts to 21 GB/hr. This means that a 500 GB backup that would have taken 8.5 hours on the old system, will take 4.5 hours on SCSI on the new system, and 3.75 hours on fibre on the new system.

Now the customer needs to decide his course of action. In order to speed up his backup, he will need to upgrade to an 8xx processor. If he is willing to upgrade to V5R1, rather than just V4R5, then fibre will be an option in addition to SCSI. If he chooses to stay on SCSI, then he only needs to buy a new #2749 adapter and he will trim his backup time almost in half. If he decides to go to fibre, he will need to purchase a new #2765 adapter instead of a #2749, and he will also need to upgrade his drive from SCSI to fibre, but he will trim another 45 minutes of this backup time, which may or may not be important in his environment.

Choice of drive - 3590E or LTO?

Although the 3590E and LTO drives have similar maximum speeds, their performance characteristics vary dramatically. In general, 3590 is an industrial strength drive that excels on small and user mix files and start/stop work such as random retrievals. In contrast, LTO was designed to provide excellence in technology at an economical price. LTO is very fast at streaming saves of large files, but is not as efficient at smaller files and start/stop work.

This section will explain the differences between these drives in more detail. Customers should understand these differences when deciding which drive technology to adopt.

Buffer flush during saves

When iSeries writes to tape, it uses IBM Standard Label formatted tapes. The objects from each library are written to a tape file, surrounded by control characters such as tape marks. Periodically throughout the save, and also at the end of each library, the system flushes the tape drive buffer so it knows that all the data to that point is safely on tape. The 3590 is able to do this buffer flush very quickly (approximately 0.4 seconds), whereas the LTO drive is considerably slower (approximately 2 seconds). For saves with a lot of small libraries, there are a large number of these buffer flushes, thus making LTO perform less favorably on source file saves. Customers who are saving large files may find that LTO is their best performer. However, if they will be saving small and average files, 3590 may be a better choice.

Note: The performance data for source file saves in Figure 10-12, “Tape Drive Performance Path” on page 260 seems to contradict this point. The reason is that there is considerable rounding error in calculating the performance benchmarks for these small file saves.

Search speed during restores

During a save, OS/400 offers API programs that can record the volume and file sequence of the libraries that are being saved to tape. Tape management systems such as IBM’s Backup, Recovery and Media Services product (BRMS) record these volume and file sequence numbers in a database to make it easier to restore the files later. Without BRMS, if you want to restore an individual object off the tape, you either need to supply the file sequence number yourself (which you probably will not know), or use the *SEARCH option to ask OS/400 to start reading at the start of the first tape in the set and check through every tape file header until it finds the object that you need somewhere in the tape set. This can be very slow.

By comparison, with a product like BRMS that can supply the volume and file sequence numbers, OS/400 mounts up the correct volume from the set, then passes the file sequence number to the tape drive. The tape drive checks the volume control header at the start of the tape, and spins directly to the proper zone of the tape. Once it is in the right neighborhood, the tape drive reads the tape file headers in that region one by one until it finds the required object.

On the 3590 tape drive, the zones that it can spin to quickly are very granular, thus meaning that very few tape file headers need to be read before the object is found. In comparison, the LTO drive has less granular zones, so more tape file headers need to be read. If a restore is jumping around the tape, picking up one tape file here, and another tape file there, the extra search time required by LTO can add up, thus making it a poor choice of drive for this type of environment.

Clearly, there is a trade-off between these two drives, especially when considering price/performance. If you want maximum performance and have large libraries and a few small libraries, 3580 is likely to be the best choice. However, if you have many small libraries, or you need to restore individual libraries or objects which are scattered throughout the tape, then 3590 would be a better choice. Apart from saving or restoring large libraries when the advantage of streaming comes into effect, 3590 still remains the best tape drive available to iSeries customers.

Bandwidth

Although Fibre Channel offers considerably more bandwidth than we have had on the iSeries in the past, it also allows more flexibility for connectivity. As a result, care needs to be taken to design configurations that have sufficient bandwidth.

Fibre Bandwidth

The #2765 I/O adapter and the Fibre Channel cable have a maximum throughput of 100 MB/s at the present time. This capacity is full duplex, which means that theoretically data could be travelling along the cable at 100 MB/s in each direction simultaneously. In the case of tape drives, that would mean there could be 100 MB/s of saves going on at the same time as 100 MB/s of restores. In practice, however, the cable probably does not have the capacity to run a full 200 MB/s, but simultaneous reads and writes would probably give more than 100 MB/s total. Recall that at V5R1, there can be at most one target per fibre adapter. This means that for the fibre 3590E and 3584 tape drives, the cable would only be used in one direction at a time, since a single tape drive cannot read and write simultaneously. In the case of the 3583, however, where all 6 drives report in via one target, it IS possible to have both reads and writes going on simultaneously, and hence it is possible to have more than 100 MB/s crossing the cable at a time.

The 3534 managed hubs and 2109 switches with Quickloop have a bandwidth of 100 MB/s, shared among all devices on the loop. Since both the 3534 and the 2109 with Quickloop are *intelligent*, then if they are zoned into multiple loops, each of the resulting loops has a bandwidth of 100 MB/second. As with direct connection, the FC-AL is also full duplex. This means that for a device like the 3583 where there are multiple devices associated with one target, the loop could potentially carry more than 100 MB/s if saves and restores are running simultaneously

In this chapter, when we indicate that 100 MB/s of bandwidth is available in each direction, that means that the rated speed is 100 MB/s. Readers should take into consideration that the actual bandwidth will be less than the rated speed.

From the chart in Figure 10-12 on page 260, we know that the fibre 3590E drives run at 22 MB/s for user mix files, and 37 MB/s for large files. Similarly, the fibre LTO drives run at 22 MB/s on user mix files, and 34 MB/s on large files.

Now let’s look at some examples:

In Figure 10-13, we direct-connect a fibre 3590E or 3584 drive. Since there is a limit of one target per #2765 I/O adapter, we know that there will be plenty of bandwidth available since the fibre cable can carry 100 MB/s each direction, but the drive will need at most 37 MB/s if it is a 3590E, and 34 MB/s if it is an LTO drive.

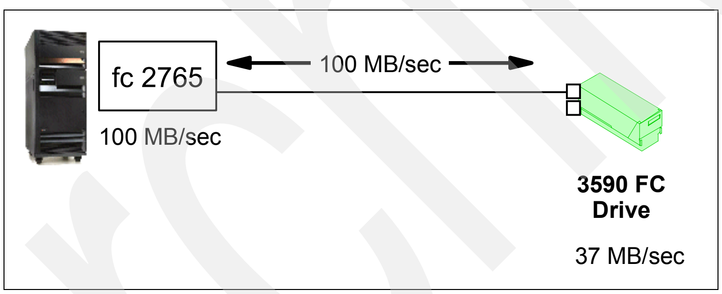

Figure 10-13 Fibre bandwidth: direct connected fibre 3590E or 3584 drive

In Figure 10-14, we direct-connected a 3583 LTO library with 6 drives. In this case, our fibre cable can carry 100 MB/s each direction. If all drives were running simultaneously doing a parallel save with large file data, they could require 6 times 34 MB/s (204 MB/s) which is more than the cable can carry. If half the drives were running a parallel save of user-mix files at 22 MB/s, and the other half of the drives were running restores, again at 11 MB/s, then the bandwidth would probably be adequate (66 MB/s for the save direction, and 33 MB/s for the restore direction).

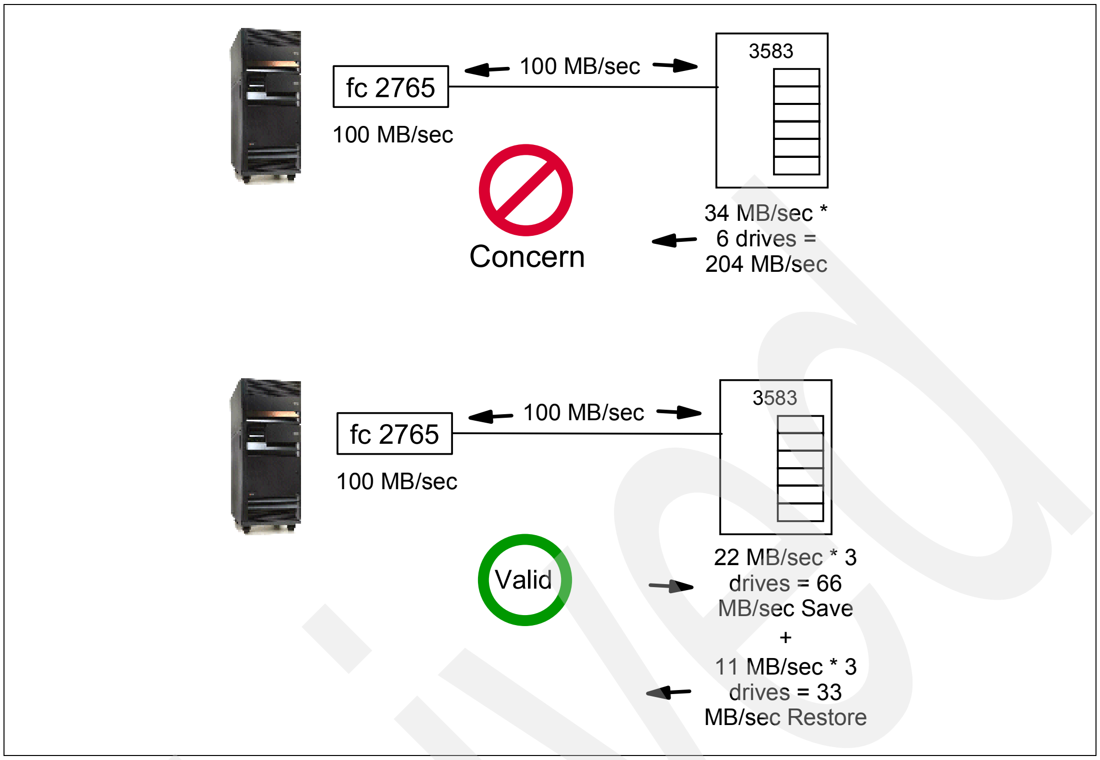

Figure 10-14 Fibre Bandwidth: direct connected 3583 with 6 drives

As a final example, consider Figure 10-15 that shows a 3583 attached to six iSeries servers via a 3534 hub. Suppose all six systems need to run their daily backup at the same time, and all have large file data (34 MB/s). The total bandwidth demand of the drives is 204 MB/s which will exceed the 100 MB/s that the FC-AL offers in one direction (i.e. saves). Zoning the hub into multiple smaller FC-ALs will not help because both zones will be using the same fibre cable between the hub and the 3583. However, this situation can be remedied by adding a second fibre cable from the hub to the 3583. Two zones are required so the iSeries hosts don't see the tape drives via more than one path. If half of the iSeries are in one zone along with one 3583 fibre port, and the rest of the iSeries are in the other zone with the other 3583 fibre port, then we have doubled the available bandwidth of this configuration. Now the available bandwidth will be just slightly beyond saturation with the nightly saves.

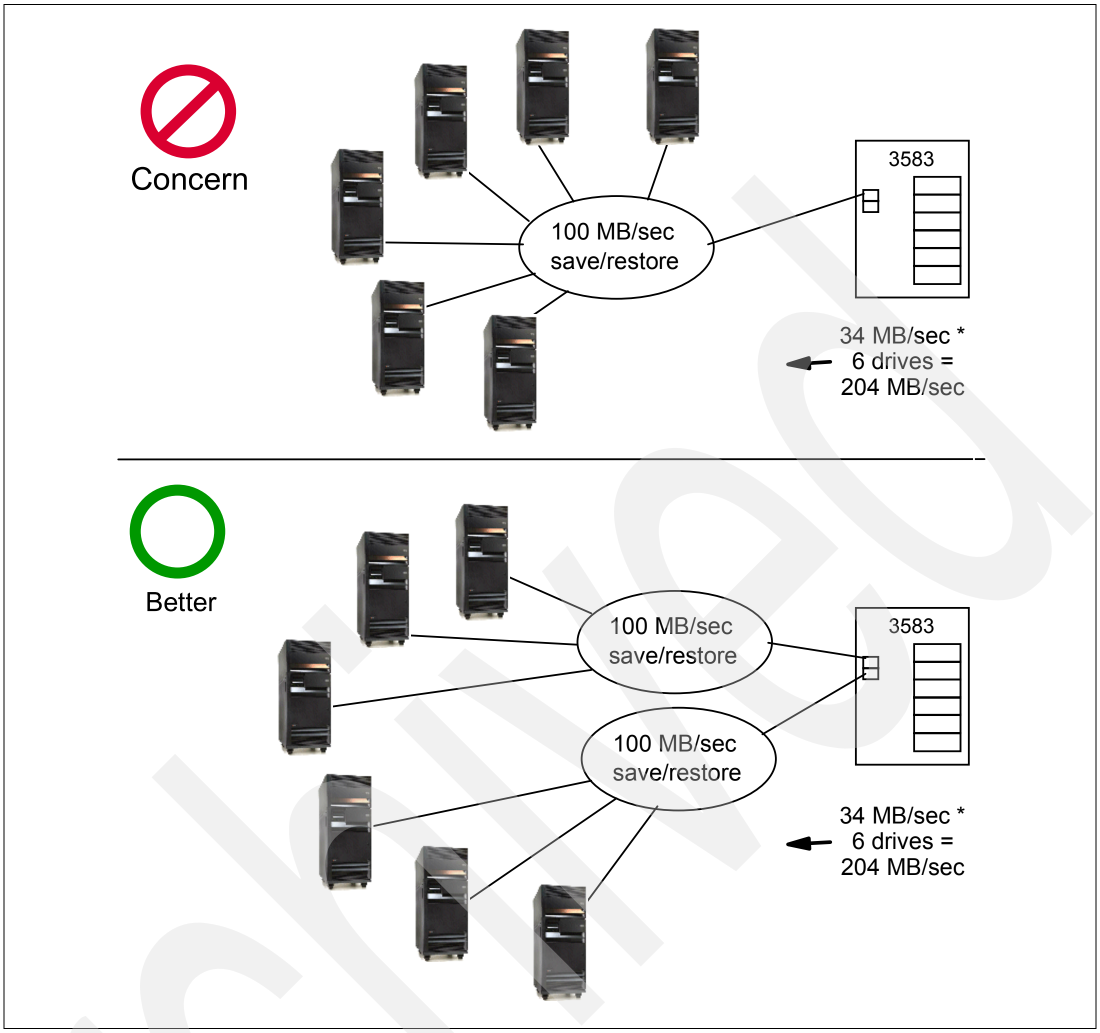

Figure 10-15 Fibre bandwidth: 3583 attached to six hosts via FC-AL

3583 SCSI bus cabling

The fibre 3583 combines an LVD SCSI tape library with an integrated SAN Data Gateway to connect the SCSI drives to two fibre ports. Careful choice of cabling between the drives and the SAN Data Gateway will maximize performance of the drives. This section will describe the 3583 configuration and cabling practices that maximize performance.

Figure 10-16 shows the back of a 3583 fibre tape library. The tape drives are on the right hand side of the diagram. To their left is the integrated SAN Data Gateway, with its two fibre ports at the top, and four SCSI ports at the bottom. Each SCSI port provides one SCSI bus. At the top left of the diagram are the ports that connect the media changer to the rest of the library.

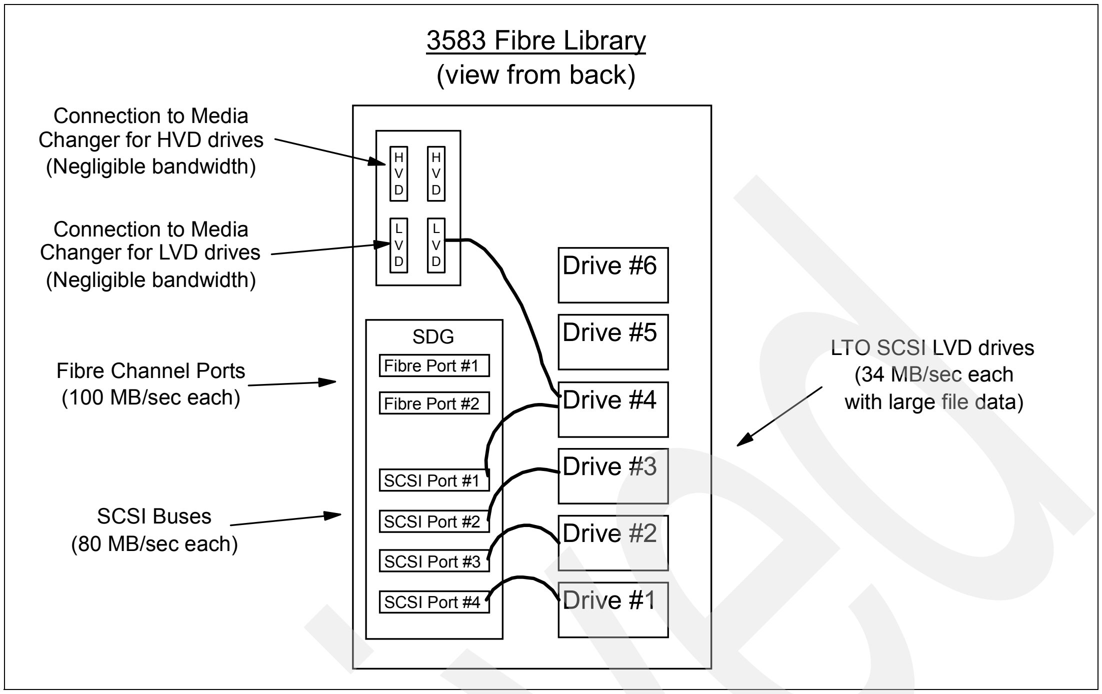

Figure 10-16 3583 SCSI bus cabling

To cable the library, 1-2 fibre cables come in to the SAN Data Gateway from the hosts. The SAN Data Gateway transfers the signal onto one of the SCSI buses. Each SCSI bus can be cabled to multiple tape drives via daisy-chaining. The media changer then needs to be connected to the library, by either connecting it to a SCSI port directly, or by daisy chaining it to one of the drives. Since the 3583 fibre library uses LVD when attached to iSeries, you would use the lower set of ports for the media changer. Note that terminator caps are required on the last drive in a daisy chain, and also on the port unused port of the pair that the media changer is using.

The bandwidths of the various items in this tape path are rated as follows:

Fibre ports: 100 MB/s per port, for each of save/restore

SCSI buses (ports): 80 MB/s per bus

LVD LTO Drives: 40 MB/s

SCSI Media Changer: negligible bandwidth required

Note that the speeds shown above are rated speeds. In real life, you will probably see less bandwidth due to overhead and other factors. For example, with large file data, the LTO drive will give closer to 34 MB/second.

When cabling the library, you should put the first four drives onto their own SCSI buses, rather than daisy-chaining to ensure you have plenty of SCSI bandwidth. When you add the fifth and sixth drives, you will need to daisy chain them to existing drives, each one to a different drive. You should avoid putting more than two drives onto one SCSI bus since three drives could exceed the capacity of the SCSI bus.

Since the media changer takes negligible bandwidth, it is typically daisy chained to a drive, rather than tying up an entire SCSI bus by attaching it directly to the SCSI port. The one scenario where the media changer might be attached to a separate SCSI bus would be if the customer was going to zone the SAN Data Gateway to restrict access to some of the drives. Since each zone needs to be able to access the media changer, and since zoning on a SAN

Data Gateway is at the SCSI bus level, the media changer would need to be on its own bus in this scenario.

10.6 Sample configurations

We now look at specific configurations in more detail for each type of tape drive, some valid and some not, to show how these configuration and performance rules work in practice. Note that although some of these configurations are not valid at V5R1, they may become valid in future releases as new code is written to enhance the SAN support.

Sample configurations for 3584 and 3590

The fibre 3590 and 3584 are very similar in their connectivity to the iSeries since each drive is a single target on the FC-AL connection. This is unlike the 3583 where multiple drives are presented via one target due to the SAN Data Gateway. Hence we will combine the 3590 and 3584 in our sample configurations. The one key difference between the 3590 and 3584 is that the 3590 has two fibre ports, whereas the LTO drives in the 3584 only have 1 fibre port each.

When planning LTO configurations, users should download the LTO Information APAR (II12621) that provides current information about this drive. This info APAR can be found at the following URL:

http://www.as400service.ibm.com/supporthome.nsf/document/10000035

Click the search key at the top of the panel, then fill in the Info APAR # (II12621).

Figure 10-17 shows a very simple configuration with a single iSeries attached to a single tape drive in the tape library, using a direct connection. iSeries #1 would have exclusive use of drive #1, and there would be plenty of bandwidth along the fibre cable to support it.

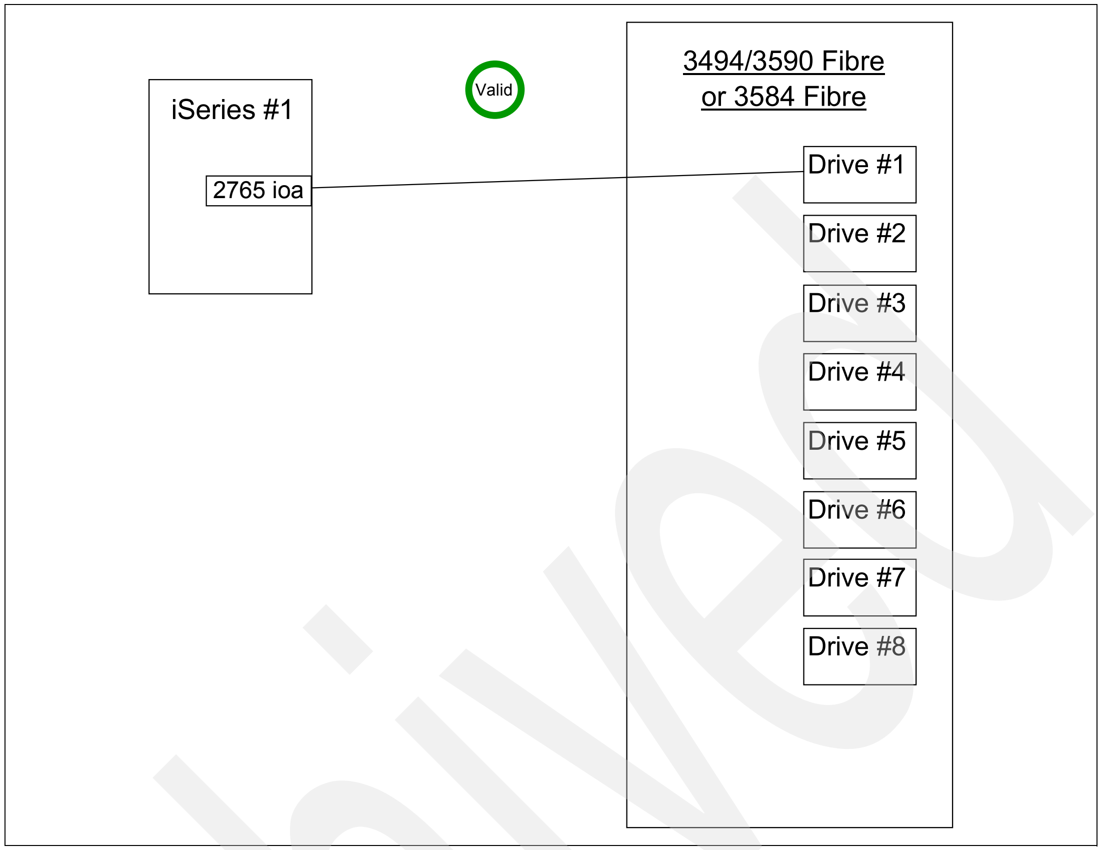

Figure 10-17 3494 or 3584 with one iSeries and one tape drive

Figure 10-18 shows a configuration that would allow the iSeries to access more than one drive in the tape library, either separately, or simultaneously. Once again, this configuration is direct-attached, and their would be plenty of bandwidth along the fibre cables to support the tape drives.

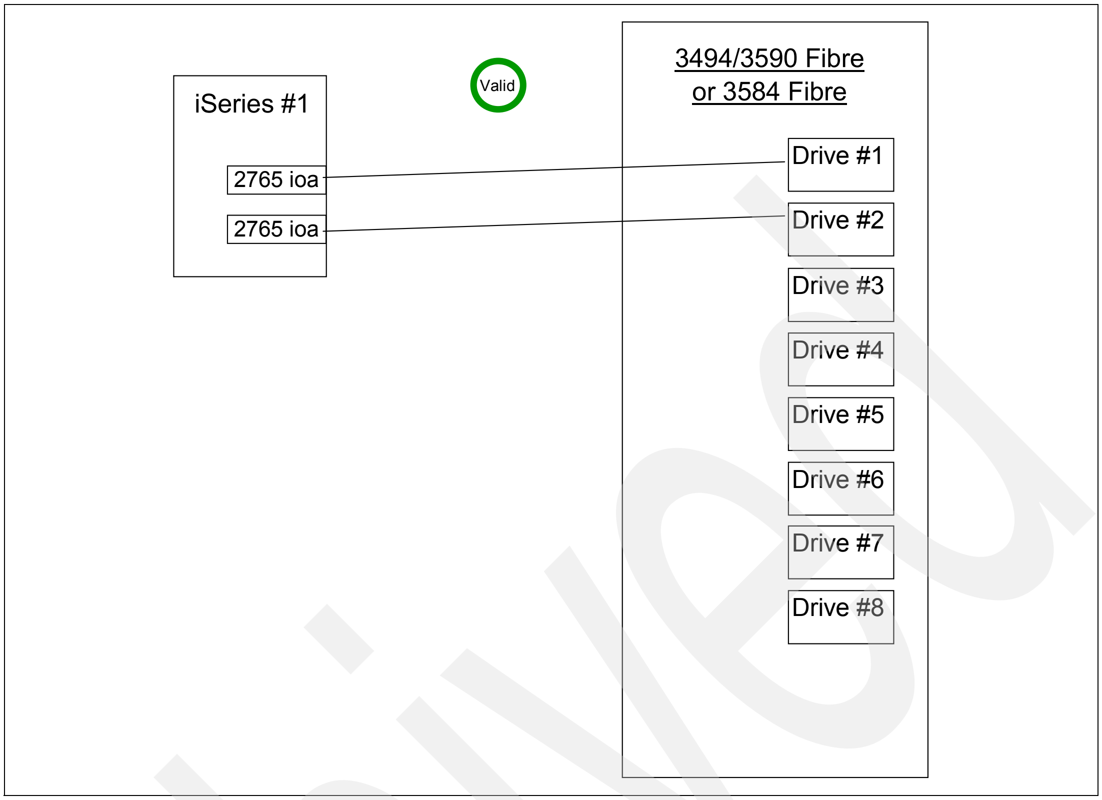

Figure 10-18 3494 or 3584 with one iSeries and two tape drives

Figure 10-19 shows two iSeries systems, each directly attached to a single drive in the tape library. Each iSeries can only access the tape drive that it is directly attached to. There would be plenty of bandwidth on each fibre cable to support the drive activity.

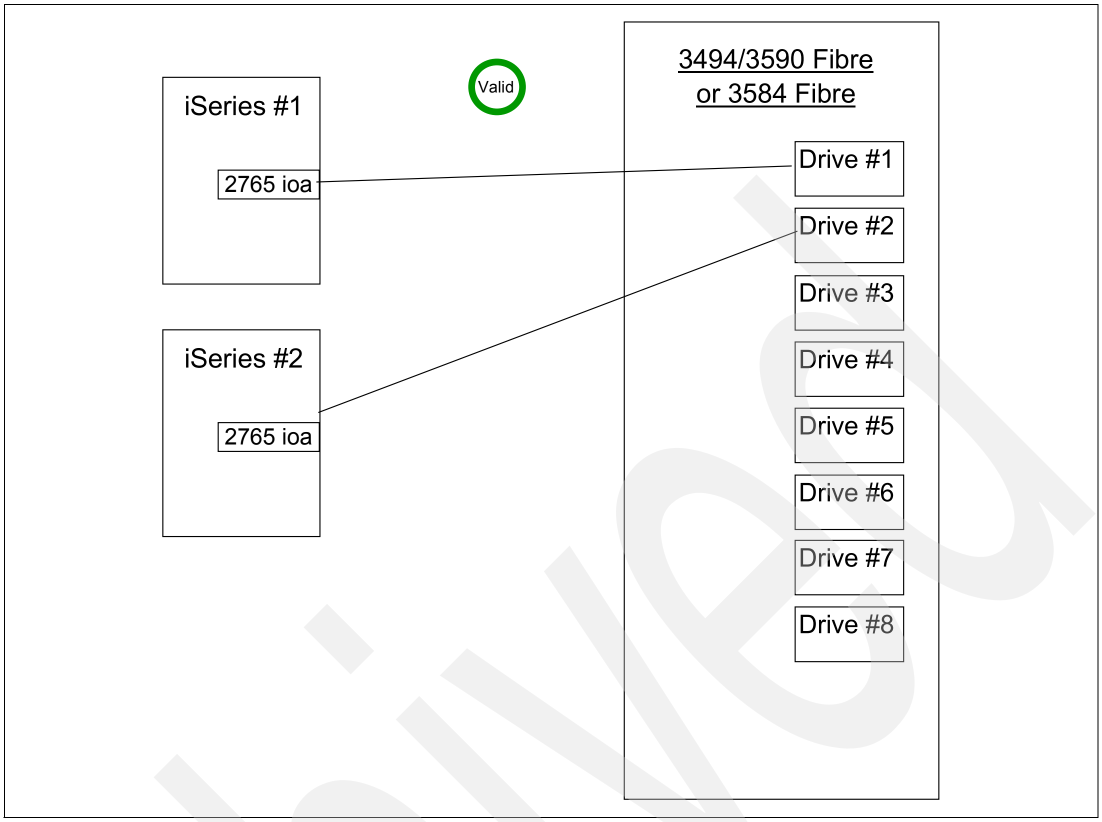

Figure 10-19 3494 or 3584 with two iSeries attached to one tape drive each

Figure 10-20 is only valid for the 3494 tape library. In this example, two iSeries systems are sharing a 3590E tape drive by attaching the systems directly to the two fibre ports on the drive. Each system has access to the drive, but only one system can be using it at a time. There is plenty of bandwidth on the fibre cables to support the drive activity. This configuration is not valid for the 3584, since the LTO drives only have one fibre port each.

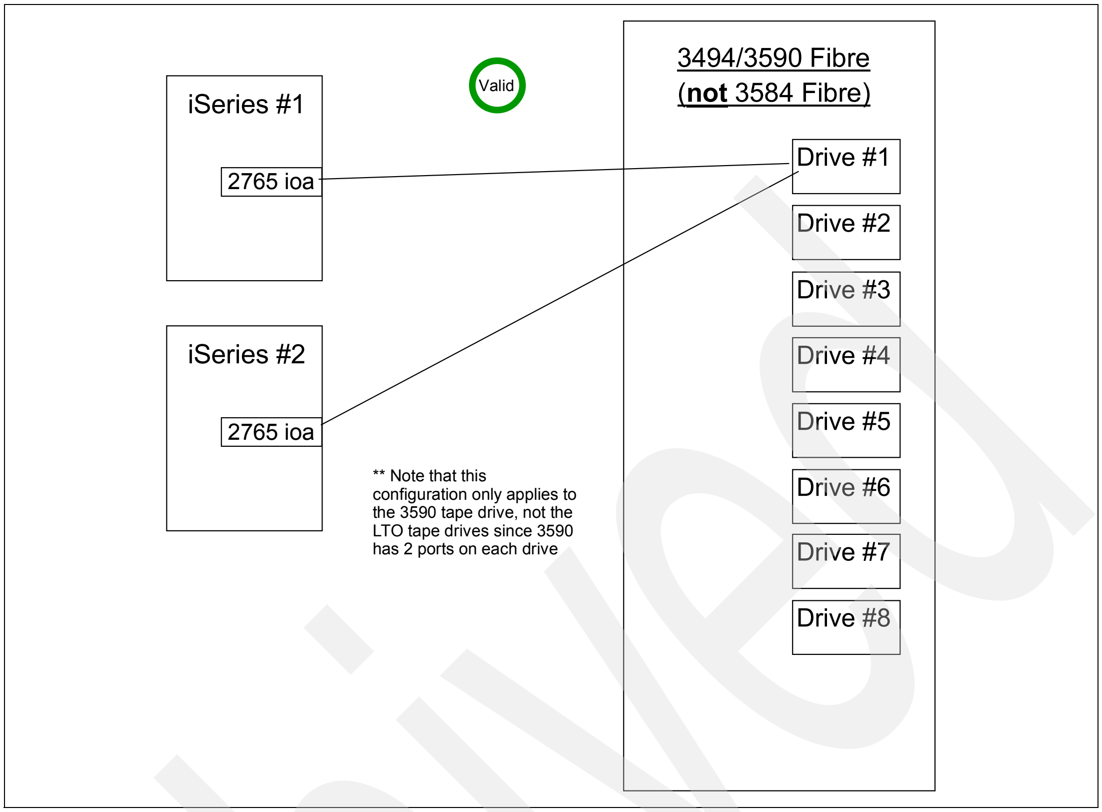

Figure 10-20 3494 with two iSeries sharing a 3590E tape drive

Figure 10-21 shows a configuration that is invalid at V5R1. In this scenario, two iSeries and two tape drives are attached to a hub without any zoning. As a result, we have one FC-AL with two tape adapters and two tape drives. The configuration is invalid because there is more than one target (i.e. tape drive) on the FC-AL. Also, each FC #2765 adapter can see two targets (i.e. tape drives) rather than one. When data and tape commands are sent from the host, it doesn't know which tape drive to apply them to. To correct this configuration, zoning would be required (see Figure 10-22).

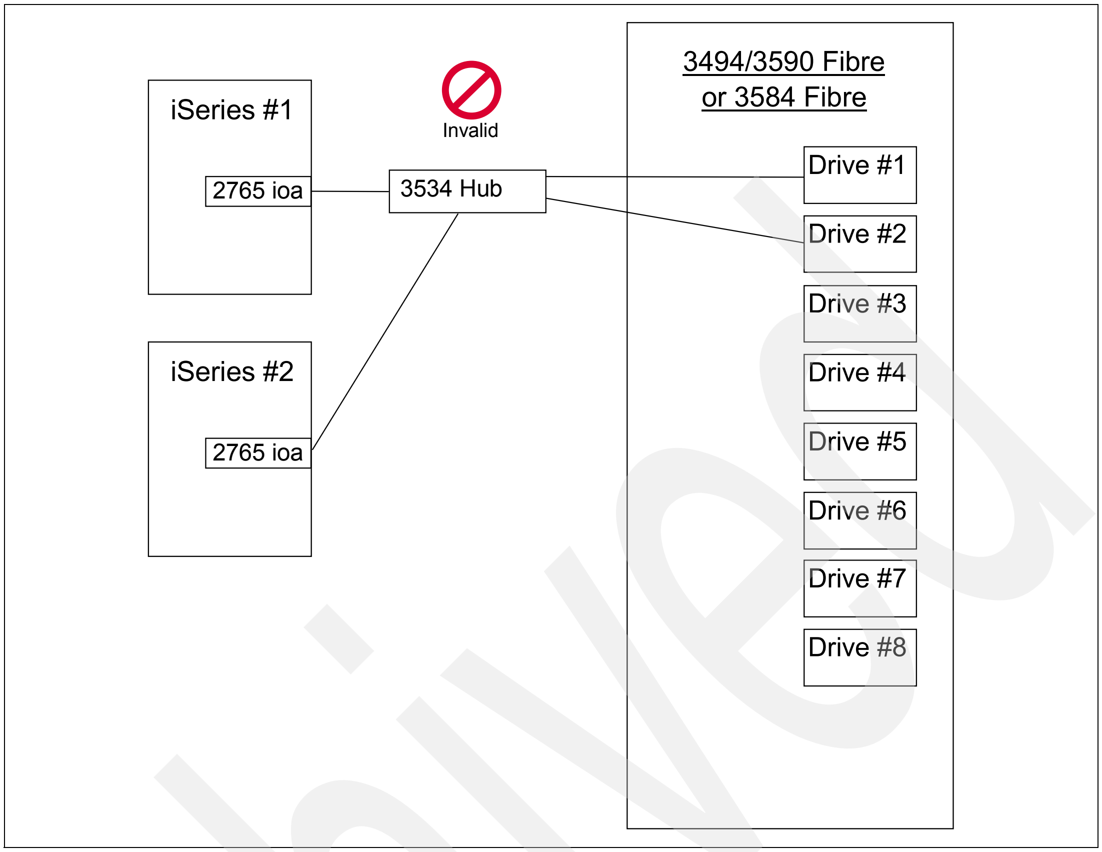

Figure 10-21 Invalid configuration with two iSeries attached to two drives via a hub

Figure 10-22 corrects the configuration in Figure 10-21 by adding zoning to limit the tape drives seen by each host. In this scenario, Zone A is a FC-AL on its own with one tape adapter and one tape drive. Zone B is an FC-AL on its own with a different tape adapter and a different tape drive. Each zone has only one target, and each adapter can only see one drive. iSeries #1 can only access drive #1 and iSeries #2 can only access drive #2. Because the 3534 is an intelligent hub rather than an unmanaged hub, each FC-AL has a full 100 MB/s of bandwidth, and hence there is plenty of capacity to run the attached drive, even if both drives are running simultaneously. In this scenario, a loop initialization primitive (LIP) is of no concern since there is only one host in each FC-AL, and hence no risk of a different host powering on during a save and impacting a save operation in progress. Note that although this configuration is valid, it is not sensible since the same thing could be accomplished at a lower cost using Figure 10-19.

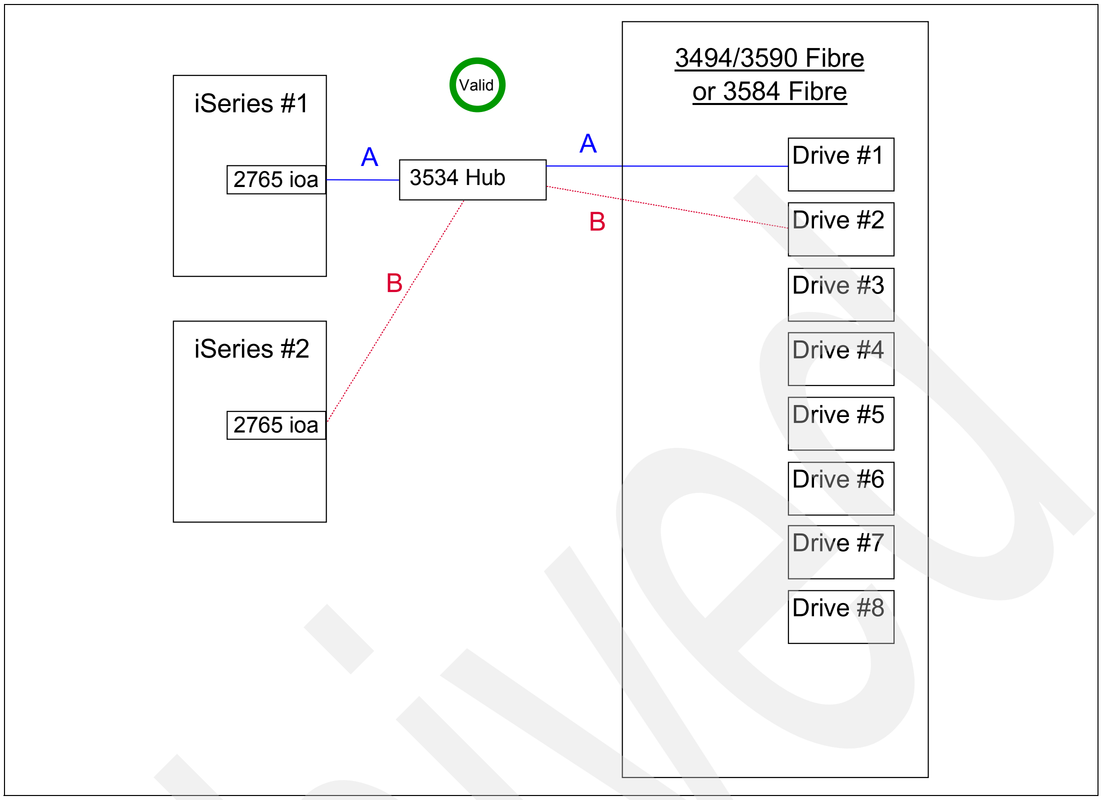

Figure 10-22 Two iSeries attached to two drives via a hub with zoning

A more sensible use of the SAN tape technology is shown in Figure 10-23. Here we use a hub to attach multiple systems to a single tape drive. Each iSeries is able to access drive #1, although they will need to schedule their time so they are not all trying to use it at once. No zoning is required since the FC-AL meets all of the rules for V5R1: there is a single target in the FC-AL, and each host adapter can only see one target. Since only one of the systems will be using the drive at a time, there is plenty of bandwidth in the FC-AL. LIP could be an issue if one of the systems powered on while another system was using a drive, but the iSeries has good error recovery to minimize the impact.

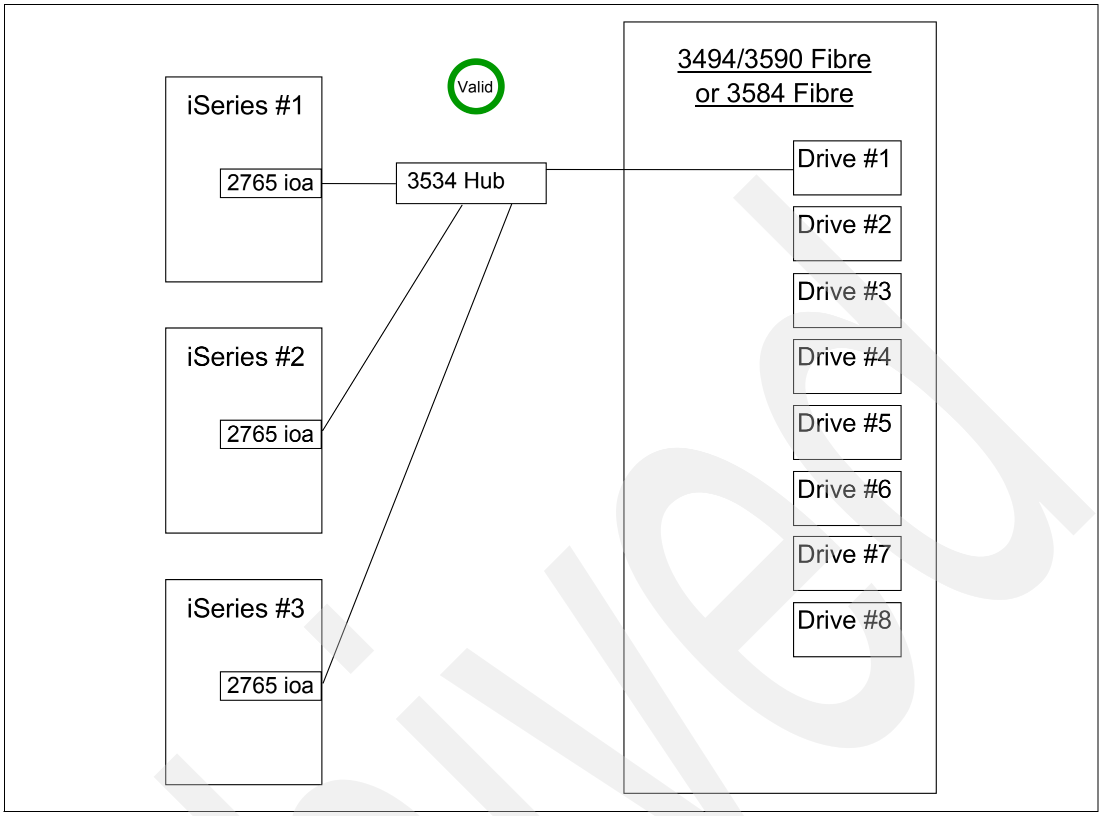

Figure 10-23 Three iSeries attached to one drive via a hub

Figure 10-24 is a configuration that considered whether overlapping zones would allow two iSeries to attach to two drives each without needing a second FC #2765 adapter for each iSeries. The idea was that zone A would have both host adapters and drive #1, and zone B would have both host adapters and drive #2. However, this configuration is not valid at V5R1, since it breaks the rule that each fc #2765 can only have one target. The zoning doesn't circumvent the problem, since each adapter is still trying to support multiple drives. When a tape command is sent to the fc #2765, it doesn't know whether to pass the command along to drive #1 or drive #2.

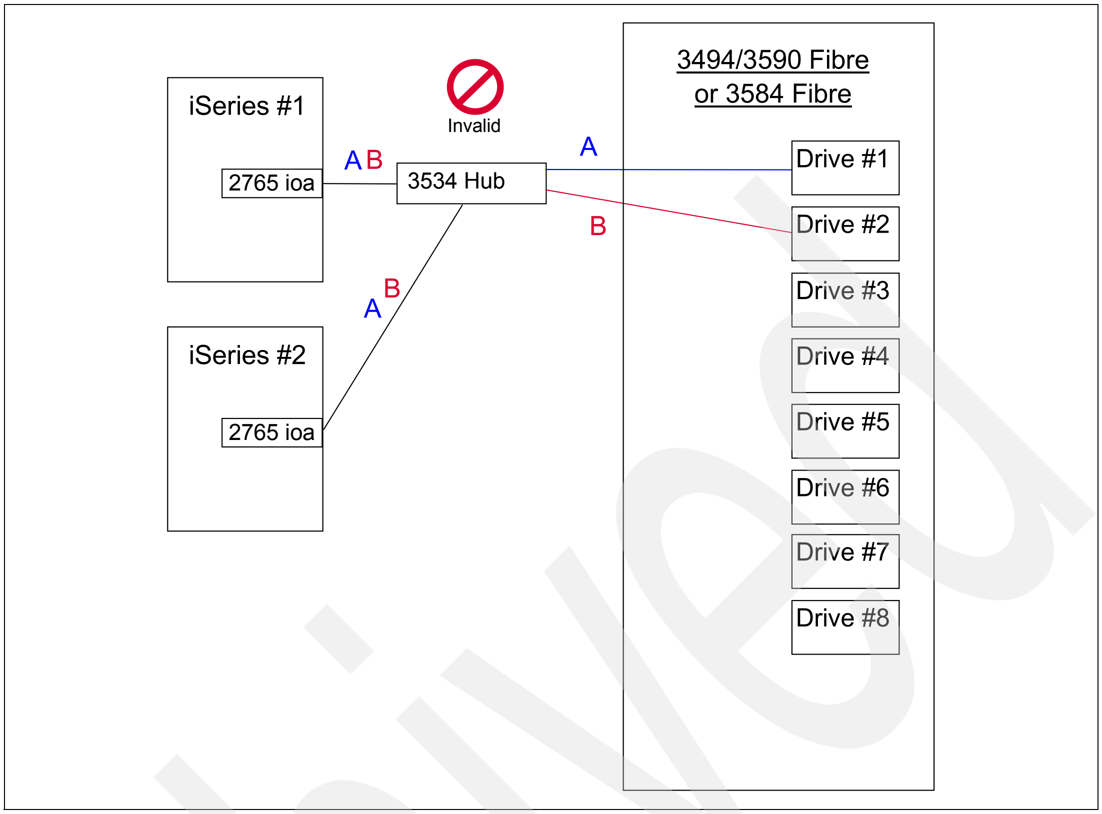

Figure 10-24 Invalid Configuration with two iSeries attached to two tape drives with overlapping zones

Figure 10-25 adjusts the configuration in Figure 10-24 so it is valid. Here we get rid of the overlapping zones, and add a second fc #2765 adapter in iSeries #2. Now Zone A has one host adapter from each iSeries and drive #1. Zone B has the second fc #2765 from iSeries #2 and drive #2. This configuration obeys the configuration guidelines. Both systems can access tape drive #1, and iSeries #2 can access tape drive #2 as well. Since the intelligent hub offers 100 MB/s in each zone, and since only one system can use a given tape drive at one time, there is plenty of bandwidth available. In zone A, LIP needs to be considered since there is more than one host, but as mentioned before, the iSeries has good error recovery and can usually recover from a LIP. Note that this configuration could also be accomplished by using the hub for zone A and direct-connecting iSeries #2 to drive #2, thus leaving ports available on the hub for other devices that require the multi-attach capability. If the drives are 3590Es as opposed to LTOs, then the hub is not required at all for this configuration since zone A could be replaced by direct connections to the two fibre ports on drive #1.

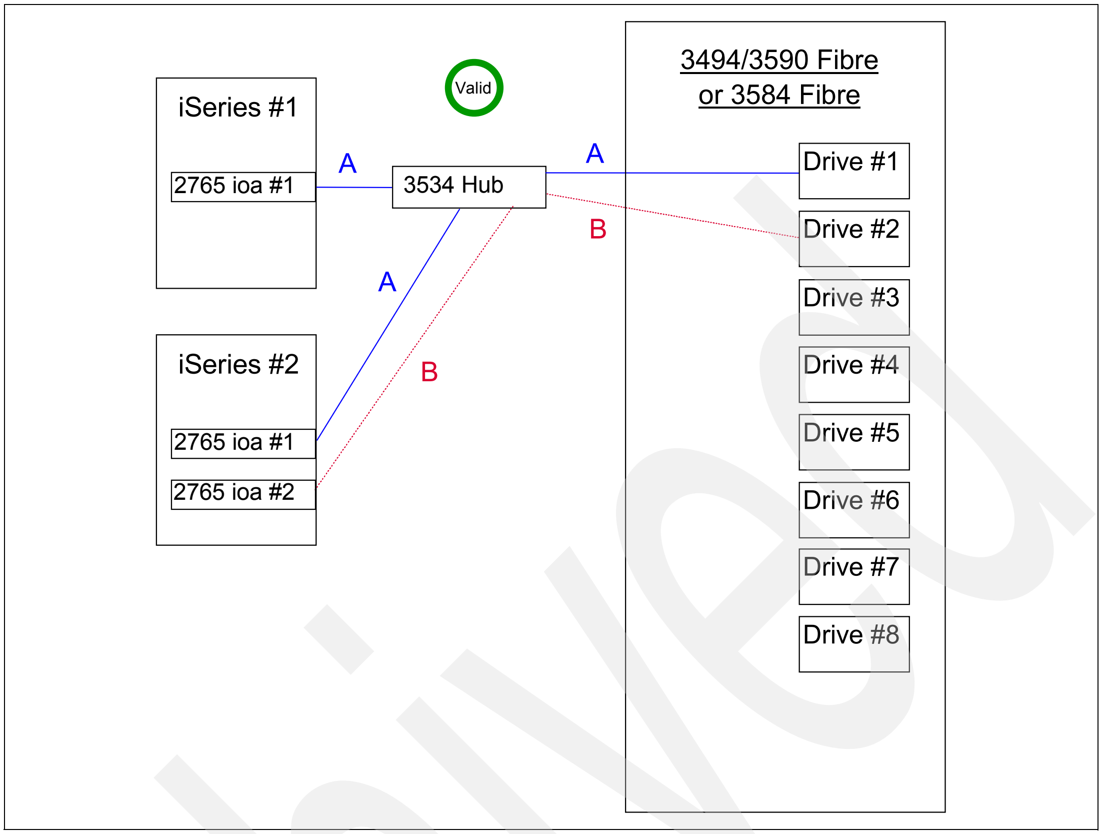

Figure 10-25 Two iSeries attached to one or two tape drives via a hub and multiple adapters with zoning

Figure 10-26 extends Figure 10-25 by adding an extra host adapter to iSeries #1 so both systems can access both drives. On the 3584, this is a useful configuration, but on 3494/3590E the same could be accomplished without the hub by direct-connecting the hosts to the two fibre ports on each 3590E. If there were more than two hosts, then a configuration similar to this would be helpful on the 3494/3590E also since more than two hosts could be attached to each drive. Note that in the future when the restriction of one host adapter per target drive is lifted, then a comparable configuration could be achieved with only one host adapter in each iSeries. The one drawback would be that if one system had multiple drives running at one time, they would be competing for bandwidth on the one host adapter.

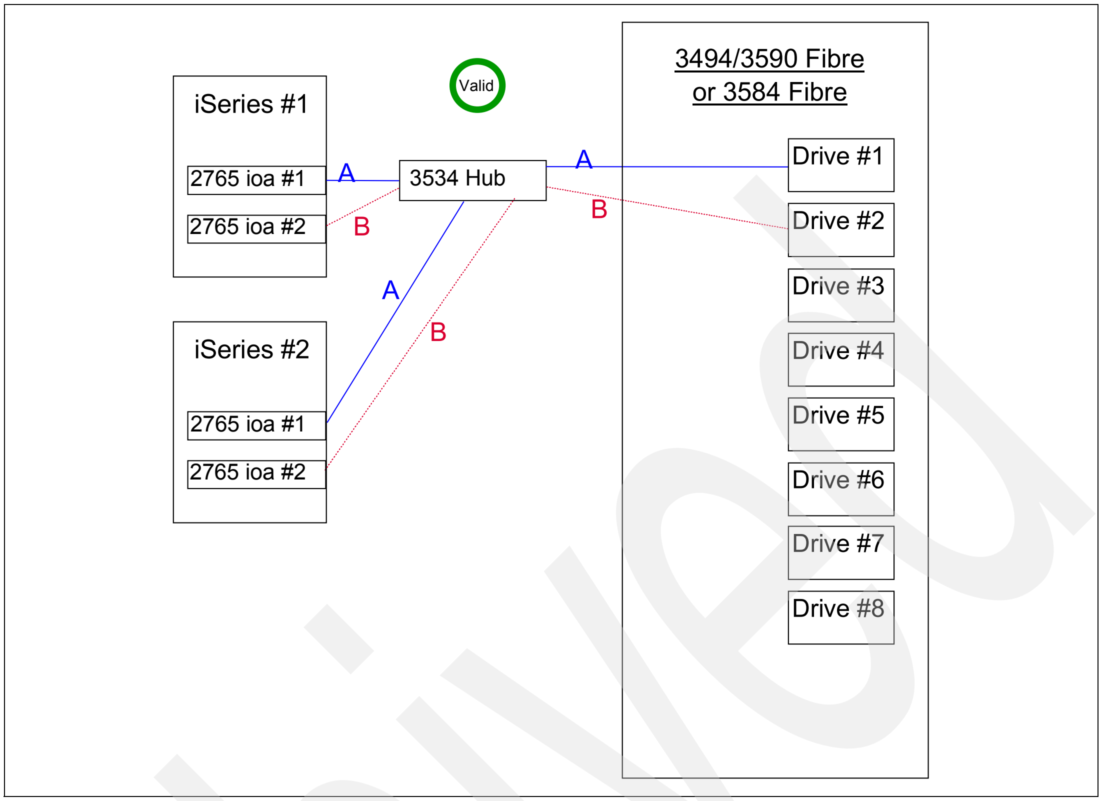

Figure 10-26 Two iSeries attached to two tape drives with a hub and multiple adapters with zoning

Sample Configurations for 3583

The 3583 fibre library introduces some different configuration scenarios since it presents itself as a single target with up to two fibre ports and up to six drives. This allows multiple drives to be attached to one host adapter since the rule of one target per host adapter can be met more easily than it can on 3494/3584 libraries where each drive is a target.

When planning LTO configurations, users should download the LTO Information APAR (II12621) that provides current information about this drive. This info APAR can be found at the following URL:

http://www.as400service.ibm.com/supporthome.nsf/document/10000035

Click the search key at the top of the panel, then fill in the Info APAR # (II12621).

Figure 10-27 shows a simple configuration where one iSeries is attached to all six drives in the 3583 via a direct connection to the fibre port. Notice that only one host adapter card is required for this connection since the 3583 behaves as a single target, rather than each drive being a single target. Bandwidth needs to be considered, however, since the link is rated at 100 MB/s in each direction, and in actual fact would only achieve part of that throughput, so if multiple drives were running simultaneously, performance could become bottle necked.

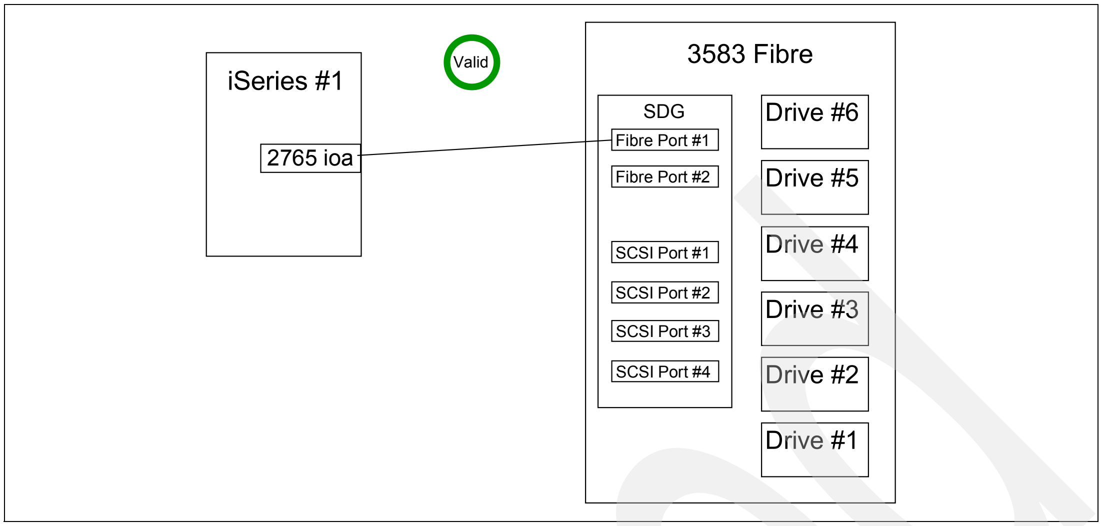

Figure 10-27 One iSeries attached to six drives in a 3583

Figure 10-28 is a configuration that considered whether more bandwidth could be provided when a single iSeries was attached to all six drives in the 3583. This configuration tried adding a second host adapter card in the iSeries and attaching it to the second fibre port. However, this configuration is invalid since the iSeries can see the same tape drives via two different paths. Theoretically, this configuration could be made to work by zoning the integrated SAN Data Gateway so some of the drives were seen via the first Fibre Channel port and the rest of the drives were seen via the second Fibre Channel port. In this configuration, the media changer would need to be attached directly to a SCSI port so it could be included in each zone to allow tape mounts and dismounts. Such a configuration is discouraged in the iSeries world due to complexity.

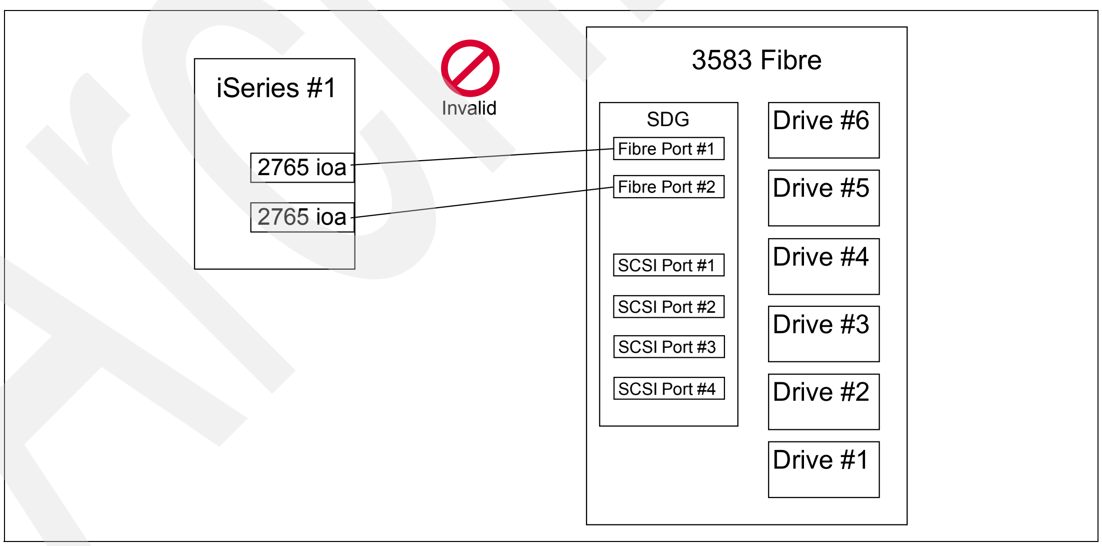

Figure 10-28 Invalid Configuration with one iSeries dual-attached to the 3583

Figure 10-29 shows how two iSeries could be given access to all six drives in the 3583, with only a single adapter in each host. Bandwidth is a consideration if a system is using multiple drives at once since each system has only 100 MB/s in each direction on its fibre link. LIPs are not a concern since each fibre connection behaves independently.

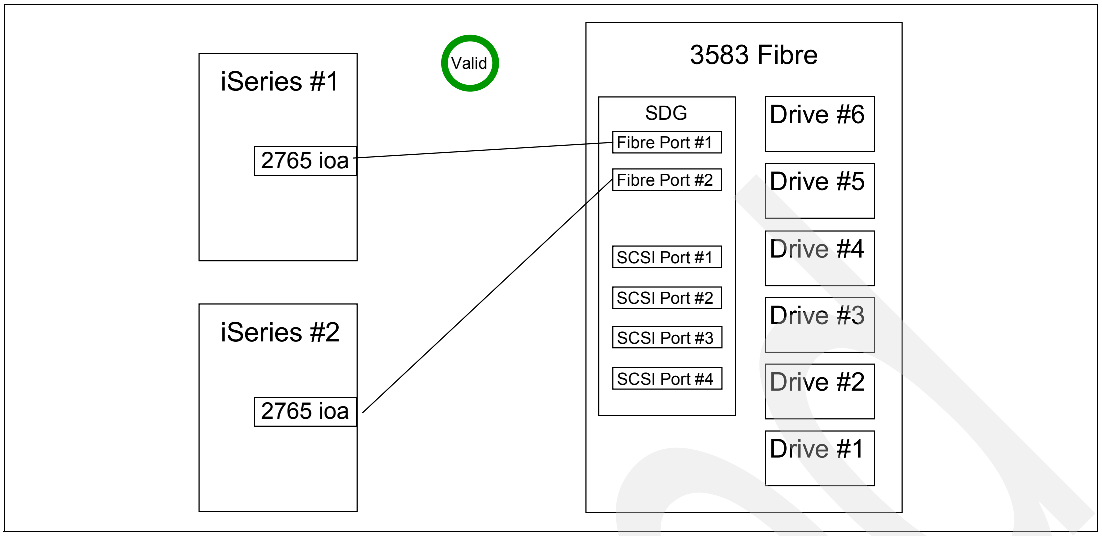

Figure 10-29 Two iSeries attached to six drives in the 3583

In Figure 10-30, we introduce a hub to the configuration. The example shown will give both iSeries systems access to all six tape drives. However, this configuration is not sensible because the same connectivity could be accomplished at a lower price without the hub as shown in Figure 10-29. In addition, Figure 10-29 offers double the bandwidth of Figure 10-30. LIP would be a consideration in this example, but as discussed earlier, the iSeries has good error recovery routines to minimize the impact of a LIP.

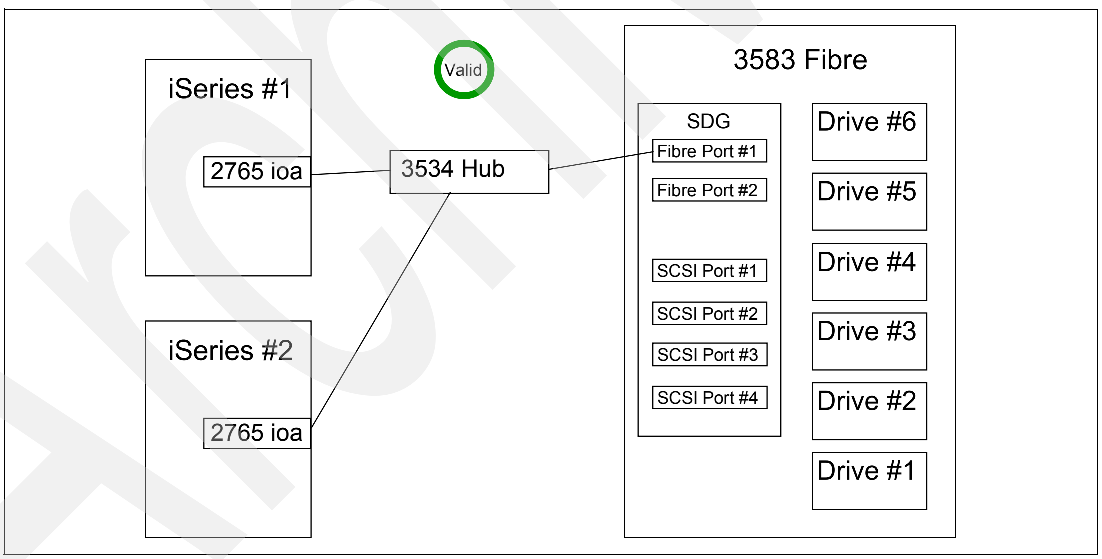

Figure 10-30 Two iSeries attached to a 3583 via a hub

Figure 10-31 shows a configuration with a hub that does make sense since it lets three iSeries systems access all six drives. This would not be possible via direct connection since only 2 hosts can direct attach. If multiple drives will be running simultaneously, performance will be a consideration since the single link between the hub and the 3583 limits throughput to 100 MB/s in each direction. LIP would be a consideration in this configuration, but the iSeries error recovery routines would minimize the impact.

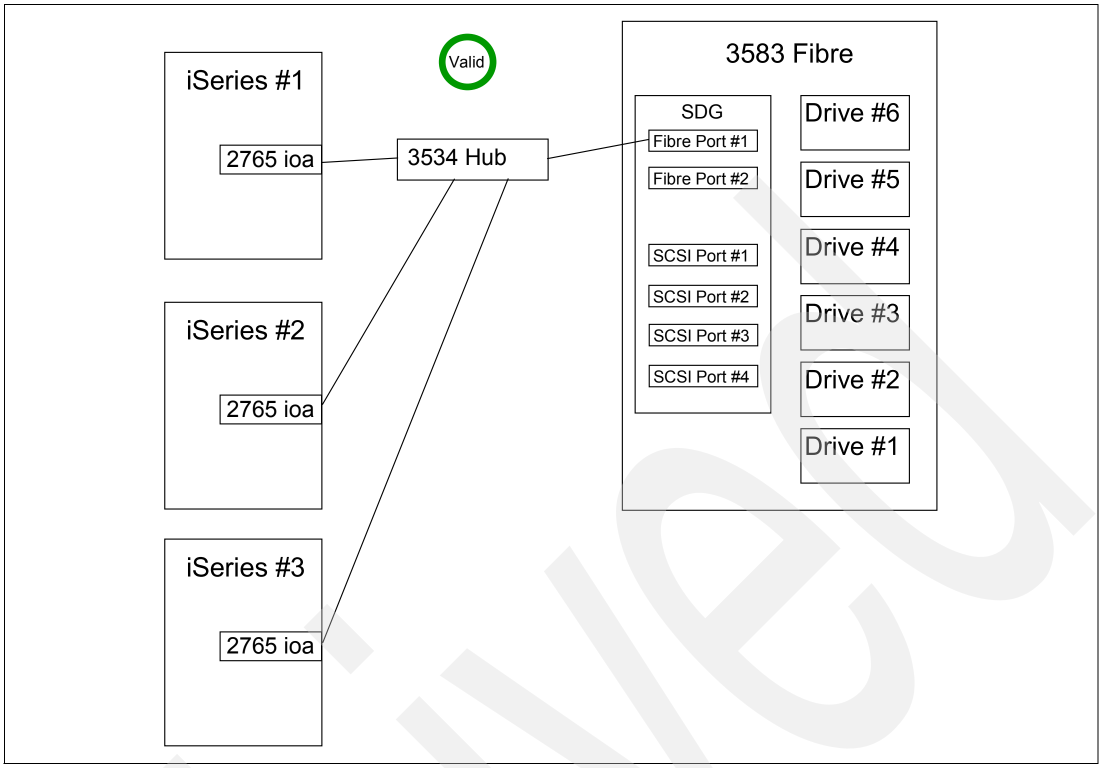

Figure 10-31 Three iSeries attached to all six drives in a 3583 via a hub

Figure 10-32 shows a scenario with two iSeries where a hub connection would be useful, since it would offer a long distance connection. As in the previous examples, both iSeries can access all six drives, although they will need to be aware of the bandwidth available when running multiple drives simultaneously. LIP needs to be considered, but the iSeries recovery routines would minimize any impact. Notice that two hubs are required, since the long-wave fibre cable that covers the 10 km distance cannot plug directly into the drive or host adapter. The distance shown is 10 km between the hubs and 11 km altogether, but longer distances could be achieved by slowing the clock speed on the hubs, although this would impact performance.

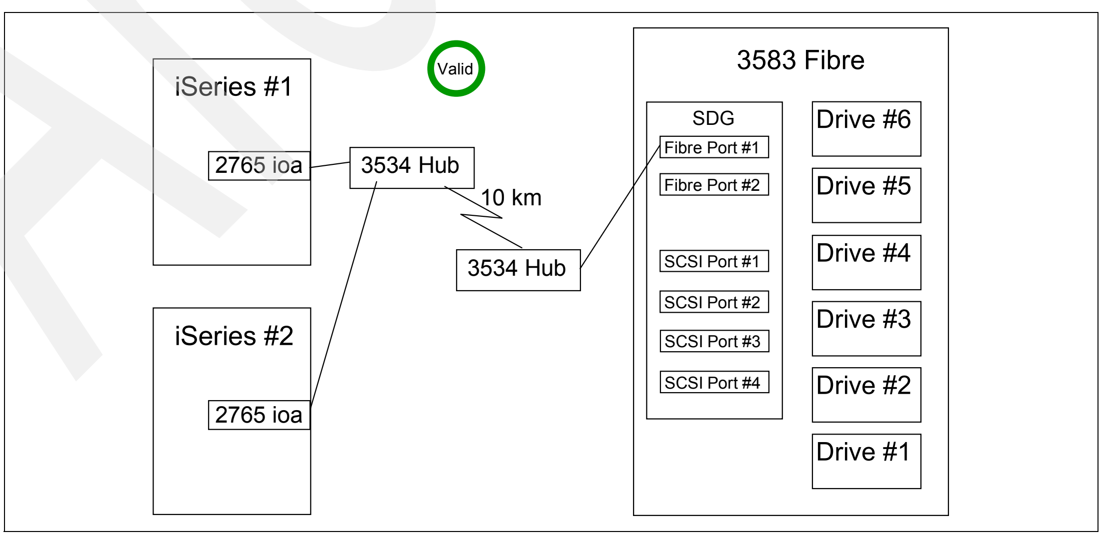

Figure 10-32 Two iSeries attached to a 3583 over a long distance

Figure 10-33 is a configuration that considered whether more bandwidth could be provided to the two attached iSeries by adding an extra link from the hub to the 3583. However, the configuration as shown is invalid because the hosts can see the same drives across two different paths. This configuration could be made valid by zoning iSeries #1 to the upper fibre port, and iSeries #2 to the lower fibre port. Note that neither of these configurations is sensible since the same connectivity and performance could be accomplished by direct-connecting the iSeries to the fibre ports as shown in Figure 10-29. However, if there were more iSeries hosts involved, then a configuration such as this with a hub and zoning would become interesting.

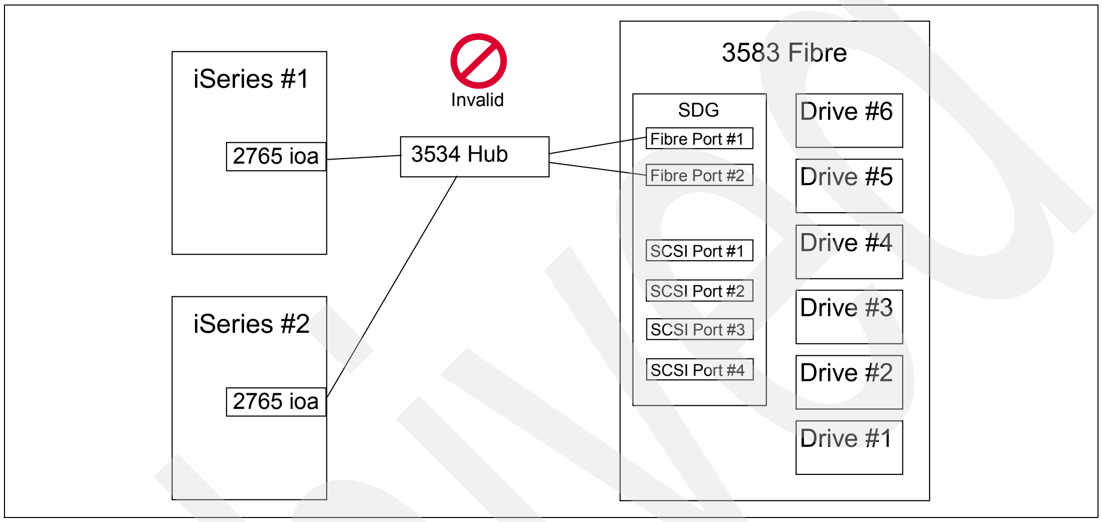

Figure 10-33 Invalid Configuration with two iSeries dual-connected to a 3583 via a hub

Figure 10-34 extends the configuration shown in Figure 10-33 by adding a third iSeries host to make the hub useful, and by adding zoning to make the configuration valid. In this example, each iSeries has access to all six drives. iSeries #1 and iSeries #2 will be sharing 100 MB/s of bandwidth in each direction, and iSeries #3 will have 100 MB/s of bandwidth all to itself. Care should be taken deciding which hosts to put in the shared zone to keep within the bandwidth capabilities. LIP will be a consideration in zone A since there are two hosts, but not LIP.

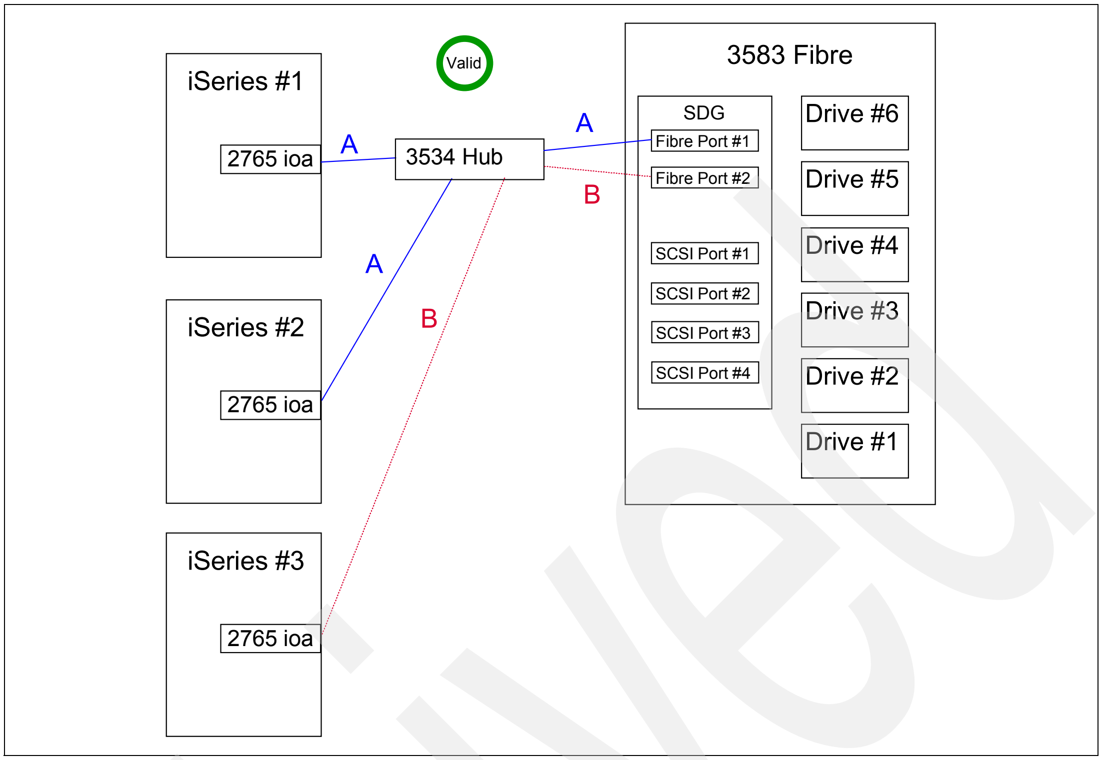

Figure 10-34 Three iSeries dual-attached to a 3583 via a hub with zoning

Figure 10-35 shows an alternate way to attach multiple hosts to the 3583 via multiple hubs. This configuration would not require zoning since each hub only has one target. It would also allow more hosts to have access to the 3583 due to the additional ports on the second hub. The two iSeries on each hub would share the 100 MB/s bandwidth in each direction provided by that hub and 3583 fibre port. LIP would be a consideration within each hub, but once again, the iSeries error recovery would minimize the impact.

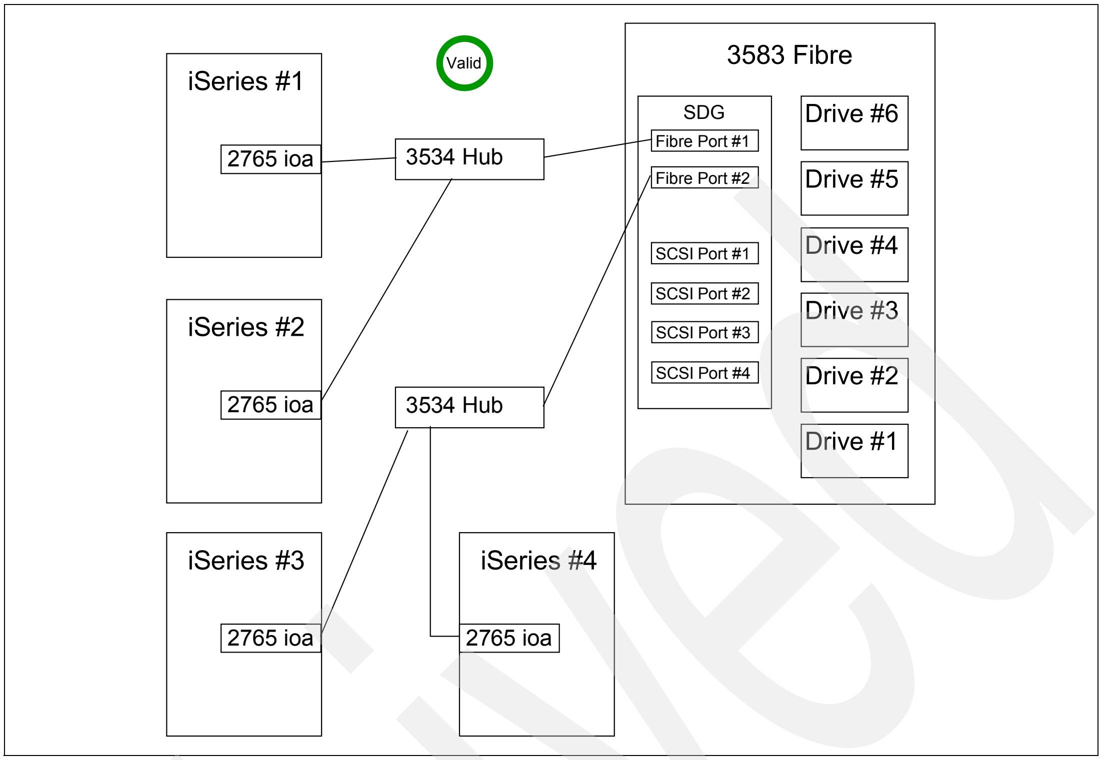

Figure 10-35 Four iSeries attached to a 3583 library via two hubs

Figure 10-36 shows a multi-platform attachment scenario for the 3583. Notice that zone A is a homogeneous iSeries zone with a single target fibre port on the 3583. Zone B is a homogenous pSeries zone with the other 3583 fibre port as its target. All hosts would have access to all drives, so it is important that the operating system and application software using the drive uses reserve/release so one host won't steal the drive from another host that already has a save in progress. In addition, controls would need to be in place to prevent one host from overwriting another hosts tapes, both among hosts of one platform, and between platforms.The iSeries hosts would share 100 MB/s of bandwidth each direction, as would the pSeries hosts. LIPs would be a consideration within each FC-AL, but the iSeries error recovery would minimize the impact within zone A.

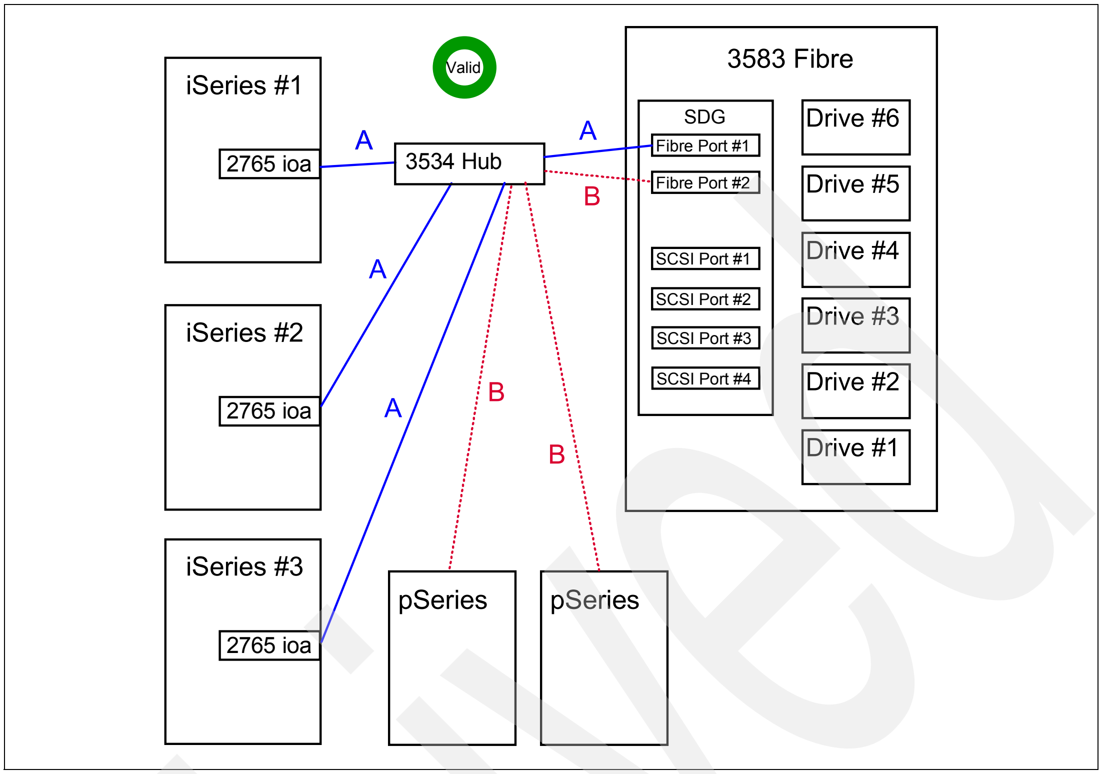

Figure 10-36 Multi-platform attachment to 3583

10.7 Implementing Fibre Channel tape

In order to implement Fibre Channel tape, the following steps should be followed:

- Select the appropriate tape technology and configuration to meet your requirements.
- Purchase the hardware, including tape drives, host adapter cards, cables, and hubs/switches with Quickloop software as required. Note that cable adapter kits may be required since the iSeries fibre cards come with the new LC adapters, whereas the older, larger SC adapters are more prominent in the cable marketplace today. You may also need to purchase tape media if you have migrated to a new tape technology.
- Plan your implementation to coordinate with other system activities.
- Set up any switches and hubs and fibre cabling that will be required. You will need to work with your LAN and SAN administrators to get the appropriate device names, TCP/IP addresses, and zoning. Note that both the 3534 hub and the 2109 switch have password protected Web browser interfaces that let you interact with the devices from the convenience of your own desk. You will want to set these up and become familiar with them. SAN education may also be in order.
- Install the tape hardware according to the documentation provided. The 3583 is a customer setup device, whereas the 3584 and 3494 will be set up by the IBM customer engineer. In addition to setting up the hardware, there are also Web browser interfaces for some of the devices. On the 3583 and 3584 there is a password protected Remote Management Utility (RMU) that lets you interact with the front panel of the libraries from the convenience of your own desk, and also look up various information about the library. In addition, the integrated SAN Data Gateway on the 3583 also has a Web browser interface that you may want to set up, although iSeries customers are not likely to use it too often.

- Connect the tape hardware to the iSeries. As with any tape drive, you should check/adjust your device names and resource names so they match and are meaningful. Instructions are included in “Changing device description and resource names” on page 286.
- Adjust any applications that access tape drives so they will use the new devices. If you have a tape management package such as BRMS, make any configuration changes that are required, such as configuring the new tape device and media class and adjusting backups to use the new device.
- If you have changed from 3590B to 3590E tape drives, you should consider reformatting your media to *FMT3590E as it expires to take advantage of the improved performance and media density. This process needs to be managed carefully since you can’t tell visually which format the tape is.
- If you plan to use your new tape drive for full system recoveries, then you need to familiarize yourself with the procedures for an alternate installation device since fibre attached tape drives cannot be used as alternate-IPL devices. Details are included in “Alternate IPL device and Alternate Installation Device” on page 294.

10.7.1 Changing device description and resource names

As we enter the SAN world, more and more devices are going to be visible to more and more systems throughout the enterprise. As a result, naming conventions become very important. We strongly recommend that every tape device in the enterprise should have a unique name, and that all systems should call that device by the same name.

In addition, it has always been recommended on the iSeries that the resource name for a device should match the device description name to avoid confusion. This is particularly important with media libraries since the drive resource names are shown when you view them in the media library (WRKMLBSTS command), and the device description names are shown when you view the standalone devices (WRKCFGSTS *DEV *TAP command).

The following sections will explain how to change the resource names and device description names.

Changing Device Description Names

Device descriptions can be renamed from the WRKDEVD screen as shown in Figure 10-37. Type option 7=Rename beside the device to be renamed and key in the new device name desired.

.

```
                          Work with Device Descriptions
                                                              System:   PRIMARY
Position to  . . . . .                Starting characters

Type options, press Enter.
  2=Change   3=Copy    4=Delete   5=Display   6=Print   7=Rename
  8=Work with status   9=Retrieve source

Opt  Device      Type        Text
 7   TAPMLB01    3590        CREATED BY AUTO-CONFIGURATION
     TAP01       6383        CREATED BY AUTO-CONFIGURATION
     TAP05       3590        CREATED BY AUTO-CONFIGURATION







                                                                         Bottom
Parameters or command
===>
F3=Exit   F4=Prompt   F5=Refresh   F6=Create   F9=Retrieve   F12=Cancel
F14=Work with status
```

Figure 10-37 Changing a device name

Changing hardware resource names

Resource names can be changed in System Service Tools (SST). You will need sufficient authority to access these panels.

Attention: SST is a powerful tool and should be used with caution. Do not use SST if you are not familiar with the facilities provided. See the warning on the Start a Service Tool screen shown in Figure 10-39.

1. Type STRSST to get to the System Service Tools Panel, then select option 1 to start a service tool.

```
                         System Service Tools (SST)

 Select one of the following:

      1. Start a service tool
      2. Work with active service tools
      3. Work with disk units
      4. Work with diskette data recovery
      5. Work with system partitions
      6. Work with system capacity










 Selection
      1

 F3=Exit         F10=Command entry         F12=Cancel
```

Figure 10-38 Start a service tool

2. There are many service tools available. We need to work with the hardware service manager. Choose option 7=Hardware service manager from the Start a Service Tool panel shown in Figure 10-38.

```
                          Start a Service Tool

Warning: Incorrect use of this service tool can cause damage
to data in this system.  Contact your service representative
for assistance.

Select one of the following:

     1. Product activity log
     2. Trace Licensed Internal Code
     3. Work with communications trace
     4. Display/Alter/Dump
     5. Licensed Internal Code log
     6. Main storage dump manager
     7. Hardware service manager




Selection
     7

F3=Exit        F12=Cancel        F16=SST menu
```

Figure 10-39 Starting the hardware service manager

3. The Hardware Service Manager provides a lot of facilities, many of which can have a severe impact on your system. We are only interested in changing the resource details. You can use Option 3=Locate resource by resource name if you know the name of the resource you want to change. However, if you are not sure, it is better to view all resources and choose the correct one from those shown. To do this, select Option 2=Logical Hardware Resources from the Hardware Service Manager panel shown in Figure 10-40.

```
                           Hardware Service Manager

Attention:  This utility is provided for service representative use only.

   System unit . . . . . . . :   9406-820 10-5GG8M
   Release . . . . . . . . . :   V5R1M0 (1)

Select one of the following:

   1. Packaging hardware resources (systems, frames, cards,...)
   2. Logical hardware resources (buses, IOPs, controllers,...)
   3. Locate resource by resource name
   4. Failed and non-reporting hardware resources
   5. System power control network (SPCN)
   6. Work with service action log
   7. Display label location work sheet
   8. Device Concurrent Maintenance

Selection
     2

F3=Exit      F6=Print configuration      F9=Display card gap information
F10=Display resources requiring attention      F12=Cancel
```

Figure 10-40 Hardware service manager

4. The resources we want to work with are on the system bus. Select Option 1=System bus resources from the Logical Hardware Resources panel shown in Figure 10-41.

```
                           Logical Hardware Resources


Select one of the following:

     1. System bus resources
     2. Processor resources
     3. Main storage resources
     4. High-speed link resources













Selection
     1

F3=Exit      F6=Print configuration      F12=Cancel
```

Figure 10-41 Selecting the system bus resources

5. You will now be presented with a list of all the logical hardware resources on the system bus. Scroll down this list until you see the IOP supporting your tape adapter and tape devices. Select the appropriate IOP with Option 9=Resources associated with IOP as shown in Figure 10-42. If you find that you have chosen the wrong IOP, you can use F12 to back out one panel and try another IOP instead.

```
                  Logical Hardware Resources on System Bus

 System bus(es) to work with . . . . . .   *ALL  *ALL, *SPD, *PCI, 1-511
 Subset by . . . . . . . . . . . . . . .   *ALL  *ALL, *STG, *WS, *CMN, *CRP

 Type options, press Enter.
   2=Change detail    4=Remove    5=Display detail    6=I/O Debug
   7=Display system information
   8=Associated packaging resource(s)     9=Resources associated with IOP

                                                               Resource
 Opt Description                   Type-Model  Status          Name
          System Bus                28AB-      Operational     LB09
           Multi-adapter Bridge     28AB-      Operational     PCI10D
         Bus Expansion Adapter         -       Operational     BCC10
          System Bus                28AB-      Operational     LB10
           Multi-adapter Bridge     28AB-      Operational     PCI05D
  9         Combined Function IOP   2843-001   Operational     CMB05
           Multi-adapter Bridge     28AB-      Operational     PCI12D
                                                                       More...
 F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
 F9=Failed resources      F10=Non-reporting resources
 F11=Display serial/part numbers      F12=Cancel
```

Figure 10-42 Selecting the IOP supporting the resource to be renamed

6. The resulting panel will show all the resources supported by the chosen IOP. Identify the resource name you want to change and select it with Option 2=Change detail as shown in Figure 10-43.

```
              Logical Hardware Resources Associated with IOP

 Type options, press Enter.
   2=Change detail    4=Remove    5=Display detail    6=I/O Debug
   7=Verify           8=Associated packaging resource(s)

                                                                  Resource
 Opt Description                    Type-Model  Status            Name
     Combined Function IOP           2843-001   Operational       CMB05
      Communications IOA             2838-001   Operational       LIN03
       Communications Port           2838-001   Operational       CMN04
      Communications IOA             2745-001   Operational       LIN05
       Communications Port           2745-001   Operational       CMN23
       Communications Port           2745-001   Operational       CMN24
      Multiple Function IOA          2768-001   Operational       DC03
      Multiple Function IOA          2765-001   Operational       DC01
  2    Tape Library                  3590-E11   Operational       TAPMLB01
        Tape Unit                    3590-E11   Operational       TAP05


 F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
 F9=Failed resources      F10=Non-reporting resources
 F11=Display serial/part numbers      F12=Cancel
```

Figure 10-43 Selecting the resource name to be changed

7. You will then be presented with the Change Logical Hardware Resource Detail screen. Key in the new name for this resource and press Enter to change the resource name as shown in Figure 10-44. You would normally choose a new resource name that matched the device description name set in "Changing Device Description Names" on page 286. This name should be unique in the enterprise.

```
                    Change Logical Hardware Resource Detail

Description . . . . . . . . . . . . . :  Tape Library
Type-Model  . . . . . . . . . . . . . :  3590-E11
Status  . . . . . . . . . . . . . . . :  Operational
Serial number . . . . . . . . . . . . :  13-00000

Current resource name . . . . . . . . :  TAPMLB01


Type changes, press Enter.

New resource name . . . . . . . . .      NEWNAME









F3=Exit      F5=Refresh      F6=Print
F9=Display detail      F12=Cancel
```

Figure 10-44 Changing the resource name

8. Press F3=Exit to return to the Logical Hardware Resource Associated with IOP screen where you should see the new name of the resource displayed as shown in Figure 10-45.

```
             Logical Hardware Resources Associated with IOP

Type options, press Enter.
  2=Change detail    4=Remove    5=Display detail    6=I/O Debug
  7=Verify          8=Associated packaging resource(s)

                                                               Resource
Opt Description                       Type-Model  Status          Name
    Combined Function IOP              2843-001   Operational     CMB05
     Communications IOA                2838-001   Operational     LIN03
      Communications Port              2838-001   Operational     CMN04
     Communications IOA                2745-001   Operational     LIN05
      Communications Port              2745-001   Operational     CMN23
      Communications Port              2745-001   Operational     CMN24
     Multiple Function IOA             2768-001   Operational     DC03
     Multiple Function IOA             2765-001   Operational     DC01
      Tape Library                    3590-E11   Operational     NEWNAME
       Tape Unit                      3590-E11   Operational     TAP05


F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources
F9=Failed resources      F10=Non-reporting resources
F11=Display serial/part numbers      F12=Cancel
```

Figure 10-45 TAPMLB01 renamed to NEWNAME

Changing the resource name in this way automatically updates the RSCNAME field in the device description. Note, it does not change the device description name itself - only the resource associated with the device descriptions which use that resource.

You should repeat this process on all systems for all tape devices to ensure that each device and associated resource names match and all tape devices and resources have unique names. One exception to this might be that you have TAP01 as the internal tape device for all systems.

10.7.2 Alternate IPL device and Alternate Installation Device

In the early days of OS/400, every system had an "alternate IPL device". This was a tape drive attached to a tape adapter in a certain card slot on the first system bus. If system problems occurred that required an IPL from tape rather than disk (e.g. a disk failure requiring a full system reload or a corrupted load source disk, etc.), then OS/400 would look for the *SAVSYS or distribution media tape in the alt-IPL tape drive and start the system from there.

Over the years, as tape drives became faster and faster, people wanted to put their fastest tape drives on buses other than the system bus to avoid contention. However, they also wanted to be able to use these drives as their alternate-IPL devices. As a result, new function was added to V4R1 of OS/400 to allow the system to use a device on the first system bus such as an optical drive or internal tape drive to install or recover enough of the Licensed Internal Code to perform a D-mode IPL. Then the system lets you choose a drive elsewhere on the system to use for the rest of your installation/recovery. This drive is called the alternate installation device.

The #2765 Fibre Channel adapter does not support the Alternate IPL function, but it can be used as an alternate installation device. Right after installing your new Fibre Channel tape drive, you should use the following procedures to prepare to use it as an alternate installation device.

Setting the alternate installation device

Before you use an alternate installation device, you need to ensure that you define it on a bus, and you must enable the device. You need to record and keep the logical address of the system bus and system card at which the alternate installation device is attached. If you have the device set up on a different bus and you do not have this information available, you will not be able to complete installations or recoveries.

If you change the configuration of your system, you need to ensure that this information is correct before you start to use the alternate installation device. Do the following to set the addresses and enable the alternate installation device:.

> **Attention:** If you are going to use this command on the primary partition of an LPAR configuration, be sure to power off all secondary partitions before running the command.

1. Use the control panel to set the mode to Manual. Then perform an IPL using the command:

   ```
   PWRDWNSYS OPTION(*IMMED)RESTART(*YES)IPLSRC(B)
   ```

2. When the IPL or Install the System display appears, select option 3 (Use Dedicated Service Tools (DST)) and press the Enter key.

The Dedicated Service Tools (DST) Sign On display appears.

3. Sign on using the QSECOFR user profile as shown in Figure 10-46.

```
      Dedicated Service Tools (DST) Sign On
                                                           System:   S101880D
Type choices, press Enter.

  DST user . . . . . . . . . . . .   qsecofr
  DST password . . . . . . . . . .









F3=Exit   F12=Cancel
```

Figure 10-46 Sign on to DST

The Use Dedicated Service Tools (DST) menu appears. From the Use Dedicated Service Tools (DST) menu, do the following:

4. Select option 5 Work with DST environment as shown in Figure 10-47.

```
      Use Dedicated Service Tools (DST)
                                                           System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Work with Licensed Internal Code
     4. Work with disk units
     5. Work with DST environment
     6. Select DST console mode
     7. Start a service tool
     8. Perform automatic installation of the operating system
     9. Work with save storage and restore storage
    10. Work with remote service support
    11. Work with system partitions


Selection
     5
```

Figure 10-47 Use DST

5. Select option 2 System devices as shown in Figure 10-48

```
       Work with DST Environment
                                                            System:   S101880D
 Select one of the following:

      1. Active service tools
      2. System devices
      3. DSP user profiles
      4. System values









 Selection
      2
```

Figure 10-48 Work with DST Environment

6. Select option 5 Alternate installation device as shown in Figure 10-49.

```
       System Devices
                                                            System:   S101880D
 Select one of the following:

      1. Printers
      2. Tape devices

      4. Optical devices
      5. Alternate installation device
      6. Console mode







 Selection
      5
```

Figure 10-49 System devices men

7. From the Select Alternate Installation Device display, type a 5 (Display details) next to the resource you want and press the Enter key as shown in Figure 10-50.

```
      Select Alternate Installation Device
                                                           System:   S101880D
 Type option, press Enter.
   1=Select   5=Display details

             Resource                        Serial
 Option      Name          Type     Model    Number
   _         TAP01         6380     001      00-08080
   5         TAP02         3590     E11      00-09090







F2=Deselect device   F3=Exit   F5=Refresh   F12=Cancel
```

Figure 10-50 Select the alternate installation device

The Display Device Details display appears as shown in Figure 10-51.

```
      Display Device Details
                                                           System:   S101880D

 Resource                           Serial
 Name          Type      Model      Number
 TAP02         3590      E11        00-09090

 Physical location:
   Location text . . . . . . :
   Frame ID  . . . . . . . . :
   Card slot . . . . . . . . :

 Logical address:
  SPD bus:
   System bus  . . . . . . . :    0002
   System board  . . . . . . :    0000
   System card . . . . . . . :    0010

  SPD bus:
   System bus  . . . . . . . :    0000
   System board  . . . . . . :    0007
   System card . . . . . . . :    0000

F3=Exit   F12=Cancel
```

Figure 10-51 Alternate installation device details

You need to have a record of the addresses assigned to the alternate installation device selected for installing and recovering your system. Record the following information from the panel in Figure 10-51:

Type/Model: __________
System bus: __________
System card: __________

Note:

a. You may want to repeat this process to record the addresses for all possible alternate installation devices. You should store this information in a safe place, such as the location where your recovery information and recovery media are stored.

b. If more than one alternate installation device is defined, only one can be enabled.

c. The alternate installation device is actually the first device that is found on an SPD bus that contains installation media. You should ensure that only one device contains valid installation media at the time of the D-IPL. This prevents you from loading the wrong version of the Licensed Internal Code.

8. Press the Enter key.
9. The Select Alternate Installation Device display appears. Type a 1 (Select) next to the resource you want and press the Enter key.
10. You should see the following message at the bottom of the display:

    Alternate installation device updated

11. Press F3 (Exit) to return to the Use Dedicated Service Tools (DST) display.
12. Press F3 (Exit) again. The Exit Dedicated Service Tools (DST) display appears.
13. Type a 1 (Exit Dedicated Service Tools (DST)) and press the Enter key as shown in Figure 10-52.

```
     Exit Dedicated Service Tools
                                                          System:   S101880D
 Select one of the following:

      1. Exit Dedicated Service Tools(DST)
      2. Resume Dedicated Service Tools










 Selection
      1
```

Figure 10-52 Exit Dedicated Service Tools

The next display you see is the IPL or Install the System display.

14. Type a 1 (Perform an IPL) and press the Enter key to complete the procedure as shown in Figure 10-53.

```
         IPL or Install the System
                                                             System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Use Dedicated Service Tools (DST)
     4. Perform automatic installation of the operating system
     5. Save Licensed Internal Code



Selection
     1

Licensed Internal Code - Property of IBM 5769-999 Licensed
Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
rights reserved.  US Government Users Restricted Rights -
Use duplication or disclosure restricted by GSA ADP schedule
Contract with IBM Corp.
```

Figure 10-53 Continue with the IPL

10.8 Migrating to Fibre Channel tape

The upgrade path from SCSI to fibre varies depending on the tape technology you currently have in use.

- The IBM 3490, 3570, 3575, 8mm and QIC drives do not have Fibre Channel counterparts. Customers wanting to upgrade to fibre from these drive technologies will need to purchase new drives that do support Fibre Channel.
- The 3590B and 3590E SCSI tape drives can be upgraded to 3590E fibre tape drives. If upgrading from a 3590B, it is cost-effective to go directly to the 3590E fibre drive, rather than doing a two-step upgrade from 3590B SCSI to 3590E SCSI, then later upgrading to 3590E fibre.
- The 3583 SCSI library can be upgraded to a 3583 fibre library by purchasing the integrated SAN Data Gateway feature (fc 8005), and doing a simultaneous one-time conversion of the drive sleds from HVD (fc 8004) to LVD (fc 8003).
- On the 3584 there is no upgrade path to convert SCSI drives to fibre drives. However, new fibre drives can be purchased and added to an existing 3584 library.

The procedure to migrate from SCSI tape to Fibre Channel tape is the same as that used to add new Fibre Channel tape drives. Details can be found in 10.7, “Implementing Fibre Channel tape” on page 285. The one challenge will be finding a suitable window to upgrade to the new technology, since the upgrade may preclude you from using your old drives, whereas a new fibre drive would often leave your old drives available for a more staged cut-over.

Part 5

Advanced topics

In Part 5 we discuss the more complex topics of remote load source mirroring, Peer-to-Peer-Remote Copy (PPRC), FlashCopy and the use of BRMS with external disk.

11

Load Source Mirroring in the ESS

The Load Source Unit (LSU) is a special disk in the iSeries. This section relates to the implementation of remote load source mirror function. This function allows the iSeries internal load source to be mirrored to an ESS unprotected LUN. This provides a degree of recoverability should a disaster befall the iSeries or AS/400 server.

This is the disk that is used to IPL the system (among other things), it is similar to a boot drive. All other “user data” can be located on external DASD units, but the LSU must be an internal drive. This is because the system cannot be IPLed from an I/O adapter (IOA) supporting external drives.

Due to the nature of iSeries Single Level Storage, it is necessary to consider the LSU as a special case. On other platforms, such as UNIX or NT, each disk volume can be identified with its contents. The iSeries is different as all storage is considered as a single large address space. The LSU is within this address space. Therefore, if you want to use facilities such as Peer to Peer Remote Copy (PPRC) or FlashCopy to do a hardware copy of the disks attached to the iSeries, *you must mirror the LSU from* the internal drive into the Enterprise Storage Server to ensure the whole single level storage is copied.

It is advisable if using PPRC and/or FlashCopy to have no other internal disks apart from the LSU as these will not be copied by PPRC or FlashCopy negating the remote copy due to all single level storage not being copied.

11.1 Mirrored internal DASD support

Support has been added to allow internal DASD, such as the LSU, to be mirrored to external DASD. This requires that the external DASD reports in as unprotected, even though in practice, it may actually be protected in a RAID-5 array within the Enterprise Storage Server. This facility for SCSI attached external DASD was added with the following PTFs:

- #6501 IOP PTFs:
  - v4r3: MF24552
  - v4r4: MF24558
  - v4r5: MF24559
- SLIC PTFs:
  - V4R3: MF24930
  - V4R4: MF24931
  - V4R5: MF24929
- OS/400 PTFs:
  - VRM430 SF57394
  - VRM440 SF57395
  - VRM450 SF62704

OS/400 V5R1 and later releases have this support included in the base SLIC and operating system.

Before implementing remote load source mirror, you should check to see if these PTFs have been superseded. In addition to these PTFs, the Enterprise Storage Server must be at the GA driver level or higher. To use this support with PPRC/SCSI attach, reference the ESS support site for more information regarding driver level support.

Table 11-1 shows the protected and unprotected model numbers for the possible Logical Unit Number (LUN) sizes on the Enterprise Storage Server.

Table 11-1 LUN sizes and equivalent model numbers

Size (GB)	Interface	Type	Protected	Un-protected	LBA (Decimal)
4.19	SCSI-3	9337	48C	48A	8,192,000
8.59	SCSI-3	9337	59C	59A	16,777,216
17.55	SCSI-3	9337	5AC	5AA	34,275,328
35.17	SCSI-3	9337	5CC	5CA	68,681,728
36.00	SCSI-3	9337	5BC	5BA	70,320,128
8.59	SCSI-FCP	2105	A01	A81	16,777,216
17.55	SCSI-FCP	2105	A02	A82	34,275,328
35.17	SCSI-FCP	2105	A05	A85	68,681,728
36.00	SCSI-FCP	2105	A03	A83	70,320,128
70.56	SCSI-FCP	2105	A04	A84	137,822,208

When using remote load source mirroring, normal OS/400 rules for mirroring apply. Both the source and target disk drives must be the same size, although they can be of different drive types and spin speeds. It is simply the capacity which must match. However, you should recognize that using remote load source mirroring in this manner is not the same as when mirroring to internal drives. Although it uses the same functions, the primary reason for using remote load source mirroring is to get a copy of the LSU into the Enterprise Storage Server so that the entire DASD space in single level storage can be duplicated by the hardware facilities, such as PPRC and FlashCopy, provided by the Enterprise Storage Server.

In our environment, we are not using remote DASD mirroring, just remote load source mirroring.

> **Note:** In a configuration with remote load source mirroring, the system can only IPL from, or perform a main storage dump to, the internal LSU. If the internal LSU fails, the system can continue to run on the other disk unit of the LSU mirrored pair but the system will not be able to IPL or perform a main storage dump until the LSU is repaired and usable. Therefore, if the system were to crash after the internal LSU failed, the system would not be able to use the main storage dump to reduce recovery time nor to diagnose the problem which caused the system to end abnormally. This is not a restriction due to the Enterprise Storage Server; rather it is due to OS/400 support for remote load source mirroring.

11.2 Setting up Remote Load Source Mirroring

To allow the LSU to be mirrored, the target disk must be unprotected, as OS/400 does not allow you to mirror any disk to another which is already protected. This must be done when you first define the LUN on the Enterprise Storage Server. Once a LUN has been specified, you cannot change the designated protection. Normally, you would only define the LUN for the LSU mirror as being unprotected. All other LUNs would be defined as protected, reflecting their true status to OS/400.

When using Remote Load Source Mirroring, we need to have one LUN defined as "unprotected" so that OS/400 can start mirroring to this LUN. To do this, select "Unprotected" as shown in Figure 11-1.

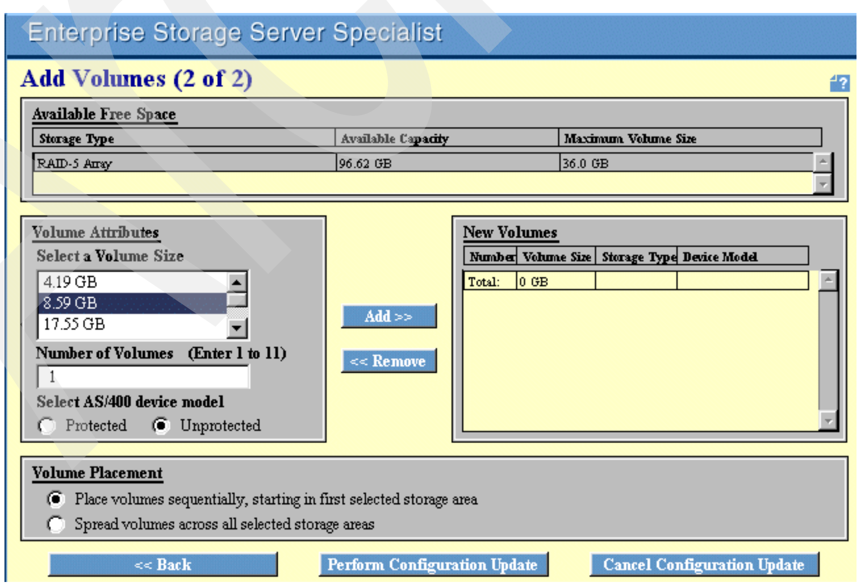

Figure 11-1 Defining an unprotected LUN

Once the unprotected LUN is defined, the ONLY place you can see this in the tabular view on StorWatch Specialist. It is a good idea to print the tabular view for the system you are setting. It is critical that you remember the serial number of the unprotected LUN. Other LUN do not matter as SLIC does not care about the LUNs, they just get assigned to single level storage.

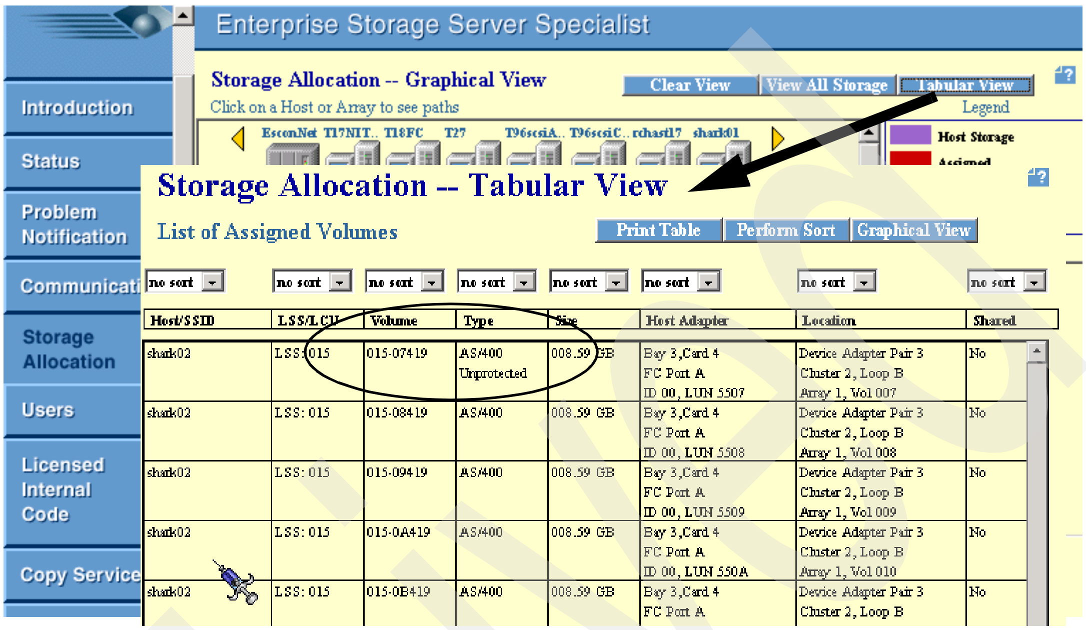

Figure 11-2 StorWatch Tabular view showing 'Unprotected type LUN'

Once the LUN to be set up as the remote LSU pair is created, this (and any other LUNs) will be identified by SLIC and displayed under non-configured units in DST/SST.

You must then use Dedicated Service Tools (DST) to setup load source mirroring on the iSeries.

1. From the DST menu, select Option 4, Work with disk units as shown in Figure 11-3.

```
                 Use Dedicated Service Tools (DST)
                                                         System:   S101880D
 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system
      9. Work with save storage and restore storage
     10. Work with remote service support


 Selection
      4

 F3=Exit   F12=Cancel
```

Figure 11-3 Using Dedicated Service Tools

2. From the Work with Disk Units menu (Figure 11-4), select Option 1, Work with disk configuration.

```
          Work with Disk Units

 Select one of the following:

      1. Work with disk configuration
      2. Work with disk unit recovery





 Selection
      1

 F3=Exit         F12=Cancel
```

Figure 11-4 Working with dIsk units

3. From the Work with Disk Configuration menu (Figure 11-5), select Option 4, Work with mirrored protection.

```
         Work with Disk Configuration

 Select one of the following:

      1. Display disk configuration
      2. Work with ASP threshold
      3. Work with ASP configuration
      4. Work with mirrored protection
      5. Work with device parity protection
      6. Work with disk compression




Selection
      4

 F3=Exit         F12=Cancel
```

Figure 11-5 Work with disk configuration

4. From the Work with mirrored protection menu, select Option 4, Enable remote load source mirroring (Figure 11-6). This does not actually perform the remote load source mirroring but tells the system that you want to mirror the load source when mirroring is started.

```
        Work with mirrored protection

 Select one of the following:

      1. Display disk configuration
      2. Start mirrored protection
      3. Stop mirrored protection
      4. Enable remote load source mirroring
      5. Disable remote load source mirroring




 Selection
      4

 F3=Exit        F12=Cancel
```

Figure 11-6 Setting up remote load source miring

You will see the Enable Remote Load Source Mirroring confirmation screen (see Figure 11-7).

```
      Enable Remote Load Source Mirroring

 Remote load source mirroring will allow you to place the two
units that make up a mirrored load source disk unit (unit 1) on
two different IOPs. This may allow for higher availability
if there is a failure on the multifunction IOP.


Note: When there is only one load source disk unit attached to
the multifunction IOP, the system will not be able to IPL if
that unit should fail.


This function will not start mirrored protection.

Press Enter to enable remote load source mirroring.
```

Figure 11-7 Enable load source mirroring confirmation

5. Press Enter to confirm that you want to enable remote load source mirroring.
6. The Work with mirrored protection screen will be displayed with a message at the bottom, indicating that remote load source mirroring has been enabled (Figure 11-8). You must then select Option `2` to start mirrored protection of the LSU.

```
         Work with mirrored protection

 Select one of the following:

      1. Display disk configuration
      2. Start mirrored protection
      3. Stop mirrored protection
      4. Enable remote load source mirroring
      5. Disable remote load source mirroring




 Selection
      2

 F3=Exit         F12=Cancel
   Remote load source mirroring enabled successfully.
```

Figure 11-8 Confirmation that remote load source mirroring is enabled

7. On the Work with mirrored protection menu, select Option 1, Display disk configuration then Option `1`, Display disk configuration status.

```
         Display Disk Configuration Status

          Serial               Resource
ASP  Unit  Number      Type Model Name        Status
  1                                           Mirrored
        1  68-09011    6717  050  DD023       Unprotected
        2  75-1044000  9337  59A  DD037       Unprotected
        3  75-2044000  9337  59C  DD038       DPY/Active
        4  75-3044000  9337  59C  DD039       DPY/Active
        5  75-4044000  9337  59C  DD040       DPY/Active
        6  75-5044000  9337  59C  DD041       DPY/Active
        7  75-6044000  9337  59C  DD030       DPY/Active
        8  75-7044000  9337  59C  DD031       DPY/Active






 Press Enter to continue.

 F3=Exit         F5-Refresh       F9-Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 11-9 Unprotected LSU ready to start remote load source mirroring

Figure 11-9 shows the internal 6717-050 LSU as disk serial number 68-09011 and our target LUN in the Enterprise Storage Server is shown as a 9337-59A, serial number 75-1044000. Unprotected Fibre Channel attached LUNs would be shown as 2105-A8x, depending on LUN size.

8. When the remote load source mirroring task is finished, you must IPL the system to mirror the data from the source unit to the target. This is done during the database recovery phase of the IPL.

Starting the mirroring process can be a long running job, depending on how many disk units are being mirrored.

11.2.1 Recovering the Remote Load Source

In order to use the disks in the Enterprise Storage Server on another system, (for example accessing the remote copies created by PPRC and FlashCopy), you must first recover the load source unit back onto an internal disk in the iSeries. We will look in more detail later at how PPRC and FlashCopy are implemented.

To recover the remote load source unit, it is necessary to perform a D-IPL and do a scratch install.

Performing a D-IPL and reloading SLIC

We advise that you do this from tape (either a *SAVSYS, a distribution tape or a tape used to save the LIC in DST) rather than use the distribution CD as that does not have any PTFs loaded. Load the tape (or CD) into the alternate IPL device and make sure the device is ready before you perform the D-IPL. If you are not sure how to perform a D-IPL, refer to the Backup Recovery book.

1. The Select a Language Group screen may appear. If so, select your country code (for example 2924 for US English) and press Enter.
2. When the Install Licensed Internal Code menu is displayed, `select Option 1,` Install Licensed Internal Code.

```
          Install Licensed Internal Code
                                                              System:   S101880D
 Select one of the following:

      1. Install Licensed Internal Code
      2. Work with Dedicated Service Tools (DST)
      3. Define alternate installation device




 Selection

      1
 Licensed Internal Code - Property of IBM 5769-999 Licensed
 Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
 rights reserved.  US Government Users Restricted Rights -
 Use duplication or disclosure restricted by GSA ADP schedule
 Contract with IBM Corp.

 Tape unit available
```

Figure 11-10 Begin installing LIC

3. At this stage, you may be prompted to specify a device to load LIC from. However, this screen may not always appear. Select the tape drive with Option 1 and press Enter.

```
        Work with Tape Devices
                                                              System:   S101880D
Type option, press Enter
  1= Select  2=Deselct    5=Display details

            Resource                            Serial
Option      Name            Type      Model     Number         Selected
  1         TAP01           6380      001       00-3088792






F3=Exit F5=Refresh    F12=Cancel
Make the device ready before making a selection
```

Figure 11-11 Selecting a tape drive to load the LIC

4. Because we are planning to recover the LSU from the remote mirror copy, we want to re-initialize the LSU before installing the licensed internal code. To do this, select Option `2`, Install Licensed Internal Code and Initialize system.

```
      Install Licensed Internal Code (LIC)

Disk selected to write the Licensed Internal Code to:
    Serial Number     Type     Model     I/O Bus      Controller     Device
     00-CA71E         6713      030         0              1            0

Select one of the following:

     1. Restore Licensed Internal Code
     2. Install Licensed Internal Code and Initialize system
     3. Install Licensed Internal Code and Recover Configuration
     4. Install Licensed Internal Code and Restore Disk Unit Data
     5. Install Licensed Internal Code and Upgrade Load Source



Selection
     2

F3=Exit           F12=Cancel
```

Figure 11-12 Initialize Load Source Unit and install LIC

In the case of using ESS Copy Services (PPRC or FlashCopy), we will be recovering the “remote mirrored” DASD to a different system. You should see a warning screen similar to that shown in Figure 11-13. However, the actual text of the warning may differ, depending on your particular circumstances (for example, LPAR, first or subsequent recovery of remote load source).

5. Press Enter to continue.

```
      Install Licensed Internal Code - Warning

 Disk selected to write the Licensed Internal Code to:
     Serial Number     Type     Model     I/O Bus      Controller     Device
      00-CA71E         6713      030         0             1            0

 Warning:
     The load source disk and its mirrored pair could not be found
     (see disk information below).


 Missing load source disk:
     Serial Number     Type     Model     I/O Bus      Controller     Device
      75-2040000       9337      59A         1             1            1


 Press Enter to continue the restore or install on the selected
 disk.

 F3=Exit           F12=Cancel
```

Figure 11-13 Missing Load Source Unit

6. You may see the following confirmation screen as in Figure 11-14. If so, press F10 to continue the install process.

```
      Install LIC and Initialize System - Confirmation

 Warning:
     All data on this system will be destroyed and the Licensed
     Internal Code will be written to the selected disk if you
     choose to continue the initialize and install.

     Return to the install selection screen and choose one of the
     other options if you want to perform some type of recovery
     after the install of the Licensed Internal Code is complete.



     Press F10 to continue the install.
     Press F12 (Cancel) to return to the previous screen.
     Press F3 (Exit) to return to the install selection screen.



F3=Exit           F12=Cancel
```

Figure 11-14 Confirm initialization of LSU and install LIC

The new Load Source Unit is being initialized. The screen in Figure 11-15 shows an estimate of the time to initialize the new LSU and the progress.

```
        Initialize the Disk - Status

 The load source disk is being initialized.


 Estimated time to initialize in minutes :    26

 Elapsed time in minutes . . . . . . . . :    14.5




 Please wait.


 Wait for next display or press F16 for DST main menu
```

Figure 11-15 Initialize LST status

When the LSU has been initialized, you will then be presented with the Install Licensed Internal Code - Status screen, as shown in Figure 11-16.

```
 Install Licensed Internal Code - Status

 Install of the licensed internal code in progress.

             +---------------------------------------------------+
 Percent     |                      20%                          |
 complete    +---------------------------------------------------+


 Elapsed time in minutes  . . . . . . . :       6.5




 Please wait.


 Wait for next display or press F16 for DST main menu
```

Figure 11-16 Install Licensed Internal Code Status

7. After the LIC has been installed, the system will begin to IPL.

If you have a non-LPAR system, you will be notified that there is a different disk configuration (Figure 11-17).

However, if you have an LPAR system, you will be warned that there is invalid LPAR configuration information on a non-reporting disk. Go to, “Recovery on an LPAR system” on page 316 to follow the steps to recover from this.

Recovery on a non-LPAR system

Attention: *Do not* press F10 to accept the error. You should check the detail of the error.

1. Select Option 5, Display Detailed Report.

```
         Disk Configuration Attention Report

Type option, press Enter.
  5=Display Detailed Report

Press F10 to accept all the following problems and continue.
The system will attempt to correct them.

Opt  Problem
 5   New disk configuration





F3=Exit   F10=Accept the problems and continue   F12=Cancel
```

Figure 11-17 New Disk Configuration Attention Report

There are a number of possible causes for the problem. These are shown in the New Disk Configuration report (Figure 11-18 and Figure 11-19).

```
         Accept New Disk Configuration

The current configuration indicates a single unit system.  You
can choose to accept it or do one of the following:

Following are the possible causes and recovery procedures:

  o  You can define a new configuration by adding units.

  o  Press F3 to exit to Dedicated Service Tools (DST) and if
     necessary, take the right option to get to
     'Use Dedicated Service Tools (DST)' display.
     On the 'Use Dedicated Service Tools (DST)' display,
      -  Select option 4, Work with disk units.
      -  Select option 1, Work with disk configuration.
      -  Select option 3, Work with ASP configuration.
      -  Select option 3, Add units to ASPs.

  o  If you are performing 'Recover mirror load source'
     utility, press F3 to exit to Dedicated Service Tools (DST).
     If necessary, take the right option to get to 'Use
                                                          More...
F3=Exit                                        F12=Cancel
```

Figure 11-18 New Disk Configuration Detailed Report (part i)

```
          Accept New Disk Configuration

     Dedicated Service Tools (DST)' display.
     On the 'Use Dedicated Service Tools (DST)' display
      -  Select option 4, Work with disk units.
      -  Select option 2, Work with disk unit recovery.
      -  Select option 16, Recover mirrored load source.

 Press Enter to accept the current configuration and continue.






                                                                        Bottom
 F3=Exit                                        F12=Cancel
```

Figure 11-19 New Disk Configuration Detailed Report (part ii)

Attention: We are performing a "Recover mirror load source" task. The report indicates that you should press F3 to Exit from the New Disk Configuration screen and return to DST. *Do not* press F10 to accept the new configuration.

2. Instead, select F3 to Exit from the detailed report screen as shown in Figure 11-20.

```
          Disk Configuration Attention Report

 Type option, press Enter.
   5=Display Detailed Report

 Press F10 to accept all the following problems and continue.
 The system will attempt to correct them.

 Opt  Problem
      New disk configuration






 F3=Exit   F10=Accept the problems and continue   F12=Cancel
```

Figure 11-20 Press F3 to Exit

3. On the IPL or Install the System screen, select Option `3`, use Dedicated Service Tools (DST).

```
         IPL or Install the System
                                                            System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Use Dedicated Service Tools (DST)
     4. Perform automatic installation of the operating system
     5. Save Licensed Internal Code




Selection
     3

Licensed Internal Code - Property of IBM 5769-999 Licensed
Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
rights reserved.  US Government Users Restricted Rights -
Use duplication or disclosure restricted by GSA ADP schedule
Contract with IBM Corp.
```

Figure 11-21 IPL or Install the System

4. Sign on to DST using an acceptable user ID and password.

> **Tip:** If you are using a distribution or SAVSYS tape, the password will be as it was when the tape was created. If you are using the distribution media, the default password for qsecofr is QSECOFR (upper case).

```
      Dedicated Service Tools (DST) Sign On
                                                            System:   S101880D
Type choices, press Enter.

  DST user . . . . . . . . . . . .   qsecofr
  DST password . . . . . . . . . .











F3=Exit   F12=Cancel
Cannot access user profiles - defaults are in effect
```

Figure 11-22 Sign on to DST

You have now completed the tasks specific to a non-LPAR environment. Go to, “Continue to the recover the load source” on page 321.

Recovery on an LPAR system

Use the following steps to correct LPAR configuration errors found when recovering your remote load source.

You will see a Disk Configuration Error Report as seen in Figure 11-23.

1. You may choose Option 5 to view the error report or simply select F12=Cancel to return to the IPL or Install the System menu.

```
         Disk Configuration Error Report

Type option, press Enter.
  5=Display Detailed Report

Opt  Error
     Unit has incorrect logical partition configuration







F3=Exit       F12=Cancel
```

Figure 11-23 Incorrect LPAR configuration warning

We need to work with Dedicated Service Tools to modify the LPAR configuration details.

2. Select Option 3=Use Dedicated Service Tools (DST).

```
         IPL or Install the System
                                                           System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Use Dedicated Service Tools (DST)
     4. Perform automatic installation of the operating system
     5. Save Licensed Internal Code



Selection
     3

Licensed Internal Code - Property of IBM 5769-999 Licensed
Internal Code (c) Copyright IBM Corp.  1980, 1999.  All
rights reserved.  US Government Users Restricted Rights -
Use duplication or disclosure restricted by GSA ADP schedule
Contract with IBM Corp.d
```

Figure 11-24 IPL or Install the System

3. Sign on to DST using an acceptable user ID and password.

Tip: If you are using a distribution or *SAVSYS tape, the password will be as it was when the tape was created. If you are using the distribution media, the default password for qsecofr is QSECOFR (upper case).

```
      Dedicated Service Tools (DST) Sign On
                                                          System:   S101880D
 Type choices, press Enter.

   DST user . . . . . . . . . . . .   qsecofr
   DST password . . . . . . . . . .










 F3=Exit   F12=Cancel
 Cannot access user profiles - defaults are in effect
```

Figure 11-25 Sign on to DST

When the Use Dedicated Service Tools (DST) menu is displayed, we need to discard the LPAR configuration found on the remote load source unit.

4. To do so, select Option 11, Work with system partitions.

```
      Use Dedicated Service Tools (DST)
                                                          System:   S101880D
 Select one of the following:

      1. Perform an IPL
      2. Install the operating system
      3. Work with Licensed Internal Code
      4. Work with disk units
      5. Work with DST environment
      6. Select DST console mode
      7. Start a service tool
      8. Perform automatic installation of the operating system
      9. Work with save storage and restore storage
     10. Work with remote service support
     11. Work with system partitions


 Selection
      11
```

Figure 11-26 Using DST

This will take you to the Work with system partitions menu as shown in Figure 11-27. We will be working with the LPAR configuration data.

5. Select Option 4 to Recover configuration data.

```
      Work with System Partitions
                                                              System:   S101880D
 Attention:  Incorrect of this utility can cause damage
to data in this system.  See service documentation


  Number of partitions  . . . . . . . :  4
  Partition manager release . . . . . :  V5R1M0 L000

  Partition identifier  . . . . . . . :  2
  Partition name  . . . . . . . . . . :  SAMUEL
  Partition release . . . . . . . . . :  V5R1M0 L000

Select one of the following:

     1. Display partition information
     2. Work with partition status

     4. Recover configuration data

 Selection
      4

 F3=Exit   F12=Cancel
 Configuration data errors detected - see Product Activity Log.
```

Figure 11-27 Working with system partitions

In our situation where we are recovering a remote load source unit which was previously copied from another partition, there is LPAR configuration information available on our remote load source unit. This is not valid for the partition we are recovering to so we need to discard it.

6. Select Option 3 to clear the LPAR configuration data from the non-configured remote load source unit.

```
         Recover Configuration Data
                                                                 System:   S101880D
Select one of the following:



     3. Clear non-configured disk unit configuration data

     5. Accept load source disk unit
     6. Copy configuration data to other side






Selection
     3
```

Figure 11-28   Clearing LPAR configuration data

You will be presented with a screen as shown in Figure 11-29, showing the disk unit containing the invalid LPAR configuration data.

7. Select Option 1=Select to identify the disk unit containing the invalid LPAR configuration data.

```
  Select Non-Configured Disk Unit for Configuration Data Clear
                                                           System:   S101880D
Type option, press Enter.
  1=Select

     I/O Resource                         ---Last Updated---  System
Opt  Description          Type-Model  Date      Time      Serial Number
 1   Disk Unit            2105-A81    03/23/01  17:57:53  5GG8M








F3=Exit    F10=Display serial/part numbers   F12=Cancel
```

Figure 11-29   Identify the disk unit containing the invalid LPAR configuration data

You will then be presented with a confirmation screen (Figure 11-30). Check that this is the correct disk unit. Use F10 to display the serial numbers if you are not certain.

8. When you are sure, press Enter to confirm.

```
  Confirm Non-Configured Disk Unit for Configuration Data Clear
                                                     System:   S101880D
Press Enter to confirm your choice to clear the logical
  partitioning configuration data residing on the disk
  units listed below.
Press F12 to return to change your choice.

I/O Resource                    ---Last Updated---  System
Description         Type-Model  Date      Time      Serial Number
Disk Unit           2105-A81    03/23/01  17:57:53  5GG8M







F10=Display serial/part numbers   F12=Cancel
```

Figure 11-30 Configuration data clear

This will return you to the Recover Configuration Screen and show a message Clear non-configured disk unit configuration data successful.

9. F12 back until you get the Use Dedicated Service Tools (DST).

You have now completed the tasks specific to an LPAR environment.

Continue to the recover the load source

The following steps will complete the tasks required to recover the remote load source unit. When the Use Dedicated Service Tools (DST) menu is displayed, you might want to check the current DASD configuration.

1. To do so, select Option `4`, Work with disk units.

```
     Use Dedicated Service Tools (DST)
                                                           System:   S101880D
Select one of the following:

     1. Perform an IPL
     2. Install the operating system
     3. Work with Licensed Internal Code
     4. Work with disk units
     5. Work with DST environment
     6. Select DST console mode
     7. Start a service tool
     8. Perform automatic installation of the operating system
     9. Work with save storage and restore storage
    10. Work with remote service support
    11. Work with system partitions


Selection
     4
```

Figure 11-31 Using DST

2. On the Work with Disk Units screen (Figure 11-32), `select Option 1`, Work with disk configuration.

```
Work with Disk Units

Select one of the following:

     1. Work with disk configuration
     2. Work with disk unit recovery




Selection
     1

F3=Exit         F12=Cancel
```

Figure 11-32 Work with Disk Units

The Work with Disk Configuration menu provides a number of options. We want to check the current disk configuration to ensure that we only have one disk unit available.

3. Select Option 1, Display disk configuration.

```
Work with Disk Configuration

Select one of the following:

     1. Display disk configuration
     2. Work with ASP threshold
     3. Work with ASP configuration
     4. Work with mirrored protection
     5. Work with device parity protection
     6. Work with disk compression




Selection
     1

F3=Exit         F12=Cancel
```

Figure 11-33 Work with Disk Configuration

From here, we want to check the status of the configured disks.

4. Select Option `1`, Display disk configuration status.

```
 Display Disk Configuration

Select one of the following:

      1. Display disk configuration status
      2. Display disk configuration capacity
      3. Display disk configuration protection
      4. Display non-configured units
      5. Display device parity status
      6. Display disk hardware status
      7. Display disk compression status




 Selection
      1

 F3=Exit         F12=Cancel
```

Figure 11-34 Display Disk Configuration

You should only see one configured disk unit (Figure 11-35). If you have more than one unit, you need to go back to DST and remove all the other units apart from the LSU.

```
 Display Disk Configuration Status

             Serial                 Resource
 ASP  Unit   Number      Type Model  Name        Status
   1                                             Unprotected
         1   00-CA71E    6713  030   DD001       Configured






 Press Enter to continue.

 F3=Exit        F5=Refresh         F9=Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 11-35 Display Disk Configuration Status

5. Keep pressing F12 to Cancel until you get back to the Work with Disk Units menu.
6. Select option `2`, Work with disk unit recovery.

```
          Work with Disk Units

Select one of the following:

     1. Work with disk configuration
     2. Work with disk unit recovery







Selection
     2

F3=Exit         F12=Cancel
```

Figure 11-36 Work with Disk Units

7. Scroll to the next page and choose option `16`, Recover mirrored load source.

```
          Work with Disk Unit Recovery

Select one of the following:

    14. Correct device parity protection mismatch
    15. Recover unknown load source
    16. Recover mirrored load source
    17. Recover from start compression failure
    18. Migrate load source disk unit data




                                                                        Bottom


Selection
     16

F3=Exit   F11=Display disk configuration status   F12=Cancel
```

Figure 11-37 Begin to recover mirrored load source

You will be presented with a confirmation screen indicating the disk configuration which will be recovered along with the mirrored load source.

8. Scroll through to check the disk configuration to be recovered.

```
Confirm Recover Mirrored Load Source

To proceed, the system must copy all data from the remote load
source to the new local load source.  Once you confirm on
this screen, the control panel will display the percent of
completion in the 'xx' places of the C6xx 4205 System
Reference Code (SRC).  After the successful completion of this
operation, the system will be IPLed to DST.

Press Enter to confirm the recovery of the mirrored load source.

The system will have the following configuration:

            Serial                  Resource
ASP  Unit   Number      Type Model  Name        Status
  1                                             Mirrored
        1   00-CA71E    6713  030   DD001       Active
        1   75-1044000  9337  59A   DD004       Active
        3   75-2044000  9337  59C   DD005       DPY/Active
        4   75-3044000  9337  5AC   DD006       DPY/Active
                                                                        More...

F12=Cancel
```

Figure 11-38 First Confirm Recover Mirrored Load Source screen

9. Verify the correct number of disk units. Press Enter to start recovery.

```
Confirm Recover Mirrored Load Source

To proceed, the system must copy all data from the remote load
source to the new local load source.  Once you confirm on
this screen, the control panel will display the percent of
completion in the 'xx' places of the C6xx 4205 System
Reference Code (SRC).  After the successful completion of this
operation, the system will be IPLed to DST.

Press Enter to confirm the recovery of the mirrored load source.

The system will have the following configuration:

            Serial                  Resource
ASP  Unit   Number      Type Model  Name        Status
        5   75-4044000  9337  5AC   DD007       DPY/Active
        6   75-5044000  9337  5AC   DD008       DPY/Active
        7   75-6044000  9337  5AC   DD002       DPY/Active
        8   75-7044000  9337  5AC   DD003       DPY/Active

                                                                          Bottom

F12=Cancel
```

Figure 11-39 Last Confirm Recover Mirrored Load Source screen

When the LSU has been successfully recovered, you will be presented with a message as shown in Figure 11-40 and the system will come up to DST from where you can continue to IPL the system and proceed as normal. You may want to check the other non-DASD hardware configuration, as this is likely to be different from the configuration on the source system.

```
Disk Configuration Information Report

The following are informational messages about disk
configuration changes started in the previous IPL.

Information

Recover mirrored load source completed successfully




Press Enter to continue.
```

Figure 11-40 Remote load source recovered successfully

12

Peer-to-Peer Remote Copy

In this chapter, we cover some of the additional features of the Enterprise Storage Server and discuss how they might be used in an iSeries environment.

12.1 Requirements and preliminary setup

In this chapter, we describe the basic hardware and software requirements for the iSeries server to set up Copy Services on an IBM Enterprise Storage Server.

This is followed by a description of how to set up the primary and secondary Copy Services servers.

For convenience, excerpts have been taken from *Implementing ESS Copy Services on UNIX and Windows NT*, SG24-5757, and *ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services*, SC26-7346, relevant to iSeries. For more detailed information please reference those books.

12.1.1 Requirements for Copy Services

The following sections describe the various hardware and software requirements for a iSeries host connected to an ESS to support Copy Services.

Complete and current information on all supported servers (including supported models), is available on the Internet at the following URL:

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

iSeries, AS/400 hardware

ESS supports the following iSeries and AS/400 servers:

- IBM AS/400 models 620, 640, 650, S20, S30, S40, 720, 730, 740 with SCSI host adapter features: FC #6501
- iSeries models 820, 830, 840, need a migration tower for FC #6501. Additional expansion units or FC 6501 must be ordered with RPQ 847 147
- New Model 820 do not support ESS SCSI connection.
- iSeries models 820, 830, 840, 270 with Fibre Channel features: FC #2766
- @server i5 Models 520, 550, 570, 595
- Hub feature: IBM 3534-1RU Fibre Channel managed Storage Hub
- Switch Feature: IBM 2109 S08/S16 only Quick loop support

iSeries, AS/400 software

The following operating system support is required for the ESS on iSeries:

- OS/400 Version 3.2, 4.1 or higher
- OS/400 Version 4.3 or higher with PTF's for *remote load source* support

ESS support for Fibre Channel attach to iSeries requires:

- OS/400 Version 5.1 or higher

Copy Services server

In order to use Copy Services you must configure one ESS to be the primary Copy Services Server. All information related to Copy Services is stored in this ESS such as volumes and their state, ESCON connectivity between ESS and much more.

On each ESS that is configured to use Copy Services there is a client running. Whenever the Copy Services configurations changes this client notifies the Copy Services server of the changes.

Optionally there could be one ESS defined as the backup Copy Services Server. In event the primary Copy Services Server is unavailable or lost, the backup server could be used for controlling the ESS Copy Services. Once the primary ESS is available again the backup server will notify all clients, and the clients switch back all communication to the primary Copy Services Server. It is recommend to have the primary and backup Copy Services Server physically located at different sites, if the primary Copy Services Server becomes unavailable for whatever reason, the backup Copy Services Server will keep Copy Services available.

The information of the primary and backup Copy Services server has to be specified on each cluster of the all Enterprise Storage Servers that are going to be used for Copy Services. We recommend to specify a backup Copy Services whenever it is possible. The primary and secondary Copy Services server however, do not need to be part of any Copy Services operation. The primary Copy Services Server could be any ESS that is connected to the network.

This task is accomplished by an IBM Customer Engineer (CE) via a service terminal connected to the ESS serial interface.

12.2 Peer-to-Peer Remote Copy

PPRC is an established data mirroring technology that has been used for many years in Mainframe environments. It is used primarily to protect an organization's data against disk subsystem loss or complete site failure. In this chapter, we describe how to configure and use PPRC.

12.2.1 Overview

When talking about PPRC, the systems where production applications run from are referred to as the primary or application site. Systems where the recovery or test of applications are performed is referred to as the secondary or recovery site.

PPRC is a synchronous protocol that allows real-time mirroring of data from one Logical Unit (LUN) to another LUN. The LUNs can reside within the same ESS, or in another ESS located at a secondary site some distance away.

PPRC is application independent, the copy functions are occurring at the disk subsystem level, the applications have no need to know of PPRC's existence.

The PPRC protocol guarantees that the secondary copy of data is up-to-date by ensuring that the primary copy will be written only if the primary system receives acknowledgement from the secondary that the data is committed to NVS.

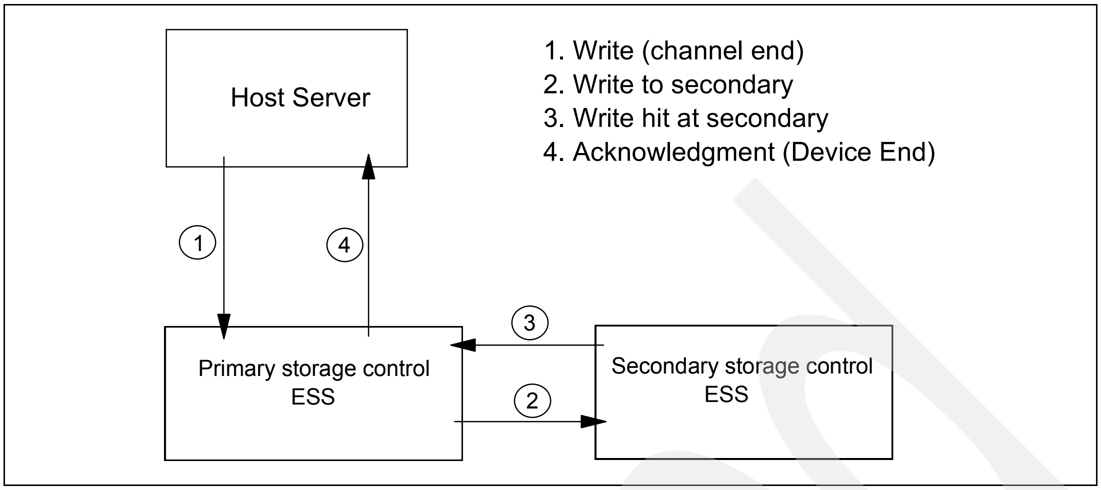

Figure 12-1 Figure 7 PPRC write cycle

1. The host server requests a write I/O to the primary ESS. The write is staged through cache into Non-volatile Storage (NVS).
2. PPRC dispatches the write over an ESCON channel to the secondary ESS. The write hits the secondary ESS's NVS.
3. The primary then receives acknowledgment of the remote write from the secondary.
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 destage to disk. Destage 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 destage 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.

12.2.2 PPRC volume states

Volumes within the Enterprise Storage Server used for PPRC could be found in one of the following states (Figure 8):

Simplex

The simplex state is the initial state of the volumes before used in any PPRC relationship.

Simplex pending

Volumes are in the simplex pending state after the PPRC copy relationship was established but the source and target volume are out of sync. In this case data still needs to be copied from the source to the target volumes. This may be the case after the PPRC relationship was just established, or in a case where PPRC reestablishes after a storage subsystem failure.

Duplex

This is the state of a volume pair that is in sync. That means, both source and target volume containing exactly the same data. Sometimes this state is also referred to full copy mode.

Suspended

Volumes are in the suspended state when the source and target storage subsystems cannot communicate anymore and therefore the PPRC pair could not be kept in sync or when the PPRC pair was suspended manually. During the suspended state, the primary volume's storage server keeps track of all updates to the source volume for reestablishment of the PPRC pair later on.

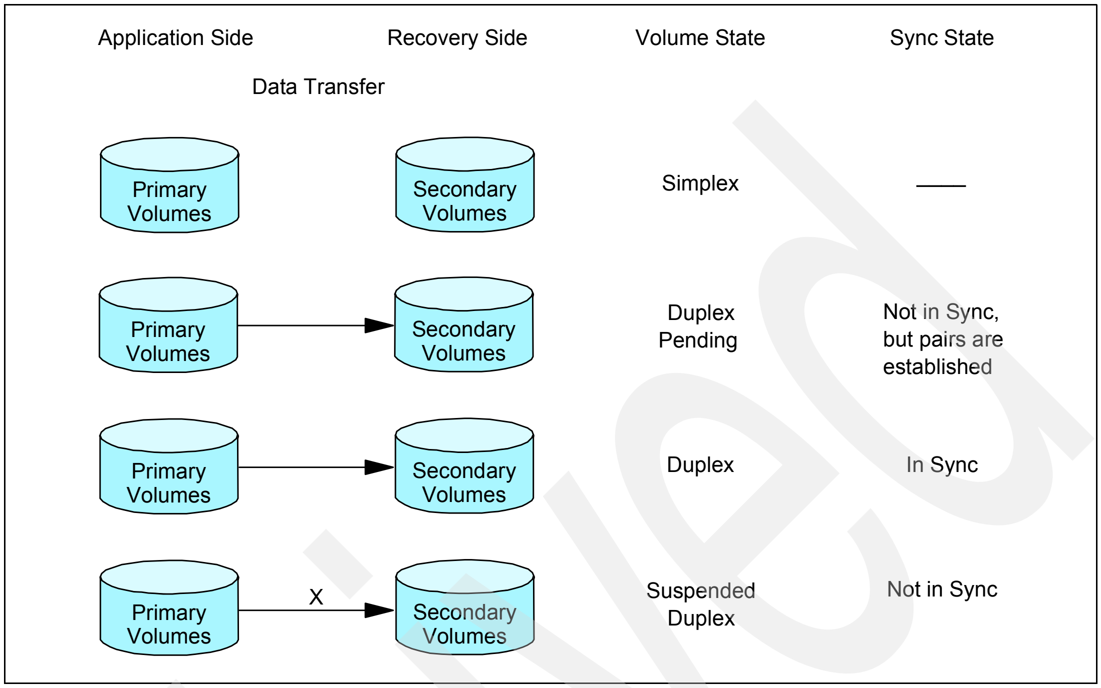

Figure 12-2 PPRC volumes state

12.2.3 Planning for PPRC on ESS

The following sections describe the important areas you should consider when planning for PPRC on an Enterprise Storage Server.

Hardware and Software Requirements

PPRC is possible only between Enterprise Storage Servers. Other disk storage units that support PPRC can also communicate to the same type of unit only.

You need to have the PPRC feature purchased, as well as PPRC capable microcode activated on all ESS units that will be used for PPRC.

PPRC operates at a volume level from one LSS to another LSS. That means you need to have the target volumes available on the secondary ESS and you need to identify the LSSs where the primary and secondary volumes are.

ESCON connections have to be configured between the units. There can be up to eight ESCON links between the subsystems. A primary ESS can communicate to up to four secondary ESS. A secondary ESS can be connected to any number of ESS primary subsystems.

All the ESS units involved in PPRC have to be connected with Ethernet and TCP/IP to the units that are the primary and backup Copy Services servers. All ESS cluster hostnames including its own have to be added to the cluster hostname list during installation.

Browser for ESS Specialist has to be installed on the machines that will be used to control PPRC with the Web interface. See the Implementing the Enterprise Storage Server in Your Environment, SG24-5420 for the browser recommendations.

Figure 12-3 Ethernet and ESCON connection between Copy Services participating ESS.

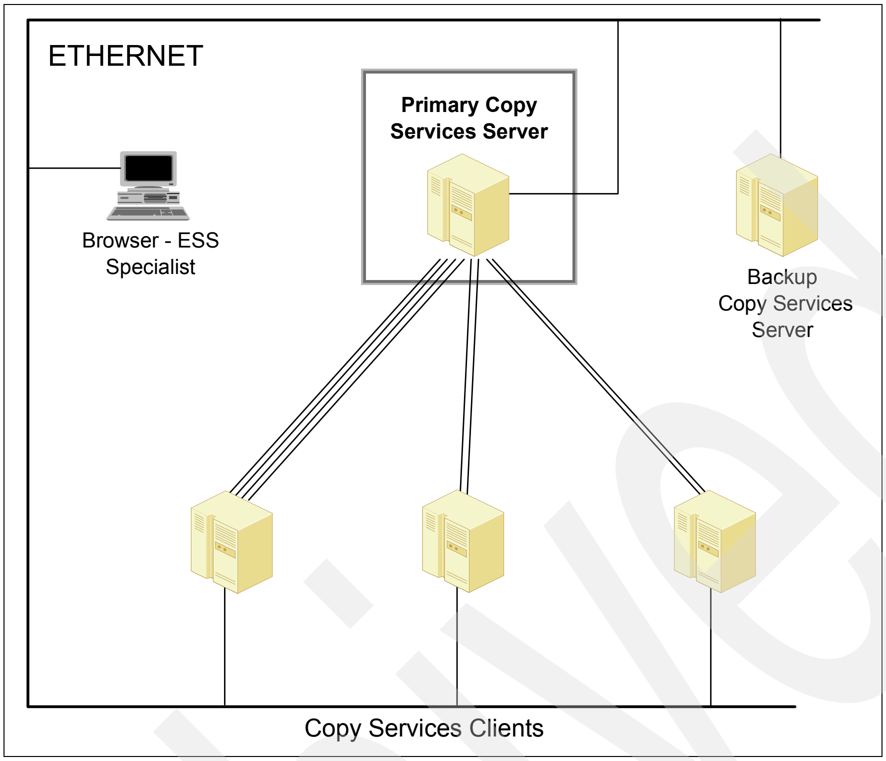

Figure 12-3 Ethernet and ESCON connection between Copy Services participating ESS

12.2.4 Resources planning

Capacity planning Considerations

Proper capacity planning is essential when implementing PPRC on ESS. Having an adequate number of links is crucial to the performance of PPRC. Although a single link may handle about 800 write I/O per second, it is best to plan for about 400 write I/O per second per link. This will allow for short term bursts of I/O will maintaining performance. For sequential writes plan for about 8 MB/s per link. For iSeries, Performance Monitor may be used to estimate write I/O rates and MB/second.

At distance, there will be a linear and predictable performance impact from PPRC. I/O service times should increase at about 1ms per 30 km for each track written. Whether PPRC at a particular distance will have acceptable performance depends on the workload characteristics and performance requirements.

Data Consistency Considerations

In disaster recovery scenarios you may be exposed to lost writes. These are the “in-flight” transactions not committed from memory to NVS, depending on if the host that went down had data in cache not yet sent to the ESS. The data that was transferred to the ESS and confirmed back as written, is in NVS (of the secondary ESS in case of PPRC) and will be committed to disks.

We recommend a full file system check on the target volumes prior to using them in any PPRC situation.

Test Plan / Disaster recovery Plan

A Disaster Recovery Plan is complex, nothing can understate the importance of rehearsals and testing of your environment. Carefully set up your PPRC tasks for establishing and terminating pairs, ensure that they are tested and well documented. Proper and complete configuration and documentation will be the difference between a recovery and data corruption.

We suggest to start using PPRC routinely to become familiar with the way it works and how it affects your systems performance, as well as to test your disaster recovery procedures regularly.

12.2.5 How to configure the ESS for PPRC

In this section, we describe the way to set up an ESS in preparation for PPRC.

ESCON

For PPRC primary to secondary unit (channel to control unit) communication a maximum of eight ESCON links using modified ESCON protocol can be configured. ESCON channels provide 160Mbps 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 ESCON director is used and the host is connected to it.

ESCON links support distances of up to 2 km with 50 micron multimode optical cables and up to 3 km with 62.5 micron multimode cables. The ESCON channel performance is a function of distance.

By using up to two ESCON Directors you can extend these distances. You can use Extended Distance Feature (XDF) ports with singlemode fiber-optic cables, the XDF maximum distance being 20 km.The total maximum distance is 103 km between the two ESS.

Various channel extenders can also be used to increase the distance between ESS servers. IBM 9729 Optical Wavelength Division Multiplexer (MuxMaster) enables 50 km distance between MuxMaster units. ESCON is used to attach ESS to it.

The IBM 2029 Fiber Saver, Dense Wavelength Division Multiplexer (DWDM), will enable up to 50 km distance between Fiber Saver units that are ESCON attached to the primary and secondary ESS.

ESS

PPRC requires logical paths to be established between the primary and the secondary ESS logical control units (or LSS). Each LSS in the primary ESS that will use PPRC requires at least one path to be set to the LSS in the secondary ESS that holds the secondary volumes (Figure 10). Each ESCON link supports 64 logical paths so even with 16 LSS defined you are able to setup a logical path from each LSS to each LSS with only four ESCON PPRC links. We always recommend you to use all available links for each LSS pair used for PPRC. That gives you maximum protection against ESCON link failure.

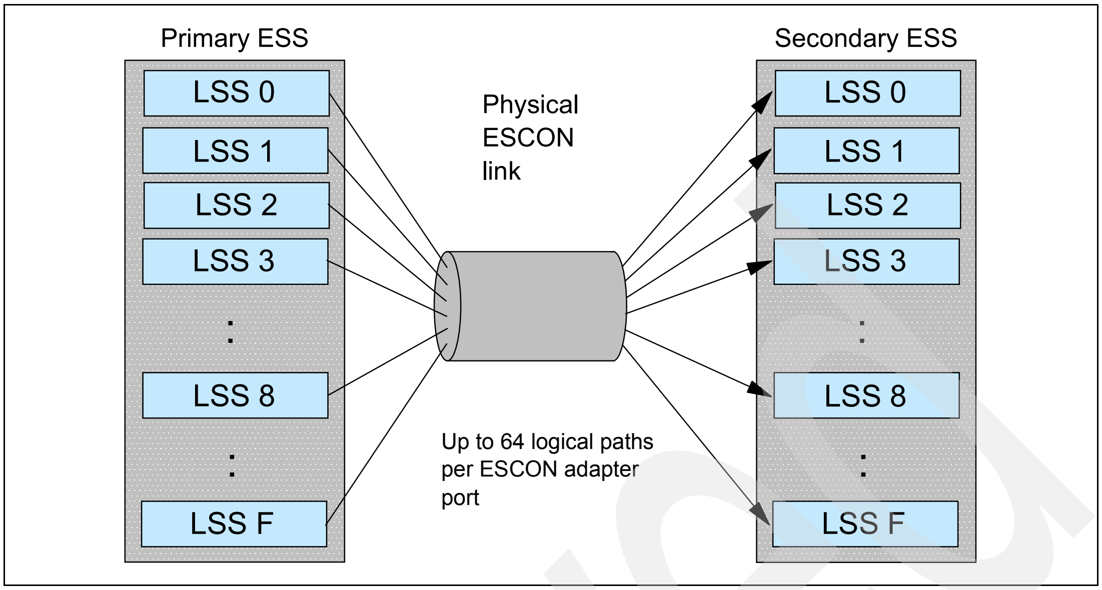

Figure 12-4 ESS logical path

12.2.6 How to set up PPRC

Recommendations

Use StorWatch ESS Specialist Copy Services to perform one-time unique tasks, like paths setup, tasks setup and PPRC invocation for data migration and workload migration. As of this writing the Command Line Interface (CLI) is not supported for iSeries.

Operational considerations

We already know that an ESS can be connected to a variety of host platforms. While configuration and management of Copy Services at the ESS is the same regardless of host platform, each platform's operating system has some differences in the way it presents and manages Copy Services volumes.

The following sections describe a range of general and specific operational considerations relating to iSeries server

Monitoring and Managing Volumes and Paths

PPRC paths, pairs and tasks can be managed by the ESS Specialist Copy Services Web panel.

ESS Specialist Web interface will allow you to manage paths, PPRC volumes and tasks. We recommend that you to setup the paths as well as preset establish, suspend and terminate tasks in advance. Then you will be able to use these tasks from the Web interface without having to configure them each time. Eliminating the potential for inadvertent data corruption do to human operational inconsistencies.

12.2.7 Moving and migrating data with PPRC

Apart from disaster recovery, you can use PPRC for tasks assisting you with data migration and workload migration.

Disk to tape copy from PPRC secondary volumes can be done in order to minimize the impact on the primary system.

1. Verify the PPRC pair status to see if volumes are in duplex
2. Make sure the data on the disk is consistent (Quiesce system, stop I/O)
3. Click on both volumes (source and target) and do a "suspend pair" the target volume is still not accessible
4. Click on the suspended target volume only and do a "terminate pair" from the target LSS which will turn the target volume to the simplex state (the secondary volume is now accessible to the host on the remote site.)
5. Recover the Backup Side (Recover remote load source unit)
6. Copy to tape
7. To re-establish the PPRC pair, click both volumes (first the suspended source, second the simplex target) and establish PPRC copy pair with the ‘copy out-of-sync cylinders only’.
8. **Note**: Use care to ensure that the secondary volume is not corrupted while it is in simplex stat.

To migrate data from one ESS to another you may implement a process similar to the following:

1. Vary secondary volumes offline.
2. Establish PPRC paths.
3. Establish PPRC pairs with COPY option.
4. Query the PPRC pair status to see if they are in duplex.
5. When copy is complete (pairs are in duplex) you can switch to secondary.
6. Stop all write I/O to primary and unmount.
7. Withdraw the PPRC pairs so volumes will return to simplex - you may delete paths if you want to.
8. If volumes are attached to the same host, you have to vary the primary volumes offline.
9. Vary the secondary volumes online, check the file system and mount Then you can start the applications using secondary volumes.

To move workload between hosts or sites and then use the original primary site as the secondary site you should perform the steps for migrating data to the other host and then the following additional actions. This procedure can also be used when you want to make a temporary move, for example when you know that the primary site is going to be unavailable due to maintenance.

1. Establish paths and pairs with NOCOPY in reverse direction.
2. Suspend the pairs.
3. Start applications on secondary volumes.
4. When the old primary site is available you can establish the pair with the RESYNC option.

12.2.8 Performance considerations

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

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

Please keep in mind that the general ESS performance considerations like the volume placement or amount of storage per host adapter still apply when planning for PPRC.

Optimized PPRC communication

There were some modifications made to the ESCON protocol used for PPRC communication of the Enterprise Storage Server. These modifications are in particular:

- A larger frame size which result in less overhead during PPRC communication.
- Less handshaking between the two communicating ESS's which makes transfer of data more efficient. The handshake was reduced from 6 down to 3 exchanges.

Number of ESCON paths between Enterprise Storage Servers

When a host system is sending a write request to a primary volume of a PPRC pair, an I/O complete will be returned to this host once the data is written to the source and target ESS (synchronous copy). This will have some performance impact upon the write I/O operations of the application.

Please make sure that you are using an appropriate number of ESCON paths for PPRC between the source and the target ESS. Increasing the number of the physical ESCON links will increase the maximum overall bandwidth for updating the targets. Using multiple ESCON connections for a PPRC pair (maximum of 8) will improve the response time of an I/O request from the host. Keep in mind that too few paths may result in a bottleneck. It is recommended that there are a minimum of four paths between the primary and secondary ESS.

Placement of the ESCON adapters used for PPRC

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

For example if there are 4 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 connection established over a single physical connections. This will be most likely the case in a real environment.

Please consider that multiple logical paths will share the bandwidth of the ESCON path(s) between each other. If there are critical PPRC pairs we recommend to 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 lower critical PPRC pairs.

Setup of the secondary ESS

For disaster recovery reason you are 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 4 secondary ESS. 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 computed 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 volumes used for PPRC pairs of multiple primaries across all available RAID arrays within the secondary ESS.

12.2.9 ESS Copy Services Web Interface

There are two different methods of using the ESS Copy Services in the Open Systems environment:

- A Web-based Interface (as part of the ESS Specialist)
- A Java-based Command Line Interface (CLI). Currently iSeries doesn't support CLI interface

In this part we explain how to use and setup the ESS Copy Services Web Interface.

Overview and requirements

The ESS Copy Services are running inside the Enterprise Storage Server. One ESS has to be defined as the Copy Services server and holds all Copy Services related information. Optionally there could be a second ESS defined to be the backup Copy Services server. On each ESS that is intended to use Copy Services there is a Copy Services client running who communicates to the server.

Access to the Copy Services is provided through an Internet browser. Using a Web browser gives the possibility to easily control the ESS copy functionality over the network from any platform the browser is supported for.

The ESS Copy Services require one of the following Internet browsers:

- Netscape Communicator 4.6 or above
- Microsoft Internet Explorer (MSIE) 5.0 or above

You enter the ESS Copy Services main menu by selecting the Copy Services button of the ESS Specialist (Figure 12-5). This connects the browser to the ESS that is specified as the Copy Services Server. If you previously have not selected any of the other ESS Specialist buttons you will be prompted for the user name and the password before starting the Copy Services Web screen.

Figure 12-5 Main Menu of the ESS Specialist

The message window shown in Figure 12-6 will be displayed while connecting to the Copy Services Server.

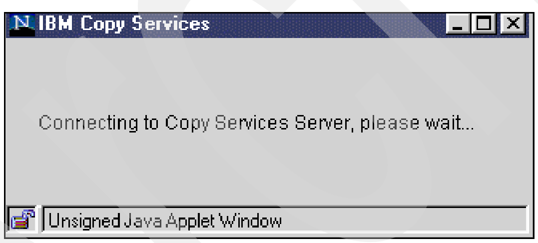

Figure 12-6 Copy services start message

Once the connection to the Copy Services server is successful the main menu of the Copy Services Web interface will be displayed (Figure 12-7). From here you can access to all Copy Services menus by selecting one of the buttons on the left side.

Figure 12-7 Mainmenu of the ESS Copy Services Web Interface

Volumes panel of the ESS Copy Services Web Interface

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 8 digit identifier (ID). This identifier is used to address each volume within the ESS.

From the *Volumes* menu you will be able to:

- Get information and status about a volume defined in a logical subsystem (LSS) of the ESS.
- Select source and target volume for a PPRC or FlashCopy task.
- Filter the output of the volume display to a selected range.
- Search for a specific volume based on its unique volume ID.
- Establish, terminate and suspend PPRC Copy pairs and optionally save the operation as a task.
- Establish and withdraw FlashCopy pairs and optionally save the operation as a task.
- Enter the multiple selection mode for PPRC and FlashCopy

Figure 12-8 shows the entry window of the *Volumes* panel. On the left side you select the source LSS and on the right side the target LSS for your copy task.

The source and target logical subsystem is specified 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 2105:14419:11.

The *Volumes* display displays all volumes defined within the LSS. Below the volume you will find its unique serial number. The color of the volume indicates if it is used in any copy relationship (source or target), or if it not part of a copy pair at all.

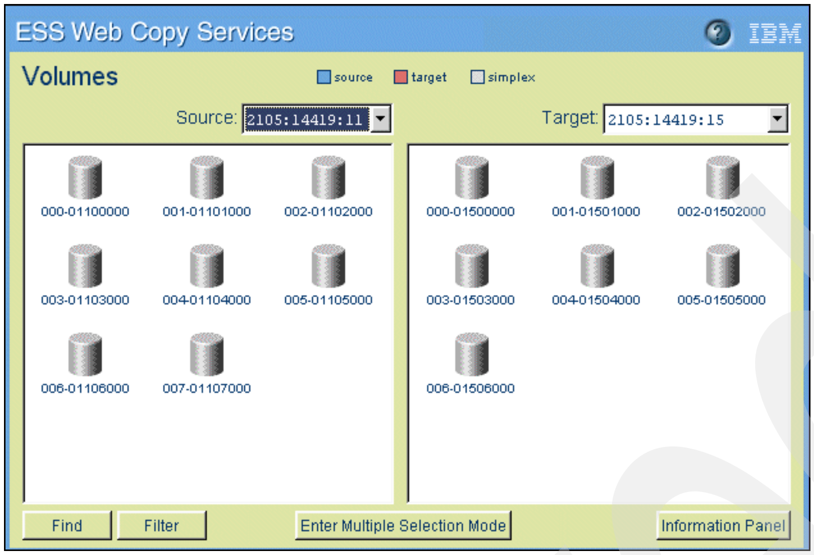

Figure 12-8 Source and target area of the Volumes menu

Note: You cannot display the same logical subsystem in both the source and the target area. Therfore, select two different logical subsystems as source and target, or only one LSS, in either the source or the target area.

Working with volumes

You can get more detailed information about a single volume by selecting the volume and clicking on the *Information Panel* button, as shown in Figure 12-9.

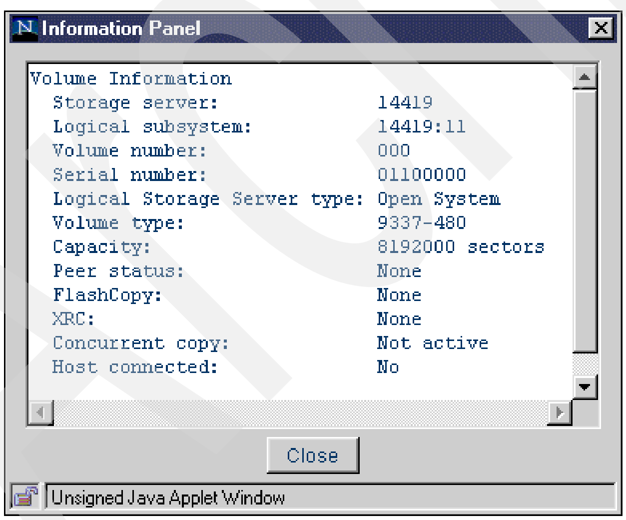

Figure 12-9 Volume Information window

With the *Find* button there is the possibility to search for a specific volume. The volume is specified with its 8 digit ID. See Figure 12-10.

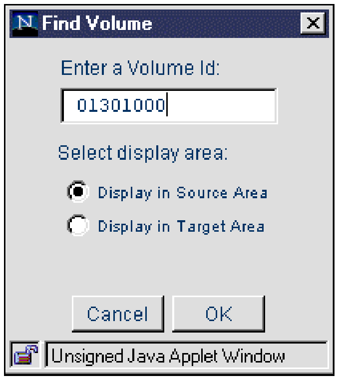

Figure 12-10 Find volume window

In addition, you can filter the output of the volume display to a selected range by clicking the *Filter* button and select the *Filter volumes* option.

In the example showed in Figure 12-11 we want to display Open Systems volumes only. The volumes should to be in the simplex state; that means they are currently not in use in any Copy relationship.

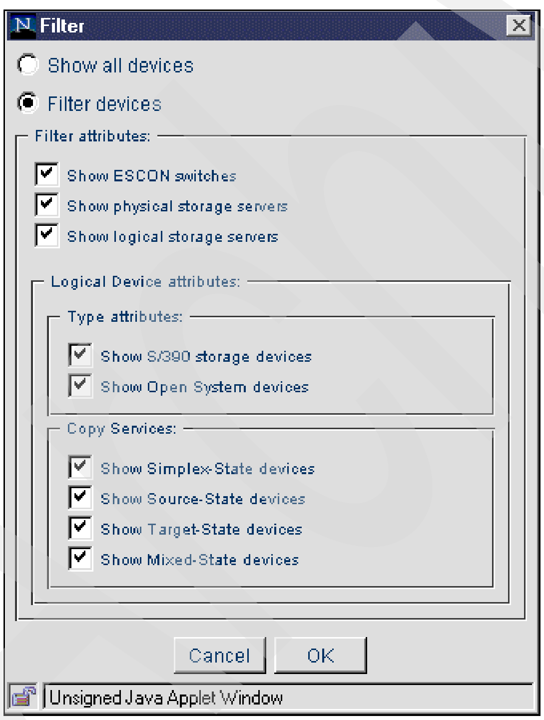

Figure 12-11 Filter volumes window

12.2.10 Storage Servers panel of the ESS Copy Services Web Interface

The *Storage Servers* panel displays the Enterprise Storage Servers and the logical subsystems 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 *Storage Servers* menu you will be able to:

- View all Enterprise Storage Servers within the storage network.
- View all logical Subsystems within the storage network.
- Get information about a logical subsystem and its status.
- View and modify the copy properties of a logical subsystem.
- Filter the output to a selected range.
- Search for a specific logical subsystem based on its unique address.
- Establish, terminate and suspend PPRC Copy pairs and optionally save them as a task.

In Figure 12-12 you will see the Storage Servers output for a selected ESS. The color indicates the state of the LSS, whether it contains volumes that are currently in any copy relationship (source, target, mixed) or not part of a copy pair at all.

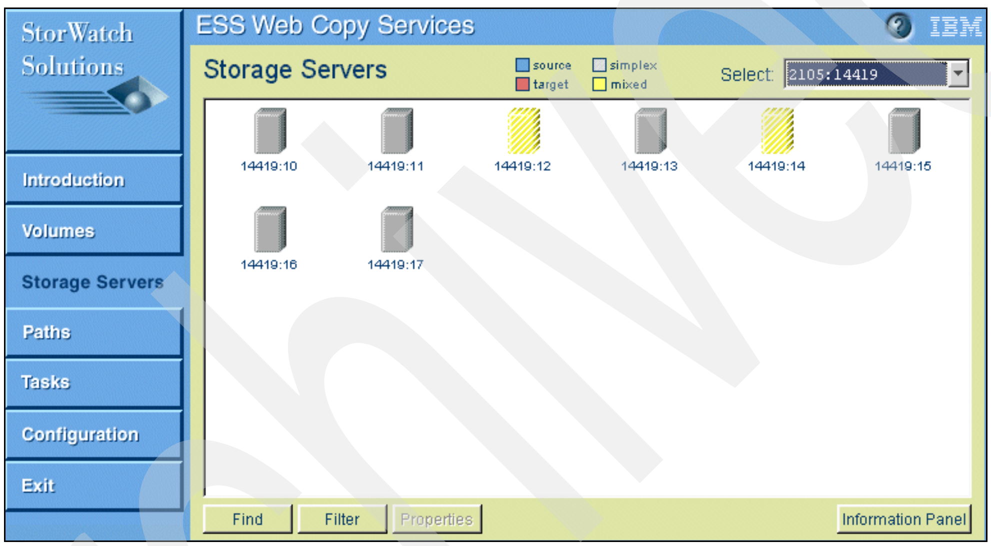

Figure 12-12 Storage Servers window

By selecting one LSS and clicking the *Properties* button you can view or change the copy properties of the entire LSS (see Figure 12-13).

iSeries and AS/400 do not support this function.

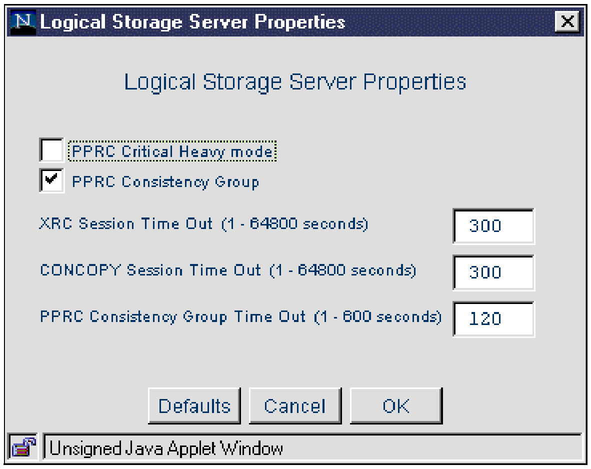

Figure 12-13 LSS properties window

PPRC Critical Heavy mode

This parameter works in conjunction with the *Critical volume mode parameter* when establishing a PPRC pair. It means that updates to the pairs of volumes involved is critical to the operation of a system or database. Look at Table 12-1 for more information.

PPRC Consistency Group

A consistency group is a set of volumes paired for PPRC. When a failure occurs on any device within such a group all I/O to volumes within the same group can be frozen by automation software. A consistency group can include multiple LSS. This option enables that function for the entire logical subsystem.

XRC Session Time Out / CONCOPY Session Time Out

Does not apply to iSeries (AS/400)

PPRC Consistency Time Out

Amount of time that write I/O is held from devices of a consistency group. This time enables automation software to detect that an error has occurred and to issue commands to freeze all other members of the consistency group.

12.2.11 Working with logical subsystems

You can get more detail information about a single logical subsystem by selecting the LSS and clicking on the *Information Panel* button. In our example in Figure 12-14 we have selected LSS 14 which is specified for Open Systems. This LSS contains **8** volumes. None of this volumes is currently part of a FlashCopy or PPRC pair.

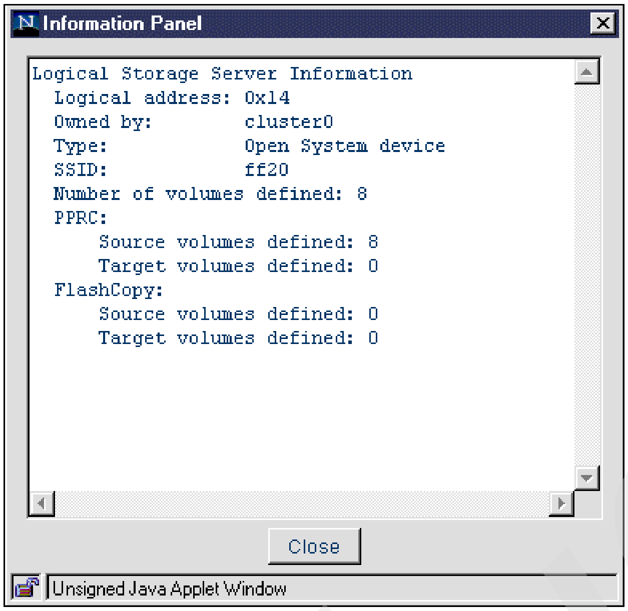

Figure 12-14 Information window of a logical subsystem

You can search for a specific LSS based on its address by selecting the *Find* button of the Storage Servers screen. Figure 12-15 shows an example where we want to find the logical subsystem 16 of the Enterprise Storage Server with the Serial Number 14419.

Figure 12-15 Find Storage Server

In addition you can limit the output of the volume display to a selected range by clicking the *Filter* button and selecting *Filter* devices option.

In our example (Figure 12-16) we want to display physical and logical storage servers only. In addition, we only want to display Open Systems devices that contain volumes which are currently part of a Copy relationship.

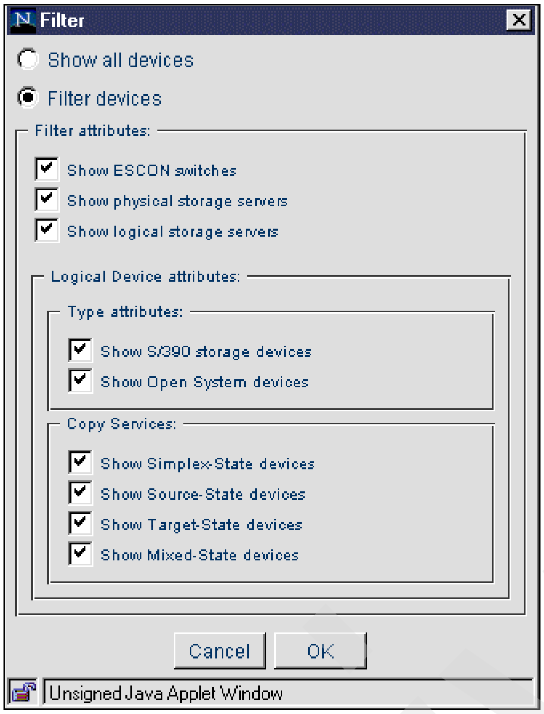

Figure 12-16 Filter logical subsystems volume

Paths panel of the ESS Copy Services Web Interface

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 Server 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.

From the *Paths* panel you will be able to:

- Establish PPRC paths.
- Remove PPRC paths.
- View information about PPRC paths.

A path is always specified between two logical subsystems within Enterprise Storage Servers. Volumes on these logical subsystem can use the paths defined to transfer PPRC data. For availability and performance reasons we recommend to define multiple physical paths between PPRC source and target. Figure 12-17 shows the entry screen of the *Paths* panel.

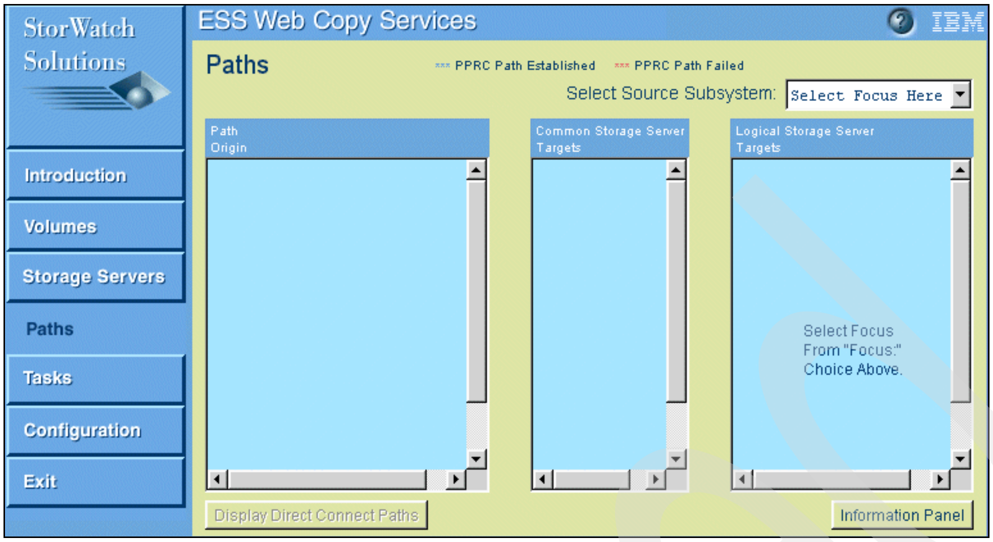

Figure 12-17 Entry screen of the Paths menu

On the top of the *Paths* panel you select the source of the PPRC path, which is done by ESS serial number and LSS number. This will show all configured ESCON adapters of the selected source in the *Path Origin* area.

The ESCON adapters are specified by their system adapter ID (SAID). shows the SAID of all ESS ESCON adapters.

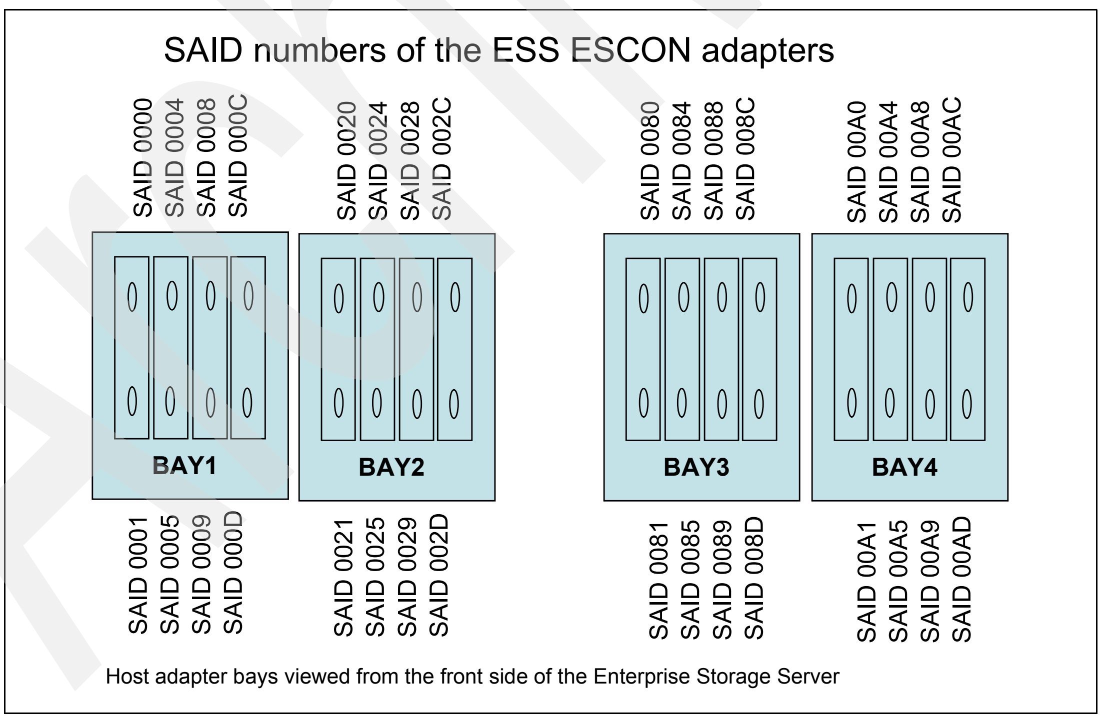

Figure 12-18 SAID numbers of the ESS ESCON adapters

Once a ESCON adapter is selected all Enterprise Storage Server that are connected to this adapter are displayed in the *Common Storage Server Targets* area. All logical subsystems that are available on a particular ESS will be listed in the *Logical Storage Server Targets* area if one of the Storage Servers is selected.

In the example shown in Figure 12-19 we have selected the ESS 2105:14419:14, that is the LSS number 14 of ESS with serial number 14419. The path origin is the ESCON adapter with the SAID0029.

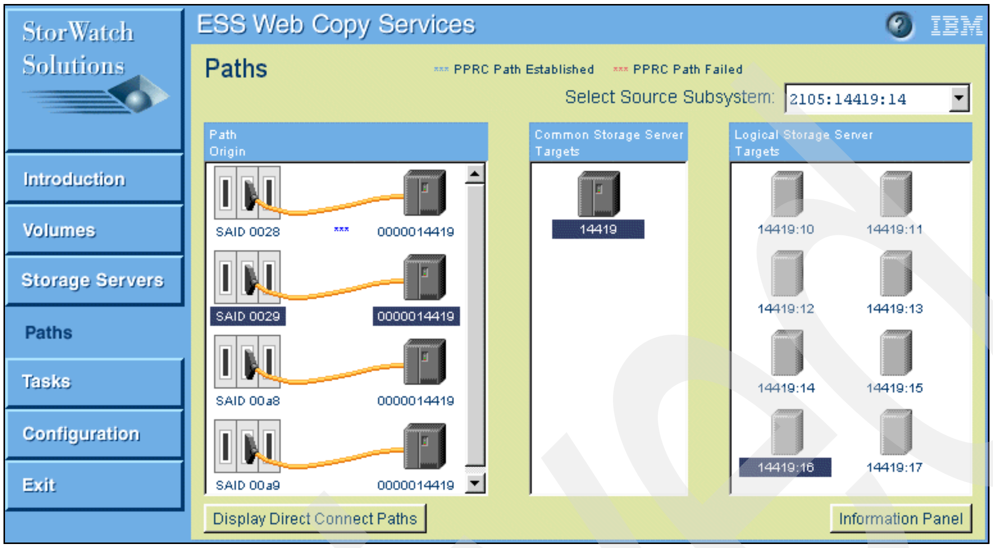

Figure 12-19 Example of the Paths panel

Getting path information

Once a ESCON adapter is selected you can get more information about the paths by clicking the *Information* button at the bottom of the path menu. The example shown in Figure 12-20 shows a path defined between source LSS14 and target LSS16 using the ESCON adapters SAID0028 and SAID00a8 on the same ESS.

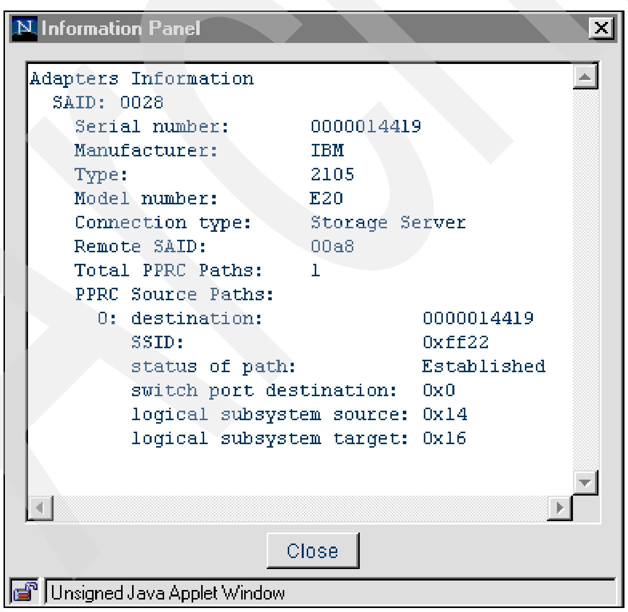

Figure 12-20 Path information panel

If there are paths defined on a ESCON adapter you will find three blue asterisks right below the adapter in the Path Origin.

Figure 12-21 shows ESCON adapters SAID0028 with and SAID0029 without defined path.

Figure 12-21 ESCON Adapter without and with defined path

Tasks panel of the ESS Copy Services Web Interface

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.

In addition, multiple tasks could be grouped together into a single task group. This could be the case if multiple PPRC copy pairs from different Logical subsystems (for example: iSeries needs more than one LSS for capacity) have to be established or if FlashCopy pairs from different Logical Subsystems have to be established at the same time in order to do a backup. All tasks within a task group will be processed in parallel.

With the *Tasks* menu you will be able to:

- View all specified tasks.
- Run previously created tasks.
- Display and modify properties of a task.
- Group or ungroup tasks.
- Remove tasks.
- Verify if the task has run successfully of failed.
- Display information about a task.

Figure 12-22 shows the Tasks panel of the ESS Copy Services. For each task, the name, a description, and the last status of execution is displayed.

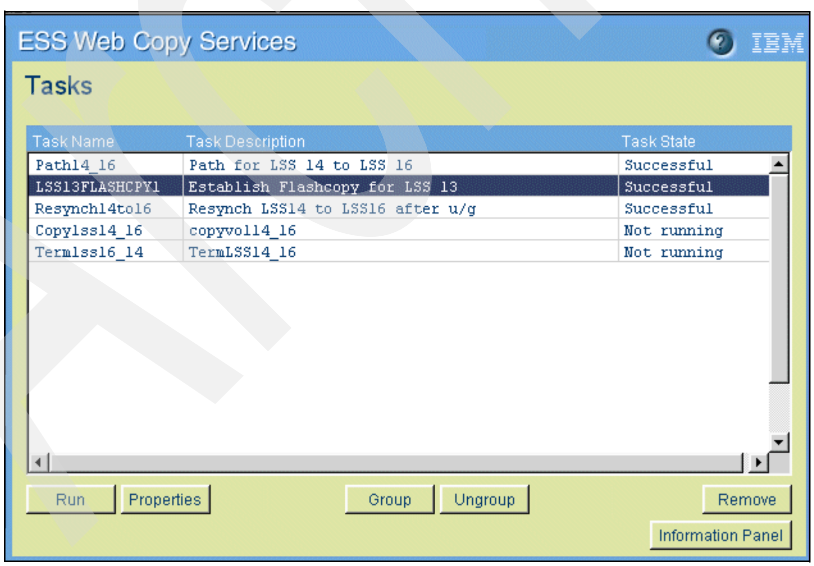

Figure 12-22 Tasks panel of the Copy Services Web interface

Working with tasks

To create a group task click the single tasks you want to group together while holding the Shift key. 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.

An example for the usage of a task group would be multiple FlashCopy pairs from different logical subsystems that needs to be issued all at the same time in order to do a backup.

You can get more information about the setup of a task. Select the task and click the *Information button* at the lower right.

In our example in Figure 12-23 we have a task named Copylss14_16. This task establish a copy pair between LSS 14 and LSS 16.

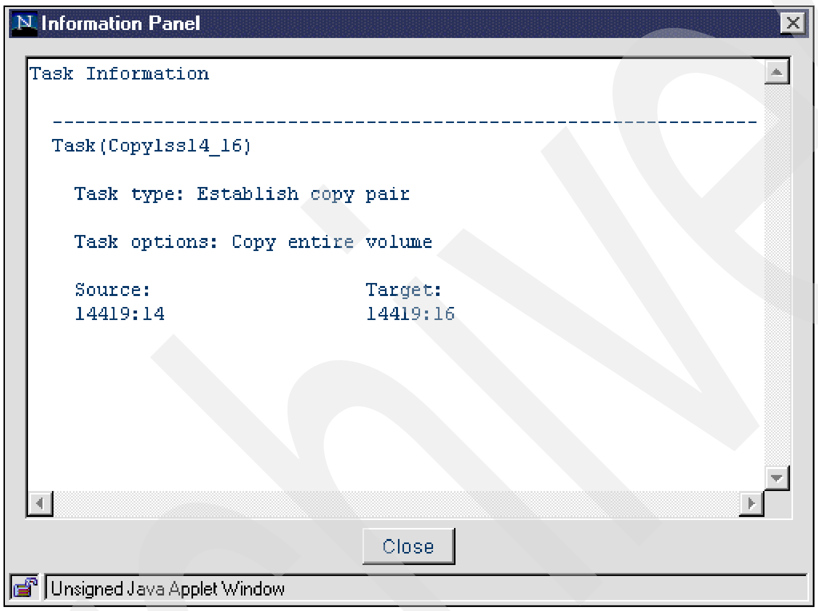

Figure 12-23 Task information panel

To display or change the setup of a task, select the task you want to modify and click the *Properties button*. You can review or change any of the parameters of this task. Once you have completed your changes you have the opportunity to place the existing task, create a new additional task or delete the task if needed.

> **Note:** It is not possible to change the properties of a grouped task. You first have to ungroup the task, make the changes and group the single task together again when finished.

We recommend to always use a task when setting up the ESS Copy Services. This will:

- Simplify the usage of data copies that have to be done periodically.
- Prevent user mistakes, once the task has been created correctly.

12.2.12 Configuration panel of the ESS Copy Services Web Interface

With the *Configuration* menu you will be able to:

- Identify the Copy Services Server.
- View the date of the last configuration update.

- Send a summary of the ESS configuration via E-mail.
- View and modify the E-mail address of the summary recipient.
- View and clear the Copy Services log file.
- Enable and disable password protection for the host commands.
- Administrate users for the host command authorization.

Figure 12-24 shows the *Configuration* panel.

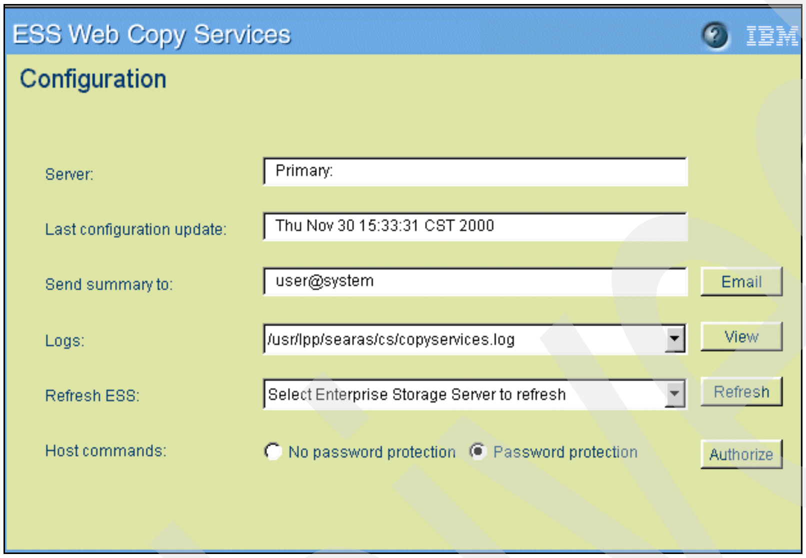

Figure 12-24 Configuration panel of the ESS Copy Services Web interface

The log of the Copy Services provides useful information. If you select the *View* button a new window with the content of the log file is displayed (see Figure 12-25).

From the log you can get various kinds of information like the execution status of tasks or the time needed to complete a Copy Services operation. If you face any problems performing a copy operation we recommend that you always have a look at the log file.

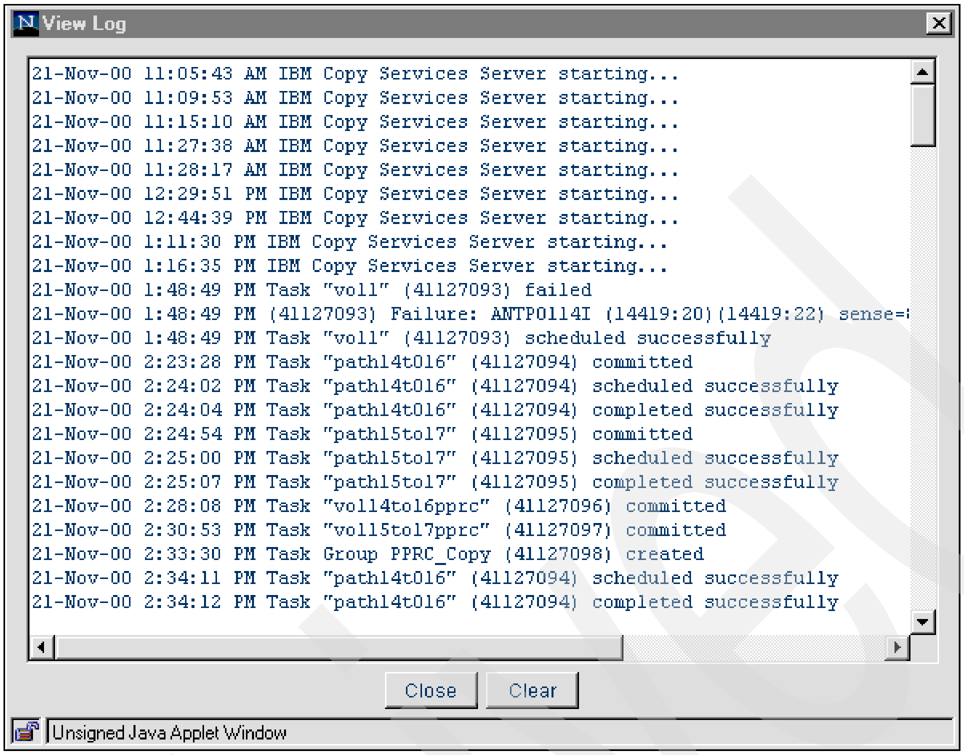

Figure 12-25 Log file of the Copy Services

Once the entries in the log file are not of interest anymore you can clear the entire log by clicking the *Clear* button.

Exit Copy Services of the ESS Copy Services Web Interface

To exit from the ESS Copy Services Web Interface, select the *Exit Copy Services* button. This will cause the current browser window to return to the main ESS Specialist screen.

12.2.13 Implementing PPRC with the ESS Copy Services Web Interface

In this part we will explain how to setup PPRC using the ESS Copy Services Web interface. In general there are two steps needed to successfully establish PPRC:

- Setting up paths between PPRC source and target.
- Establishing the PPRC pairs, either single volumes or entire logical subsystems.
- For ESS Copy Services Messages refer to *Web Interface Users Guide for ESS Specialist and ESS Copy Services*, SC26-7346.

Using Task Maps

If you have used ESS Copy Services before, see Table 12-1 to go directly to the desired task, if not select *Copy Services* button of the ESS Specialist (See Figure 12-5, “Main Menu of the ESS Specialist” on page 338). These task maps refer you to the sections that describe the tasks.

Table 12-1 ESS Copy Services task map for SCSI or fibre-channel attached hosts

Copy Services Task	Related Section
Setting up paths for PPRC	“Setting up paths for PPRC” on page 352

Copy Services Task	Related Section
Establish PPRC pairs	"Establish PPRC pairs (single pairs)" on page 357
Terminate PPRC pairs	"Terminate PPRC pairs (single pairs)" on page 362
Establish and terminate multiple PPRC pairs	"Establish and terminate multiple PPRC pairs at the same time" on page 364
Suspend and resume PPRC pairs	"Suspend and resume PPRC pairs" on page 365
Manually suspend PPRC pairs	"Manually suspend PPRC pairs" on page 366

Please make sure you are aware of the 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 have to the same size as the source or larger
- One primary ESS can have up to four secondary ESS

> **Note:** You must demount the target volume from all hosts systems before establishing a PPRC pair. Be aware that the PPRC process is a destructive operation to the target and will overwrite the data on the target volume.
>
> For iSeries, AS/400 you must power down the backup system or you have to unassign the target volumes.

There are three different ways establishing a PPRC pair:

- From the *Volumes* panel (based on volumes)
- From the *Storage Servers* panel (based on entire logical subsystems)
- From the *Tasks* panel (once a task for PPRC is created)

Setting up paths for PPRC

Before you can establish any PPRC pairs you first have to setup the paths between the source and the targets. The paths are needed for communication between 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 configuration two or even more Enterprise Storage Servers are required.

Use the Paths panel of the ESS Copy Services Web Interface to setup paths for PPRC (Figure 12-26).

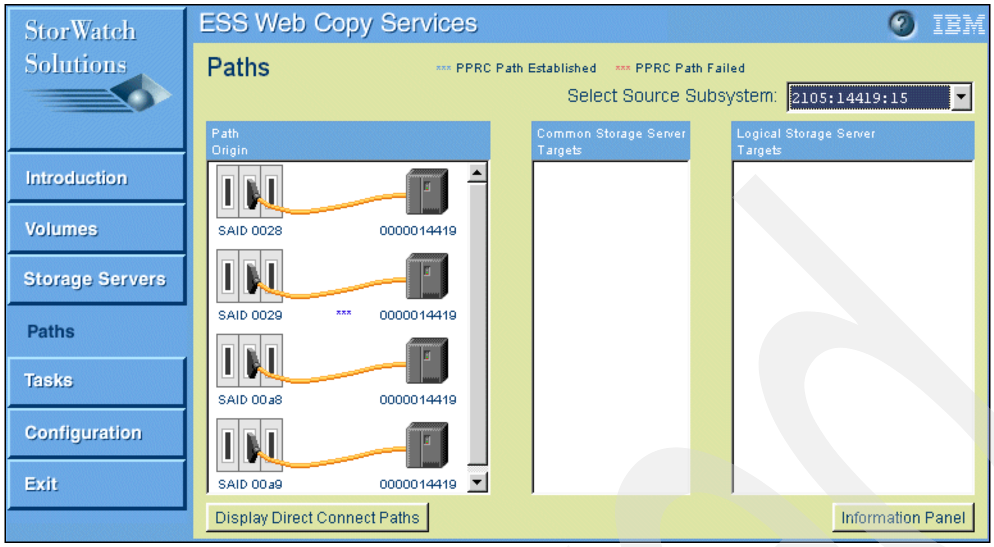

Figure 12-26 Paths window of the Copy Services Web interface

Select the source of the PPRC relationship. This is done with the drop-down menu of the Select Source Subsystem box. All available ESCON adapters for the source will automatically be displayed in the *Path Origin* area (see Figure 12-27).

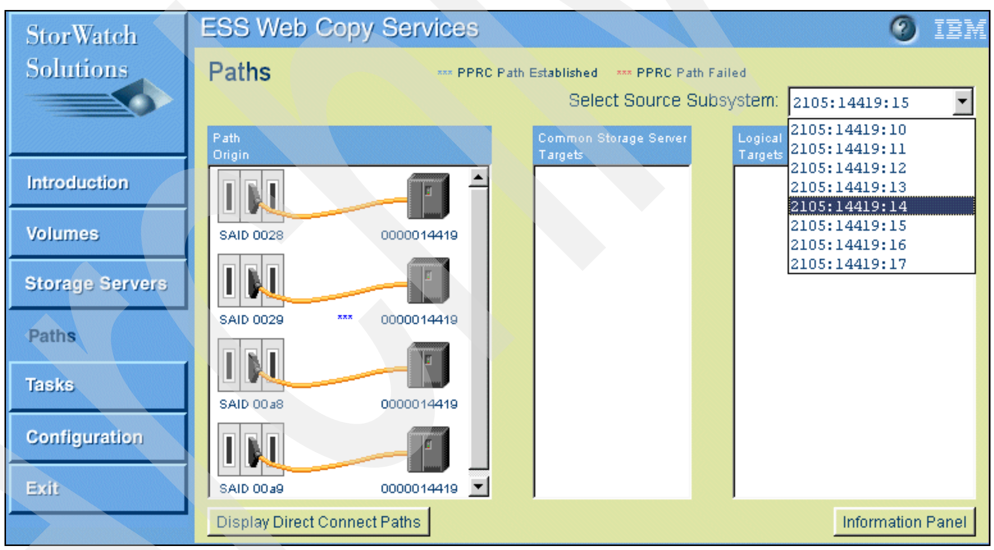

Figure 12-27 Setup PPRC: Source selection

Next select the ESCON adapters you want to use for the PPRC which will automatically show the Enterprise Storage Server that are connected to this ESCON adapter. Select one ESCON adapter with a left-click. Multiple adapters could be selected with a right-click after the first ECSON adapter was selected. If you have chosen the wrong adapter(s) just left-click the correct ESCON adapter again to delete the selection.

Note: For best performance you should establish a logical path from each ESCON port to each LSS (select multiple adapter).

The Enterprise Storage Server that are connected to the adapter(s) will be automatically displayed in the *Common Storage Server Target* area.

Figure 12-28 of our example shows that we have selected the ESCON adapter with system adapter ID SAID0029.

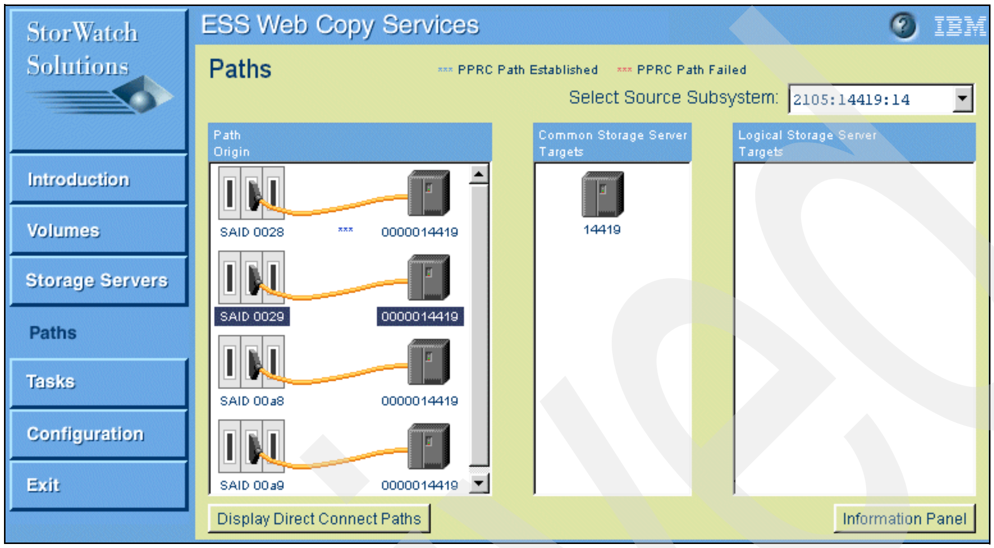

Figure 12-28 Setup PPRC: ESCON adapter selection

Next you left-click the target Storage Server. All Logical subsystems available on the target ESS will be displayed in the *Logical Storage Server Targets* area (see Figure 12-29).

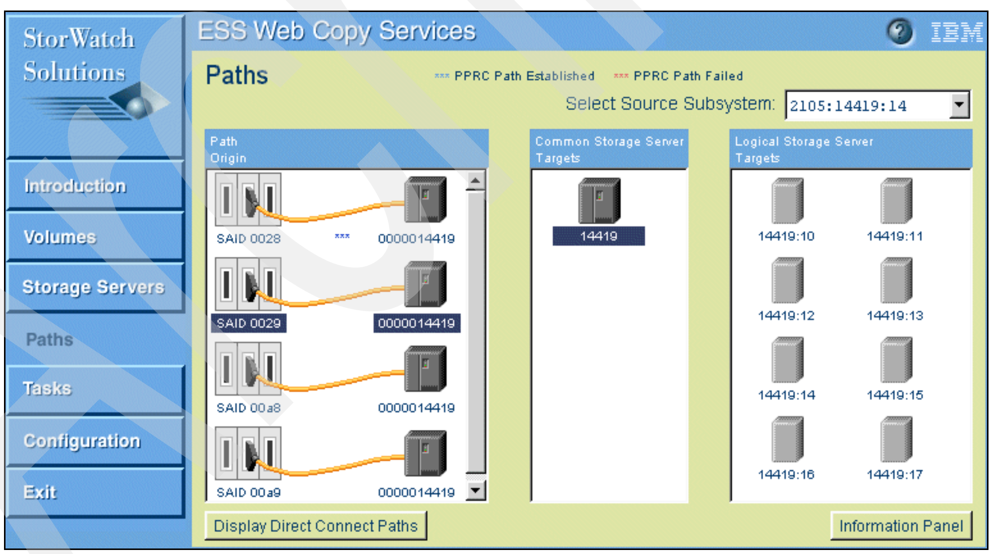

Figure 12-29 Setup PPRC: Target ESS selection

Within the Logical Storage Server targets area select the target LSS of your PPRC path. Select the target LSS with a left-click. Multiple LSS could be selected with the right-click after the first target LSS was selected. If you have chosen the wrong target just left-click the correct LSS again to delete the selection.

In the example shown in Figure 12-30 we have selected LSS 16 to be the target logical subsystem.

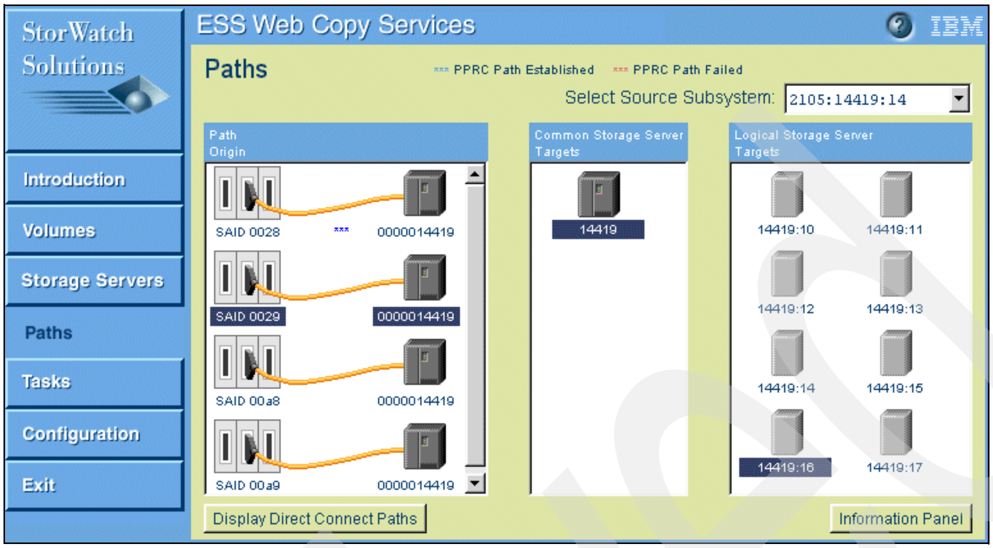

Figure 12-30 Setup PPRC: Select target LSS

Once the target and source of the PPRC path have been selected, right-click one of the highlighted target LSS to bring up the Task Wizard (Figure 12-31). Select the *Establish Path* option and click Next.

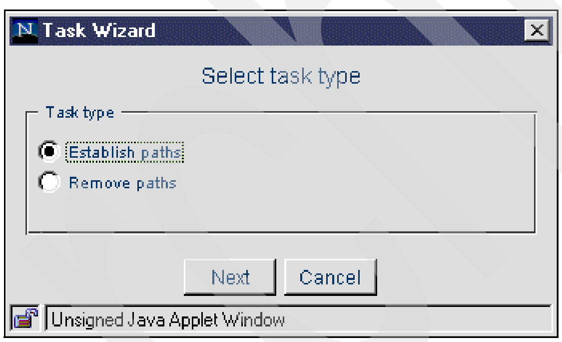

Figure 12-31 Establish PPRC path

Within the next window you can specify the path options (Figure 12-32).

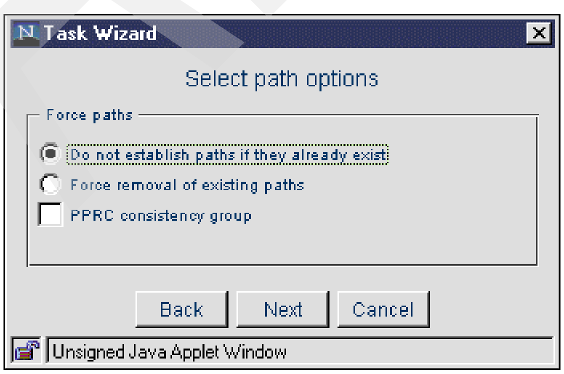

Figure 12-32 Establish PPRC path options

Do not establish path 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 the path will not be executed.

Normally select this option to establish path
For iSeries, AS/400 PPRC Consistency Group is not supported

Force removal of existing paths

If this option is checked and there are already paths defined between the selected source and the target these paths will be removed prior to establish the new paths.

From the next window you can either *Save, Run* or *Cancel* the path task, shown in Figure 12-33. 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 to specify 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 needed to establish a PPRC pair.

Select *Save* to save the definition, and then go to Task Menu to highlight and *run* saved definition.

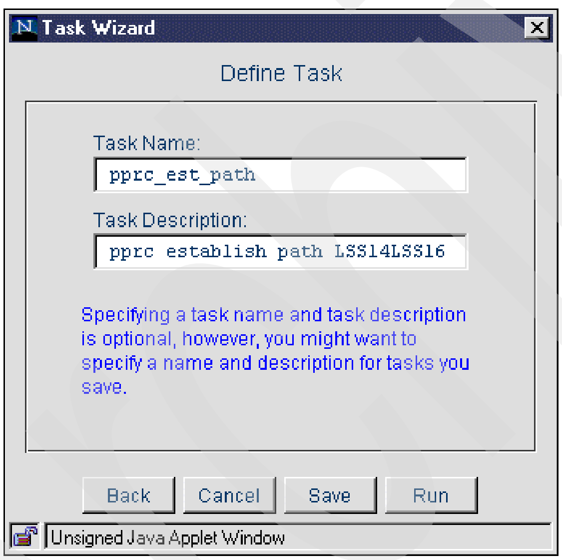

Figure 12-33 Define task window

Once the path successfully was established you will see three blue asterisks directly below the ESCON adapter. This is shown in Figure 12-34.

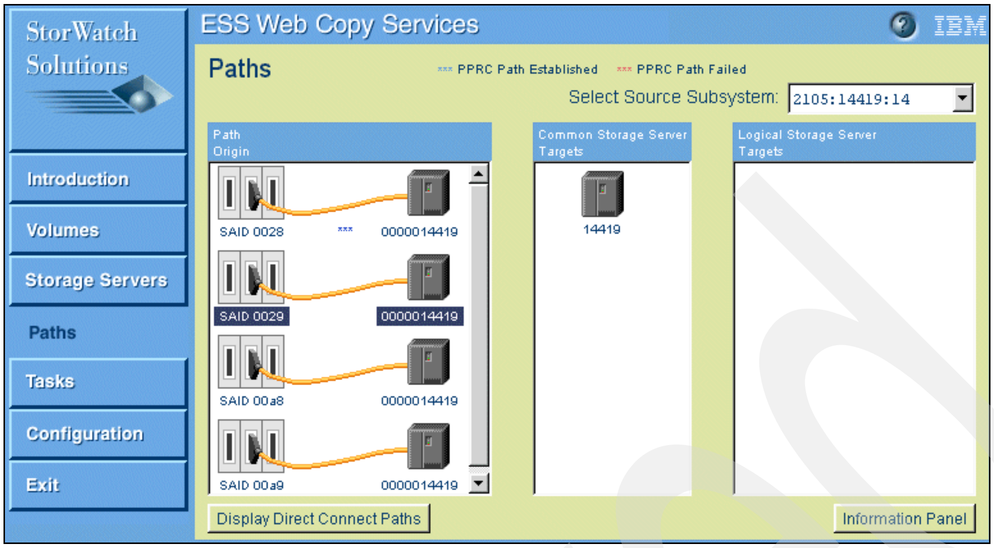

Figure 12-34 Path successfully established (SAID0028)

To get the path information for an ESCON adapter select the adapter and click the *Information Panel*. A window will be displayed showing all path information for the selected adapter. Figure 12-35 shows the path we have created between LSS 14 and LSS 16.

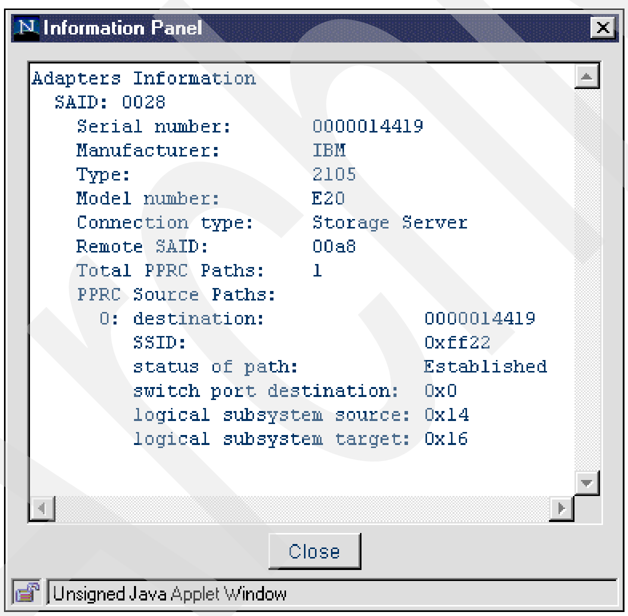

Figure 12-35 Path information window

Establish PPRC pairs (single pairs)

Use the *Volumes* panel to establish single or multiple PPRC pairs (See also “Establish and terminate multiple PPRC pairs at the same time” on page 364). On the left side you 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.

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 12-36 we have selected 2105:14419:14 as source and 2105:14419:16 as target of our PPRC pair.

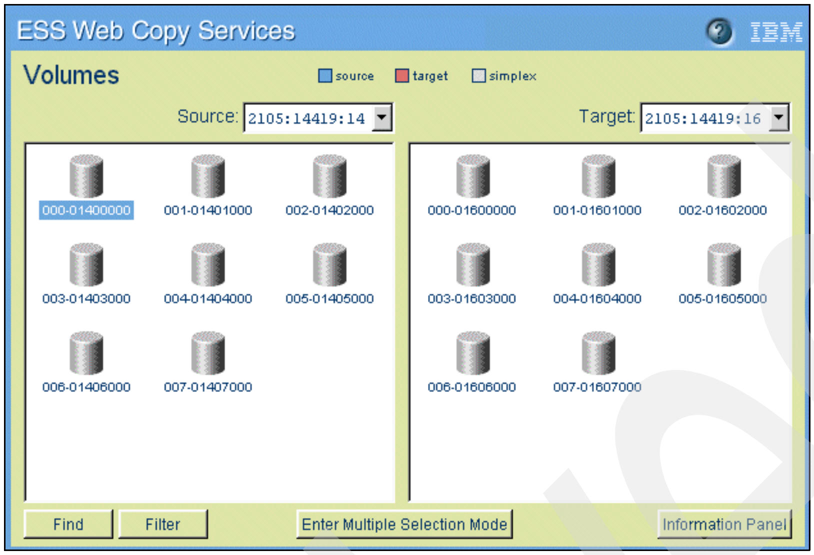

Figure 12-36 Select PPRC source and target LSS

You always need to have two components to establish a PPRC pair, a source and a target. With a left-click you select the source volume and with a right-click the target. If you have selected the wrong source or target volume just left-click the correct source volume again to clear the selection (Figure 12-37).

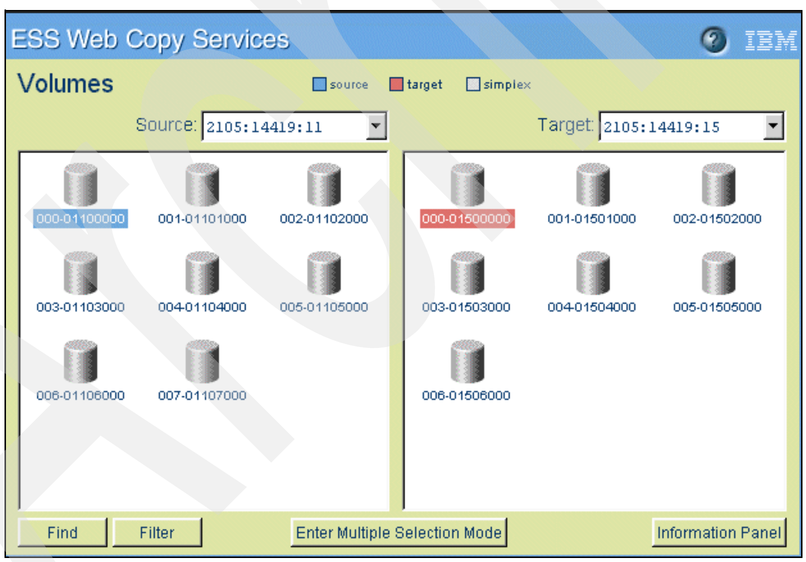

Figure 12-37 Selecting source and target volume

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 12-38). Select the *Establish* PPRC copy pair option and click *Next.*

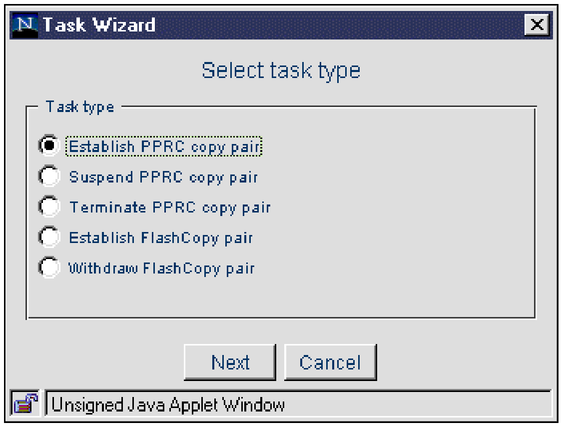

Figure 12-38 Establish PPRC copy pair

Within the next window you can specify the copy options of the PPRC pair (Figure 12-39). **Click Next** when you have finished the selection.

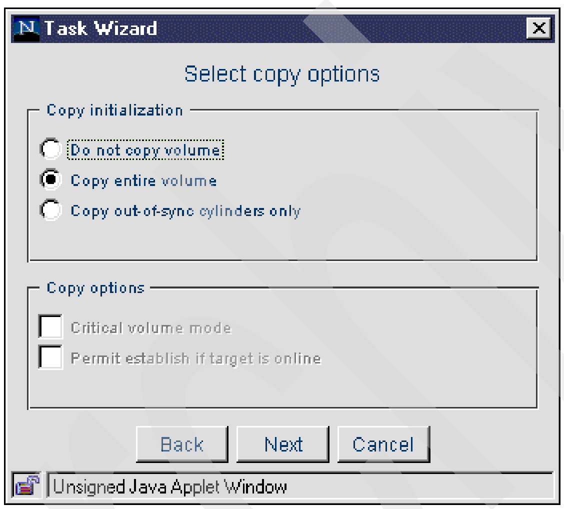

Figure 12-39 PPRC copy options

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 are in sync, that means contain exactly the same 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 re synchronize a PPRC pair.

Critical volume mode

This check-box works in conjunction with the Critical heavy mode check-box of the LSS (see Table 12-2).

Permit if target is online

If this option is checked, the PPRC operation will be performed even if the target volume is mounted on the host system. This is the case, for example, if a UNIX file system is open on the target.

iSeries server volumes cannot be online.

Table 12-2 Critical Heavy mode of PPRC pairs

Critical Volume Checkbox	**Critical heavy not checked**	**Critical heavy checked**
not checked	When secondary volume cannot be updated the pairs suspends. Updates to the primary are allowed	When secondary volume cannot be updated the pairs suspends. Updates to the primary are allowed
checked	When secondary volume cannot be updated the pair suspends. The primary volume is write inhibited only after the last path to the secondary is lost.	When secondary volume cannot be updated the primary and secondary volumes are write inhibited.

From the next window you can either *Save, Run* or *Cancel* the copy task, shown in Figure 12-40. 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 to specify a name and description. This will help with the interpretation of the Copy Services log file later on.

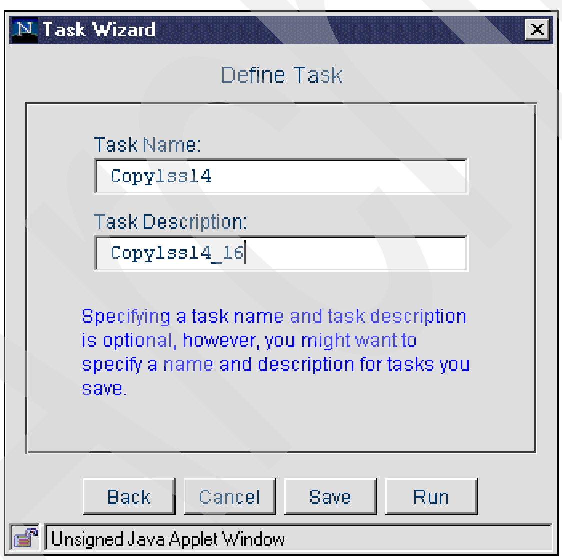

Figure 12-40 Establish PPRC task window

During the short period the PPRC relationship between source and target is established the source and target there is a solid colored triangle displayed within the volumes (Figure 12-41). Once the relationship is successfully established the source and target volume will change its color, indicating the status of the volume (Figure 12-42).

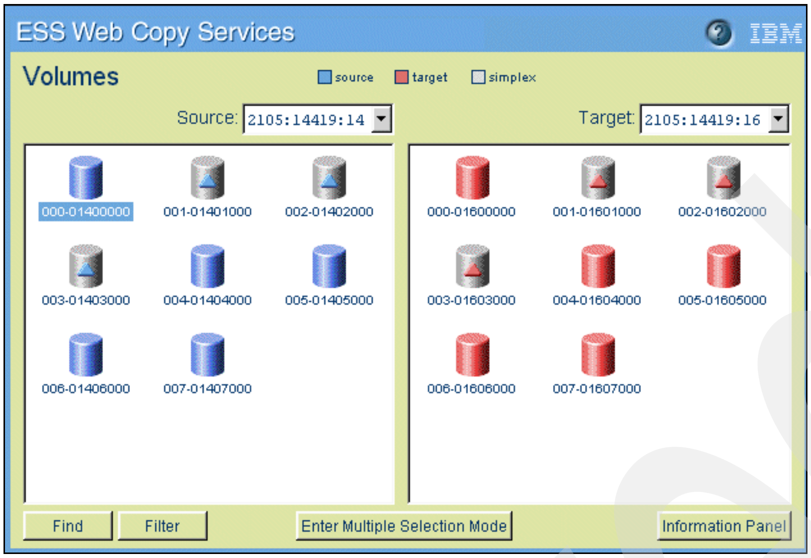

Figure 12-41 PPRC relationship in progress

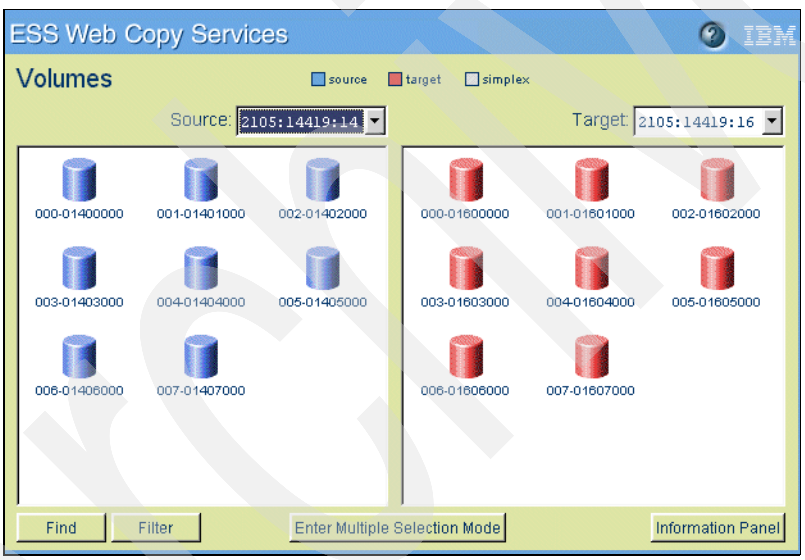

Figure 12-42 PPRC in full copy mode

Select a volume and click the *Information* button to retrieve more information about the status. If the source of a PPCR pair is selected the number of out-of-sync cylinders that are still left to copy are displayed in addition. That are the tracks that needs be copied from the source to the target to achieve full copy mode(Figure 12-43).

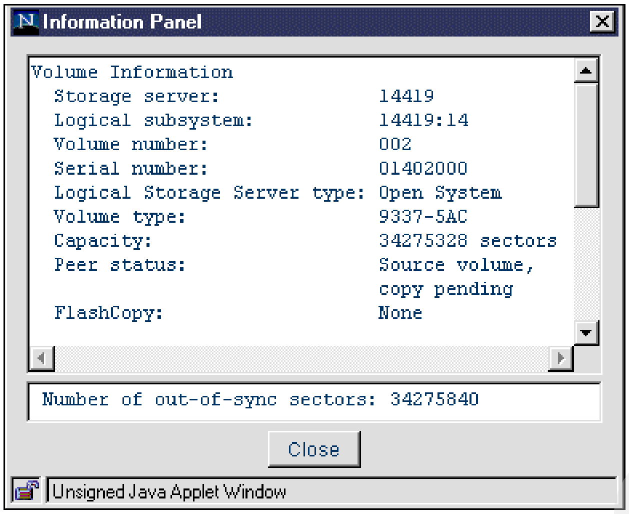

Figure 12-43 Informations window of a PPRC source volume

Note: Once a PPRC pair is established the target volume will not be accessible on any host system until the PPRC relationship will has been terminated.

Terminate PPRC pairs (single pairs)

To end the PPRC relationship you have to manually terminate the PPRC pair. (See also "Establish and terminate multiple PPRC pairs at the same time" on page 364). After a PPRC pair was terminated the target volume will be accessible from a host system again.

To terminate a PPRC pair select the pair you want to terminate by left-click the source and right-click the target (Figure 12-44).

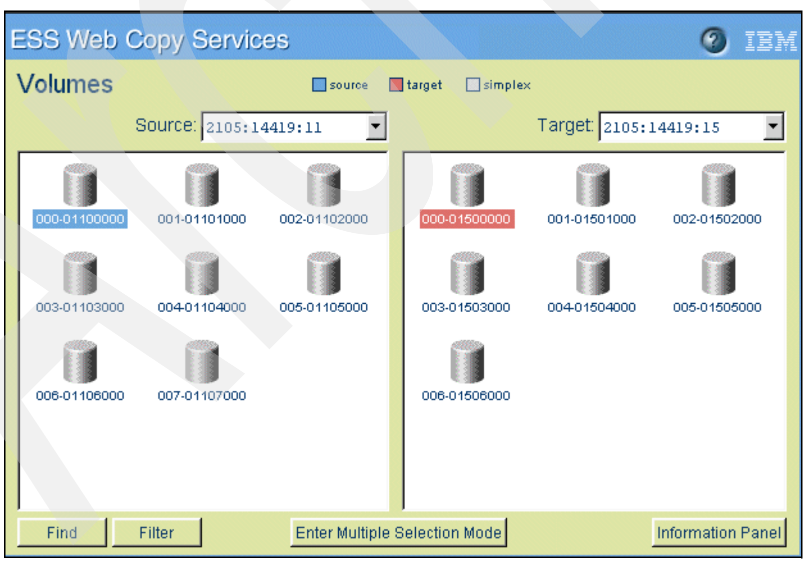

Figure 12-44 Terminate PPRC pair

Next right-click one of the target volumes of the pair to bring up the window, shown in Figure 12-45. From here select *Terminate* PPRC copy pair.

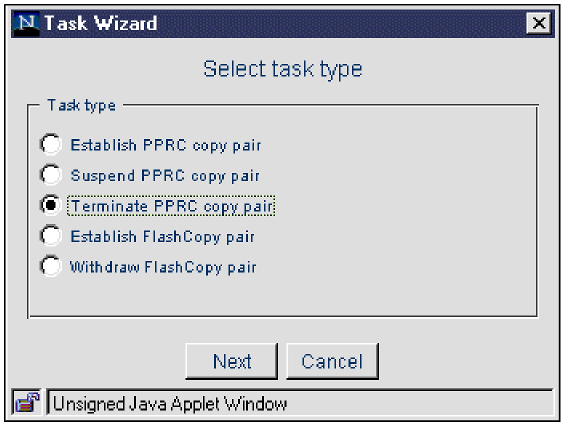

Figure 12-45 Terminate PPRC pair task wizard

Within the next window you have to choose from which ESS you have to schedule the task, the source or the target storage server (Figure 12-46). This means which Enterprise Storage Server should execute the task. An example would be to schedule a termination of a PPRC pair with the target storage server in case the source ESS would not be available.

If *Schedule task with target* logical subsystem is selected to terminate pair, pair will be suspended. Target volume goes to simplex mode and source volume to suspended mode.

If *Schedule task with source logical subsystem* is selected to terminate pair, pair will be terminated. Both Target volume and source volume goes to simplex mode.

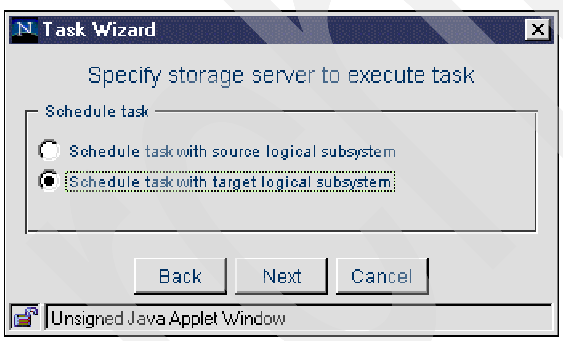

Figure 12-46 Terminate PPRC pair schedule task window

From the next window you can either *Save, Run* or *Cancel* the copy task, shown in Figure 12-47. 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 to specify a name and description. This will help with the interpretation of the Copy Services log file later on.

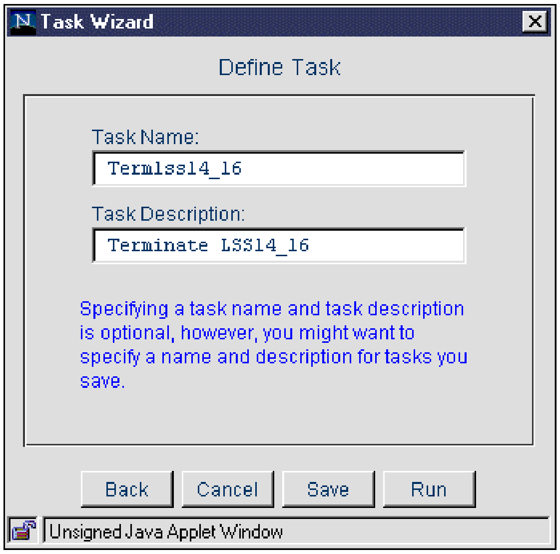

Figure 12-47 Save task window

After the Save. Go to the Tasks panel, click the task and select Infopanel to verify actions

During the short period of terminating the PPRC pair you will see a gray triangle within the source and target volume.Once the pair has been terminated successfully the volumes will be in the simplex state again.

Establish and terminate multiple PPRC pairs at the same time

There are two possibilities to create multiple PPRC pairs at the same time:

- Selecting the entire source and target LSS from the *Storage Servers* menu
(For example: Only iSeries, AS/400 volumes for PPRC within the source and target LSS)
- Using the *Multiple Selection Mode* from the Volumes menu

When using the Storage Servers panel to select entire logical subsystems just treat the LSS like a single volume from the Volumes panel when performing a PPRC operation.

Storage Server menu may only be selected when

1. All volumes in the LSS are iSeries, AS/400 volumes and should be copied
2. The source and target volumes numbers are the same

Next we will give an example of the usage of the *Multiple Selection Mode.*

Go to the *Volumes* menu and click the *Enter Multiple Selection Mode* button at the bottom of the window. Select one pair at a time beginning with a left-click for the source and a right-click for the target. Once you are finished with the selection right-click again on the last target volume which will start the Task Wizard. Continue setting up the PPRC task as described in the previous section.

In the example shown in Figure 12-48 we have selected 8 PPRC pairs. After running the task for these pairs the copy process for all pairs is started at the same time, shown in Figure 12-49.

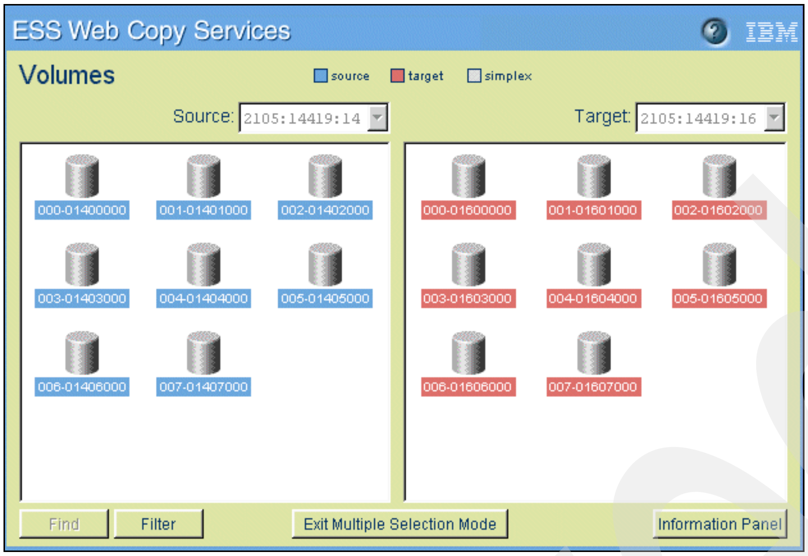

Figure 12-48 Establish multiple PPRC pairs

Figure 12-49 Multiple pairs in full copy mode

The selection of multiple PPRC volume pairs for termination is done the same way.

Another alternative to select multiple volumes for PPRC at the same time is provided through the *Storage Server* panel. Select the source LSS with a left-click. From the drop-down menu at the top of the window select the target ESS and right-click the target LSS. A second right-click on the target LSS starts the task wizard for the PPRC copy task. When selecting an entire LSS all volumes within this LSS are automatically selected for the PPRC copy task.

Suspend and resume PPRC pairs

There are two possible reasons for a PPRC pair to be in the suspended mode:

- The source and target of a PPRC cannot communicate to each other anymore. This could be the case if the source or target volume is not available or the ESCON link in between the pair is offline.

- The administrator manually have suspended the PPRC pair. If a volume is in the suspended mode it is not accessible from the host side even if the volume is connected to a host adapter.

Manually suspend PPRC pairs

The process of suspending a PPRC pair is similar to the termination process.

12.2.14 Setup for PPRC test

- Define PPRC target volumes for LSSxx (refer to Chapter 6, “Creating iSeries storage in ESS” on page 127)

  Add backup or dummy host
  Configure host adapter (SCSI/FC)
  Configure disk group (if necessary)
  Add target Volumes to backup/dummy host
  The Backup System must be power down or you must have unassigned target volumes
- Establish physical ESCON link within the same ESS or between source and target ESS
- Establish paths between source and target LSS’s (Logical link between source and target LSS.
  Selection from ESCON port, source LSS, ESS and target LSS.
  (Save and run the task)
- Establish a relationship between source and target LSS or multiple source and target volumes.
  Define the copy pair task. (Save and run the task)
- Wait until the source and target volumes are synchronized (Duplex, full copy mode)
- Quiesce or force a error condition on the AS/400 (primary or source side)
- Terminate the relationship between source and target LSS or volumes.
  Define the terminate task from target LSS.(Save and run task)
- D-IPL the Backup host (Tape or CD)
- Install SLIC (Opt.:2)
- Recover the remote load source unit (Opt.: 16).
- IPL the backup host.(Change the system IPL values, IPL type “2”, Autoconfig “off”)
- Change the HW configuration tables
- ReIPL Change IPL type “2” to “0”

13

FlashCopy

In this chapter, we cover some of the additional features of the Enterprise Storage Server and discuss how they might be used in an iSeries environment.

13.1 Overview

FlashCopy provides an instant or point-in-time copy of an ESS logical volume. With systems which recognize individual volumes, such as NT or Unix systems, it is quite easy to identify those volumes which should be copied and those which do not. However, with iSeries, this is not the case, as all disk storage is considered as part of single level storage, along with main memory. Indeed, even if you define a User Auxiliary Storage Pool (ASP) there is still a link from the User ASP to the System ASP. Within this document, we will refer to “volumes” meaning the whole DASD configuration within the ESS for a single iSeries system.

FlashCopy functions gives an (almost) instantaneous copy or “view” of what the original volumes looked like at a specific time when the copy was initiated. This is a so called T_0 (time-zero) copy.

Support is provided from the Web-based interface which is part of the StorWatch ESS Specialist Copy Services functions. As of this writing, support for the Command Line Interface (CLI) has not been made available for OS/400.

When FlashCopy is invoked, it creates a “map” of the source volumes. This process takes only a few seconds and when complete, you have access to the T0 copy of the source volumes. As soon as the map of the target volumes has been created, you can begin to read and write to both the source and target volumes.

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 creating a database for data-mining purposes. The copy looks exactly like the original source volume and is instantly available. Figure 13-1 shows this concept.

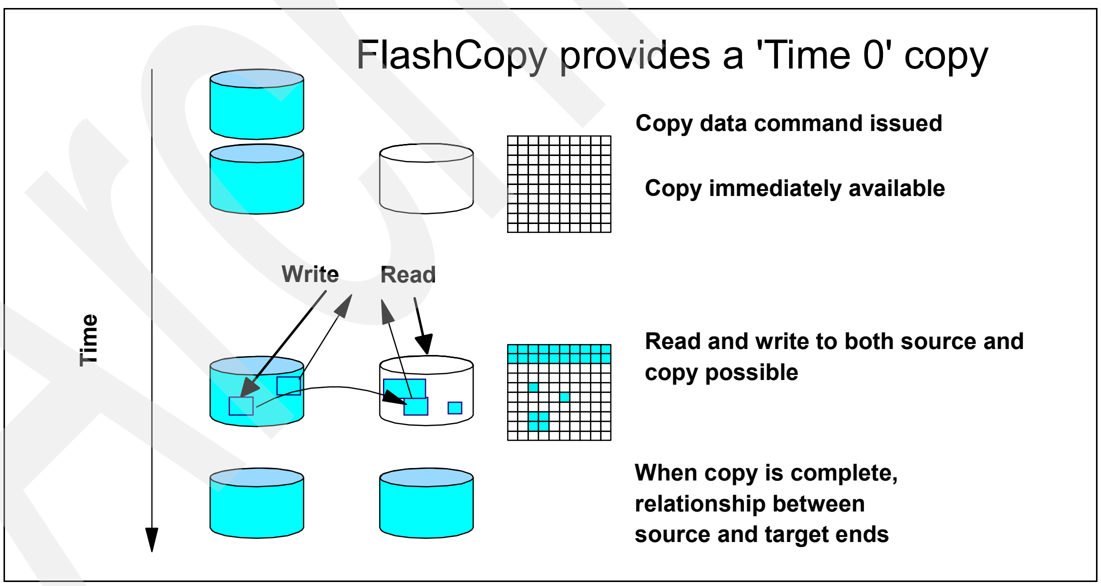

Figure 13-1 Implementation of FlashCopy

13.2 FlashCopy implementation

FlashCopy is only possible between disk volumes. FlashCopy requires a target volume to be within the same logical subsystem (LSS) as the source. When you set up the copy, a relationship is established between the source and the target volume. Once this relationship is established, the volume copy can be accessed. See Figure 13-1 for an illustration of FlashCopy.

Once the relationship between source and target volumes is established, an optional background task copies the tracks from the source to the target. You can suppress this background copy task using the NOCOPY option. This may be useful if you need the copy only for a short time; such as when you are using the flashcopy to make a backup to tape. The NOCOPY option can be activated through the ESS Specialist.

The source and target volumes are available immediately for full read and write access. If an application wants to update the target before the source data has been copied to the target, the data is first read from the source, written to the target, then updated by the application. The target volume's bitmap is updated to indicate that the copy has been made.

This is also true if an application wants to update the source volumes before the data has been copied to the target. The data is first copied from the source to the target, and then updated by the application. The ESS keeps track of which data has been copied from source to target. If an application wants to read some data from the target that has not yet been copied, the data is read from the source and copied to the target. All subsequent requests are then satisfied from the target volume.

A source volume and the target can be involved in only one FlashCopy relationship at a time. The relationship ends when the physical background copy task has completed. If you had started FlashCopy from the StorWatch ESS Specialist with the NOCOPY option, you must then withdraw the pair (a function you can select) to end the relationship.

In addition, you cannot create a FlashCopy on one type of operating system and make it available to a different operating system. On other platforms, you can make the target available to another host running the same type of operating system. However, this is only possible on iSeries if the remote load source mirror and whole disk configuration are used by the "receiving system". In this case, the receiving system would lose its own identity and take on that of the source system.

13.2.1 To copy or not to copy?

Due to the whole iSeries disk space being regarded as a single entity, we feel it is better to use the NOCOPY function (defined as "*Do not perform background copy if checked*" - see Figure 13-2) if your main reason for using FlashCopy is to take backups on another system. This is due to length of time to copy the entire contents of the primary set of disks and therefore more sensible to use the instant T_0 copy and start the backups immediately. All read access would then be satisfied by the source volumes.

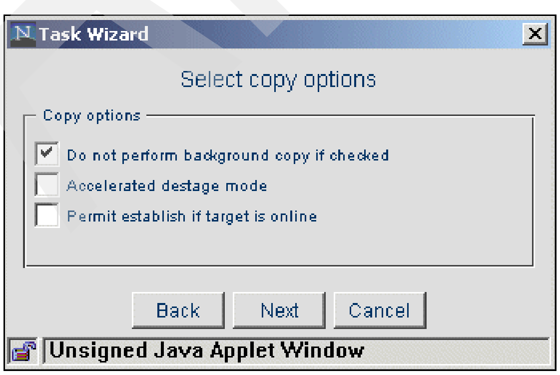

Figure 13-2 Choosing the NOCOPY option

13.3 Implementing FlashCopy with ESS Copy Services Web Interface

In this section, we explain how to set up FlashCopy using the ESS Copy Services Web Interface.

Please make sure you are aware of the requirements of the FlashCopy functionality:

- The source and target volume have to be in the same LSS.
- The target volume has to be the same size as the source volume or larger.
- A volume can be only in one FlashCopy relationship at a time.

> **Note:** We recommend that you discontinue the use of target volumes from all host systems prior to performing a FlashCopy. Be aware that the FlashCopy process is a destructive operation, and all existing data on the target volumes will be lost.

There are two different ways of establishing a FlashCopy pair:

- From the **Volumes** panel
- From the **Tasks** panel (once a task for FlashCopy is created)

13.3.1 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.

You always need to have two components to establish a FlashCopy pair: a source and a target. With a left-click you select the source volume and with a right-click the target. If you have selected the wrong source or target volume, just left-click 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 13-3). Select the **Establish FlashCopy pair** option and click Next.

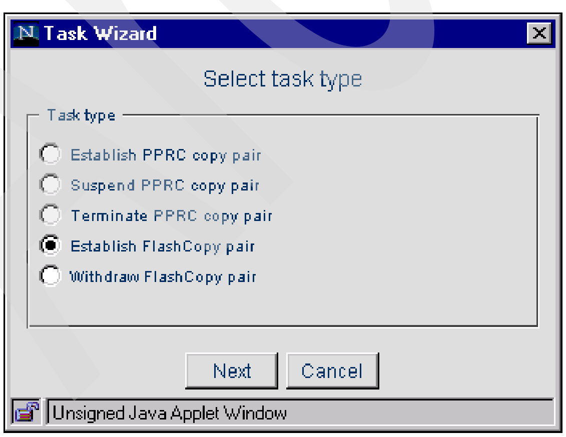

Figure 13-3 Task Wizard window

Within the next window you can specify the copy options of the FlashCopy pair (Figure 13-4).

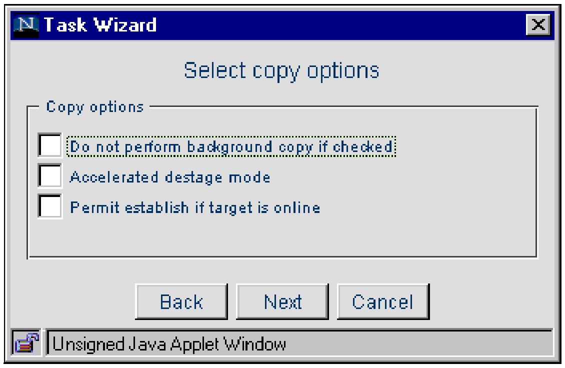

Figure 13-4 Select copy options window

Do not perform background copy if checked

If this option is checked, only the tracks that are modified on the source volume are copied to the target volume. The 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. Once all data is copied, this relationship ends automatically.

Accelerate destage mode

If this option is checked, I/O of the FlashCopy process gets a higher priority than other I/O requests at the same time. Therefore the data from the FlashCopy process that is staged into cache will be destaged to the disk sooner than with the normal destage algorithm.

Permit if target is online

If this option is checked, the FlashCopy will be performed even if the target volume is in use of an operating system. This is the case, for example, if a AIX volume group is active on the target.

From the next window you can either **Save**, **Run**, or **Cancel** the copy task, shown in Figure 13-5. Once a task is saved, it can be executed from the Task panel at any time. Optionally, a name and description of the task can be specified. Even if you do not want to save the task you have created, we recommend that you specify 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.

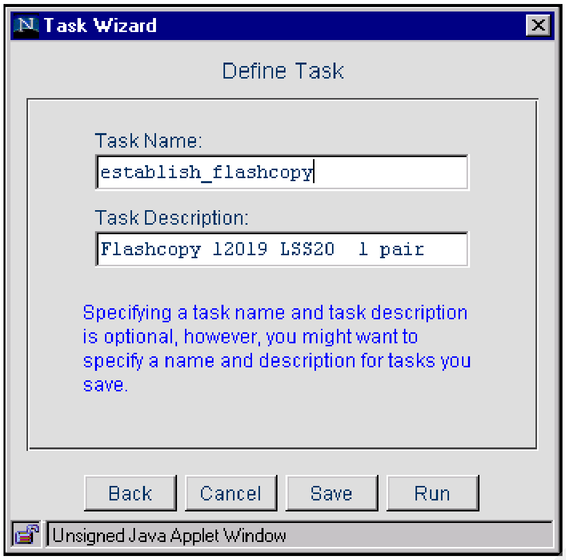

Figure 13-5 Define task window

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. Two triangles within the source and target volume will be displayed. The color of the triangles defines whether it is a source or a target. The legend of the color is shown at the top of the volumes panel.

During the short period where the copy pair is created, only one of the triangles is filled. Once the relationship has been established successfully, both of the triangles will be solid-colored. This is illustrated in Figure 13-6.

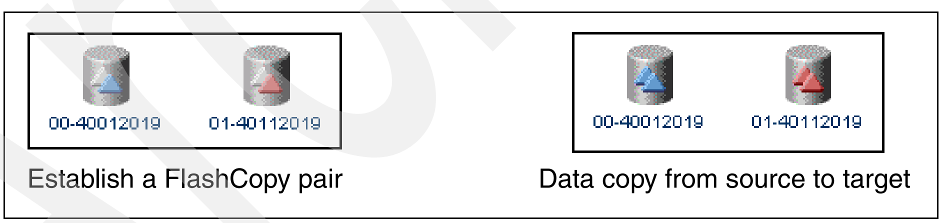

Figure 13-6 FlashCopy Volume display

13.3.2 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 will also see how many tracks still have to be copied to the target volume (Figure 13-7).

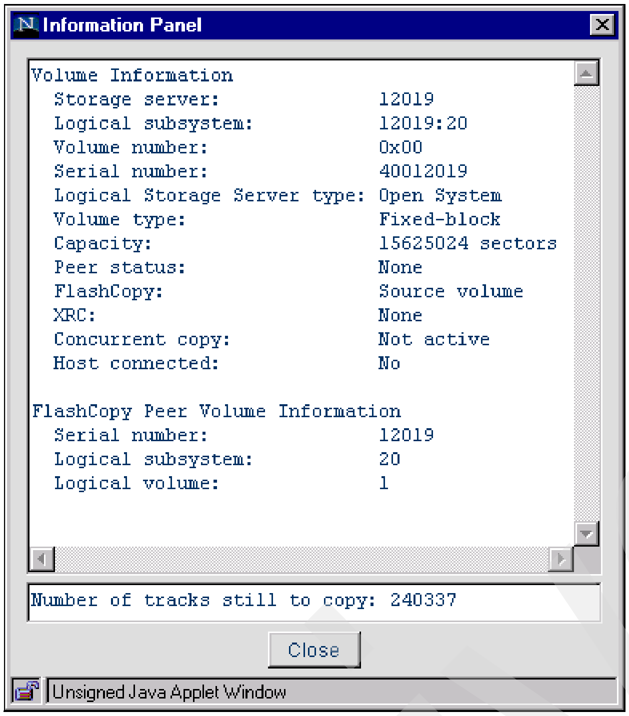

Figure 13-7 Information panel of a FlashCopy source

13.3.3 Withdrawing a FlashCopy pair

In the following cases, you need to withdraw a FlashCopy pair:

- If a FlashCopy pair is not needed anymore, but it has not yet finished the background copy.
- If a FlashCopy pair that was created with the NOCOPY option is not needed anymore.

To withdraw a FlashCopy pair, select either the source or the target volume with the left-click and start the Task Wizard by right-clicking one of the volumes of FlashCopy pair. This will bring up the window shown in Figure 13-8.

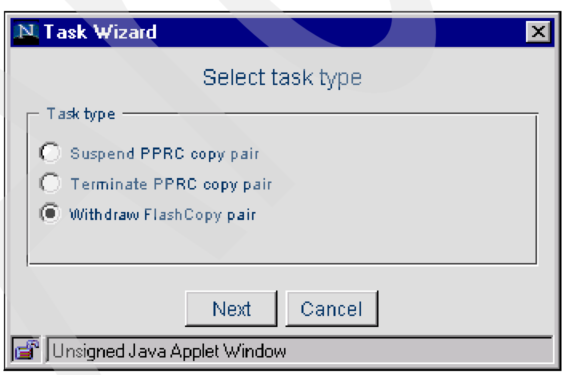

Figure 13-8 Withdraw a FlashCopy pair

Select the Withdraw FlashCopy pair option and decide whether to Save, Run, or Cancel the task (Figure 13-9).

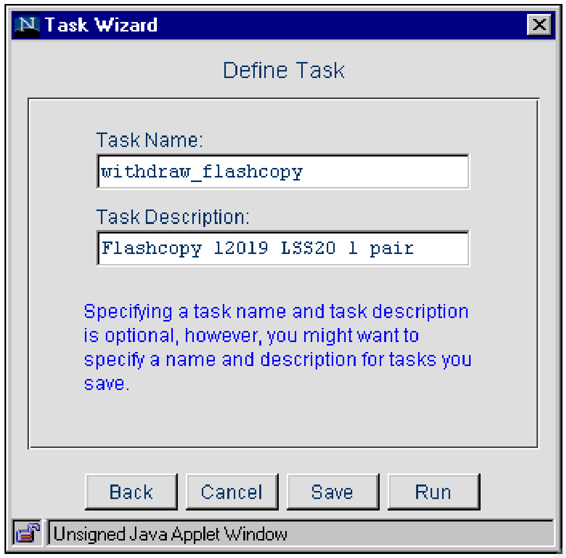

Figure 13-9 Define task window

13.3.4 Selecting multiple volumes for a FlashCopy task

In some cases you may want to establish or withdraw multiple FlashCopy pairs at the same time. Therefore you have the possibility to define multiple FlashCopy pairs within the volume panel. Click the **Enter multiple selection mode** button at the bottom of the volumes panel. Select one pair at a time beginning with a left-click for the source and a right-click for the target. Once you are finished with the selection, right-click again on the last target volume, which will start the Task Wizard. Continue setting up the FlashCopy task as described in the previous section.

In the example shown in Figure 13-10 we have created 6 FlashCopy pairs within the same LSS. After running the task for these pairs, the copy process for all pairs is started at the same time.

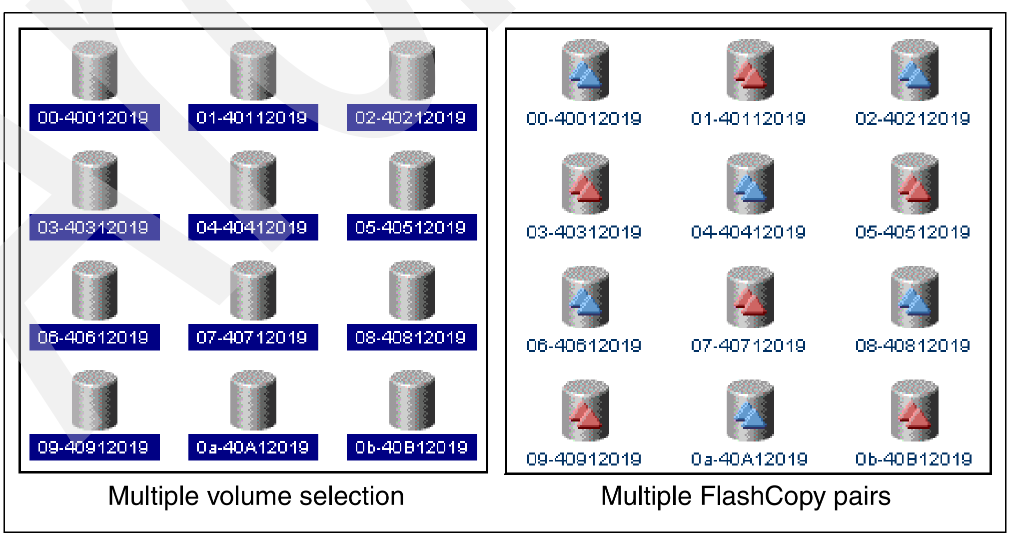

Figure 13-10 Multiple volume selection and FlashCopy pairs

If you have made a mistake during the volume selection, you have to exit the multiple selection mode by clicking **Exit multiple selection mode** and enter the multiple selection mode again afterwards.

> **Note:** The multiple selection mode is limited to one LSS within the source and target area. Once you have entered this mode, the drop-down menu of the selection area will be disabled.

13.3.5 Configuration tips

If you are creating a FlashCopy task that involves multiple source and target volumes, there is a quick way to create the task to withdraw the FlashCopy pairs. See Figure 13-11.

At the Tasks panel:

1. Click the task you created that establishes the FlashCopy relationships.
2. Click the **Properties** button.

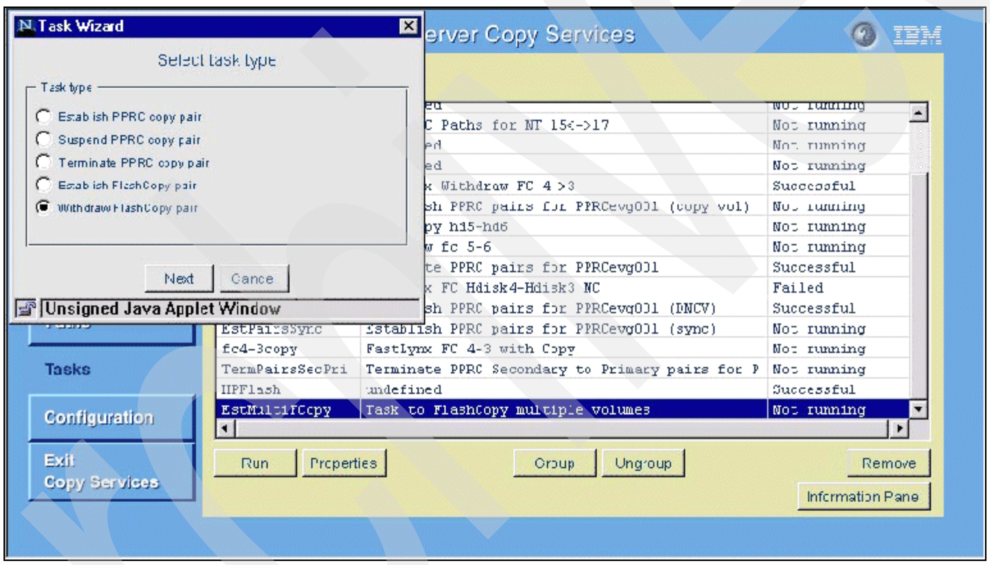

Figure 13-11 A quick way to withdraw FlashCopy pairs

3. Click the **Withdraw** FlashCopy Pair button, then click Next.
   See Figure 13-12.

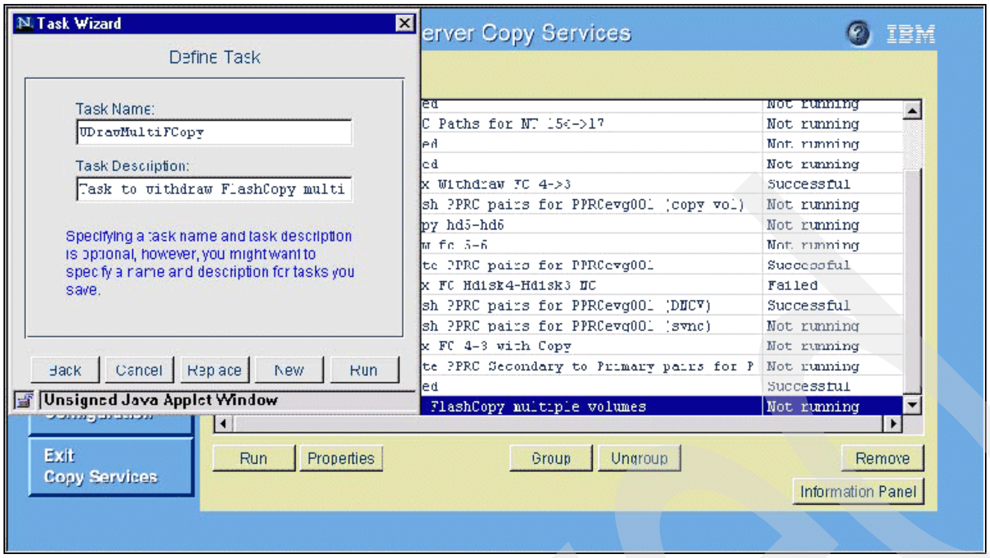

Figure 13-12 Withdraw FlashCopy Pair

4. Type the new Task Name and Task Description, then click **New**. See Figure 13-13.

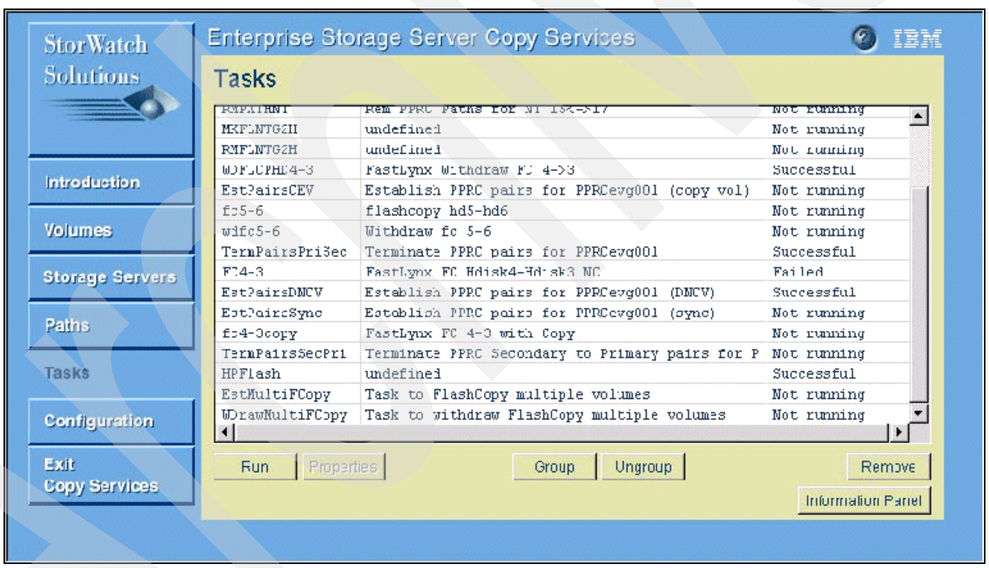

Figure 13-13 Naming the task to withdraw

13.3.6 Using FlashCopy with PPRC

One good example of combined PPRC with FlashCopy is the split mirror solution for backup and recovery. This solution has been tested by SAP for SAP R/3 and S/390 with ESS and its description may also help you to design your own split mirror solutions for other applications.

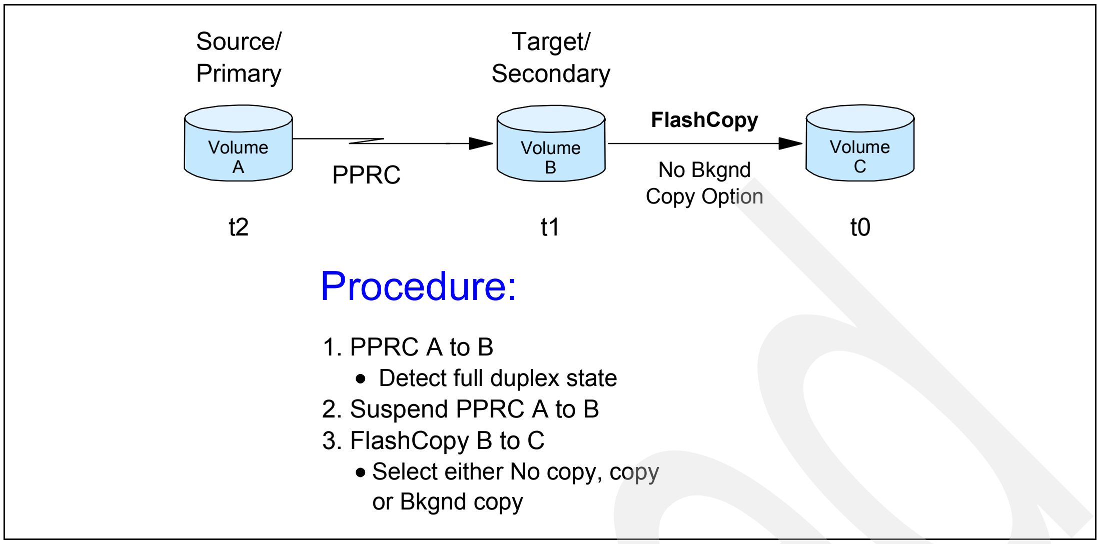

Figure 13-14 PPRC with FlashCopy

The basic idea of split mirror solution is managing a consistent copy of the database in a certain point of time on a secondary site (the t_0 copy) while production continues on a primary site on the t_2 copy. On the secondary site, a Flash Copy of the consistent t_1 copy is created, referred to as t_0 copy, that enables off-line backup to tape and application testing, see Figure 13-14. PPRC pairs are normally suspended, the mirror is split, preventing propagation of logical errors to the t_1 secondary copy. At certain checkpoints the pairs are resumed and synchronized, that means the database is copied to secondary. This allows restart of application on last known consistent copy or roll forward to a certain point of time using log files. The database active and archive logs are copied constantly, the volumes they are at are not suspended.

In case of application logical error that is more frequent than a total site failure the production resumes on the secondary database while the primary can be analyzed. In case of a hardware failure you can switch quickly to the secondary site.

The implementation planning should include:

1. Identification of logical to physical volume mapping
2. Preparation of both the primary and secondary site ESS's
3. PPRC ESCON channels setup and paths setup
4. PPRC tasks setup
5. Recovery site hosts setup
6. Testing

Identification of mapping will include creating a list of physical volumes that contain the database and the logs and checking they are placed on physical volumes properly in a way they can be suspended and resumed. For example, the database cannot be on the same physical volume the logs are on.

The preparation of the ESS's means purchasing the PPRC features and installing ESCON adapters, checking that disk space is available on the selected LSS.

You need to have PPRC-dedicated ESCON channels available between the primary and secondary sites.

Tasks have to be setup using the Web interface and then they are ready to be executed. You also need to group the tasks into task groups to be able to carry the point in time critical operations (establish, suspend, resync, FlashCopy) into task groups.

The host server on the remote site has to be setup so it can takeover in case of a primary site failure.

All the PPRC tasks, backup operations and site takeover have to be tested regularly.

Before you start using split mirror routinely, you have to perform the initialization steps, either with the Web interface.

1. Establish the paths between primary and secondary ESS
2. Establish the pairs for all necessary volumes (establish, copy entire volume)
3. Query the pairs if establish is complete
4. Suspend write I/O to the primary t_2 volumes
5. Withdraw the pairs between primary and secondary ESS (suspend pairs)
6. Resume write I/O to the primary t_2 volumes
7. FlashCopy the t_1 volumes on the secondary ESS to t_0 copy volumes on the
8. Resynchronize the primary t_2 and secondary t_1 volumes (establish, copy changes only)
9. Dump or backup the t_0 copy to tape

Repeated routine steps will include:

1. Create safety copy if resynchronization fails (suspend all, recover t_1 and FlashCopy t_1 to t_0)
2. Resynchronize the primary t_2 and secondary t_1 volumes
3. Query the status of resynchronize process
4. Suspend write I/O to the primary t_2 volumes
5. Withdraw the pairs
6. Resume the write I/O
7. FlashCopy the t_1 volumes to t_2
8. Resynchronize the primary and secondary volumes for BSDS, active and archive logs.

Dump or backup the copy to tape.

14

Taking backups from a FlashCopy

In this chapter, we discuss how FlashCopy can be used to allow backups to be taken while the production system or partition continues with normal production workload, thus minimizing the downtime necessary for taking backups.

> **Note:** Throughout this chapter we will refer to partition, but this could be a separate system. However, because FlashCopy is most likely to be used on an LPAR system for performing “offline backups”, we will use the term partition.

For the latest information on taking backups with FlashCopy, refer to the iTC Web site at:

`http://www-1.ibm.com/servers/eserver/iseries/service/itc/`

Also refer to the offering document at:

`http://www-1.ibm.com/servers/eserver/iseries/service/itc/pdf/Copy-Services-ESS.pdf`

First of all, we look at a typical setup and the steps necessary to do the backups. Then we look at how this solution can work with the IBMs Backup Recovery and Media Services to provide a fully managed backup and recovery environment.

14.1 Steps required to take backups

Due to the requirement to have the LSU internal to the iSeries, using FlashCopy to assist with backups means that we must recover the mirrored load source before we could do a backup of second copy of data to tape. With the advent of Logical Partitioning (LPAR), and particularly the advanced functions introduced with V5R1 on 8xx hardware, we can now take advantage of these facilities to minimize the downtime for backups.

As we discussed in Chapter 11, “Load Source Mirroring in the ESS” on page 303, recovering the mirrored load source requires that you perform a D-IPL to re-load LIC onto the internal LSU drive. With LPAR and the Enterprise Storage Server, this is quite feasible, especially since it is no longer necessary to have a minimum of one processor per partition.

Consider the setup shown in Figure 14-1. Here, we have a two partition iSeries model 820. The primary partition P0 (Blue) runs the production workload. It has an internal LSU, disk B1. This is mirrored to an equivalent sized LUN in the Enterprise Storage Server, LUN B1’. This is shown as step 1 in Figure 14-1. The remainder of the disk storage for the Blue partition is in the Enterprise Storage Server (B2 to B7).

The secondary partition P1 (Red) has the single internal LSU, disk R1. It is a minimum configuration having a single IOP, an IOA for the LSU, an IOA for the console, an IOA for attachment to the ESS and an IOA for the tape drive. Clearly, if additional devices are required, then additional IOPs and IOAs may be needed to support them.

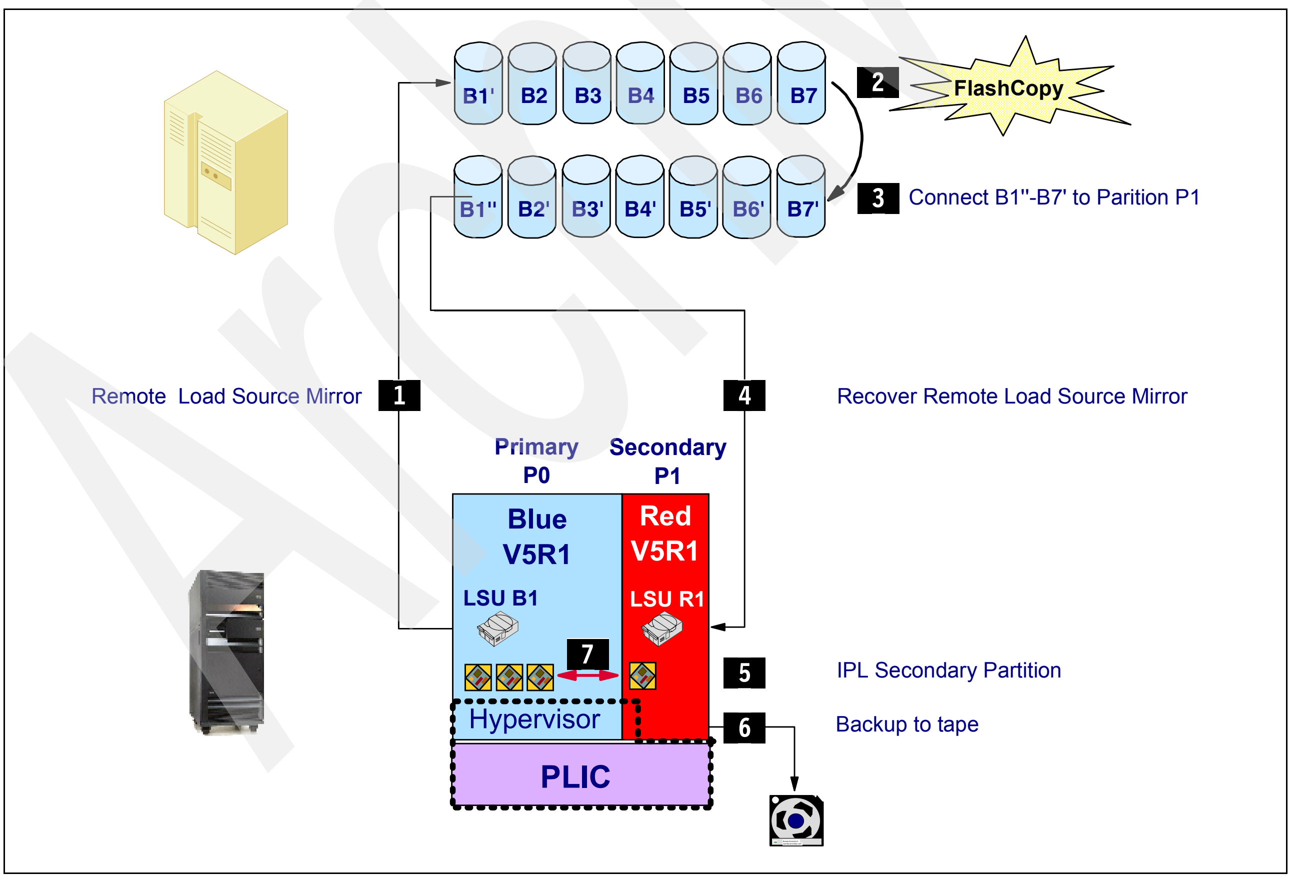

Figure 14-1 Using LPAR with FlashCopy to perform backups

Copy Services has been setup in the Enterprise Storage Server to allow FlashCopy of LUNs B1' - B7 to LUNs B1" - B7'.

Under normal daily operations, the production workload is running in partition P0. Partition P1 is dormant and is using no CPU or memory resources.

1. When backups need to be taken, partition P0 is brought to a restricted state using the command

   `ENDSBS SBS(*ALL)`

   When the restricted state has been reached, all user data for partition P0 has been flushed from memory to disk.

2. At this stage, FlashCopy can be initiated on the Enterprise Storage Server. This is shown as **2** in Figure 14-1. When FlashCopy has completed, in partition P0, issue the command:

   STRSBS SBS(*controlling_subsystem_name*)

   then resume normal operations in this partition.

3. On the Enterprise Storage Server, connect the LUNs B1"-B7' to the secondary partition P1 (step **3**).

4. In secondary partition P1, force DST and recover the remote load source, as described in 11.2.1, "Recovering the Remote Load Source" on page 310. This is shown as step **4** in Figure 14-1.

5. After recovering the LSU for P1, the partition will IPL (step **5**) and make the target LUNs from the FlashCopy available to P1.

6. You can now perform backups of the copy of data created from P0 by the FlashCopy (see step **6**). Because LPAR supports dynamic re-allocation of CPU and memory resources in V5R1, partition P1 will use some of the cycles and memory previously available to P0 (shown as 7 in Figure 14-1) to do the backups. The amount of these resources used can be controlled when you initially define the logical partitions.

7. After the backups have completed, it is necessary to set P1 back to the state ready to restart the process again for the next day's backups. To do this, force DST in P1 and perform a D-IPL. Now you are ready to start again for the next cycle of backups, or perhaps for another logical partition on the same system.

14.2 FlashCopy cloning considerations

Due to the fact that FlashCopy creates a duplicate of the entire DASD configuration on an iSeries, everything defined on the source partition is duplicated on the target partition, with only a few exceptions (such as some system values relating to the hardware configuration *and the hardware resource information.* Consequently, it is reasonable to consider the target partition to be a clone of the source partition after the FlashCopy has taken place. Because of this, you should be aware of the following operating considerations.

14.2.1 Job schedule entries

The entire job schedule will be duplicated, whether you use the standard OS/400 scheduler, the IBM Advanced Job Scheduler or another third party scheduler. To avoid jobs which should be run on the source partition running on the target, you should ensure that the job scheduler does not run automatically when you start the target partition. Conversely, you should define the jobs to be run on the target partition to be held.

Standard OS/400 job scheduler

If you are using the standard OS/400 scheduler, you should use the following command on the source system/partition immediately before bringing the partition into a restricted state. This will ensure that all jobs are held when the "clone" partition is started.

```
HLDJOBSCDE JOB(*ALL) ENTRYNBR(*ALL)
```

When you restart the source partition after going to the restricted state, you should release the jobs which were held, using the RLSJOBSCDE command.

> **Attention:** If there are any entries in the schedule which were to be run on the target partition, you should take care that these jobs are not released otherwise they will also be run on the source partition.

After starting the target partition, you should release any entries to be run on the target (for example, your backup jobs). It is likely that you would select these jobs manually, ensuring that the other jobs intended to run on the source partition are not released. Use the RLSJOBSCDE command to do this. If you have a regular set of backup jobs, you could use a CL program to release them rather than having to release each one individually.

IBM Advanced Job Scheduler for OS/400

If you use the IBM Advanced Job Scheduler for OS/400 you can use the following command on the source partition to stop any jobs being submitted from the schedule immediately before bringing the system to a restricted state:

```
ENDJS OPTION(*MONITOR)
```

When you restart the source partition, you should restart the monitor using the command:

```
STRJS OPTION(*MONITOR)
```

Depending on how your schedules are setup, you can use the WRKJOBJS screen and select the jobs to be released. Alternatively, it may be better to define your jobs as members of an Application and use the command

```
WRKJOBJS APP(application-name)
```

to identify the jobs to be released. You should use a similar technique on the target partition if you have any backup jobs scheduled to be run.

14.2.2 Hardware resource names

It is highly unlikely that the hardware resource names associated with the tape drives in the source partition will match those on the target partition. In this case, you should either change the device descriptions after you have done the FlashCopy or create a CL program to do this for you.

14.2.3 Last save date/time

When an object is backed up, the object header for that object is updated with (amongst other things) the date and time of the save. Due to FlashCopy being a one-way operation (source to target), these updates are not reflected back onto the source copy of the data. Consequently, when the source is next copied to the target using FlashCopy, the date and time of the last save is not available for incremental (SAVCHGOBJ) saves.

However, this is not really an issue, as the reason for using incremental backups is to minimize the downtime of the production partition and this is not an issue when taking backups off a copy of the data. Indeed, taking full backups instead of incrementals is advantageous as when recovering, only a single restore is required rather than the more usual method of restoring the latest full backup followed by the most recent incremental(s). Doing a full save and restore will minimize the recovery time after a failure.

14.3 Using BRMS with FlashCopy

Backup Recovery and Media Services for OS/400 (BRMS) is IBM's strategic solution for performing backups and recovering AS/400 and iSeries systems. BRMS has a wealth of features, including the ability to work in a network with other AS/400 and iSeries systems to maintain a common inventory of tape volumes.

Using BRMS with ESS FlashCopy provides a challenge as FlashCopy actually creates a clone of the source system onto a second set of disk drives which are then attached and used by another system (or LPAR partition). In this chapter, we will explore how BRMS can be used to perform backups and recoveries from a secondary LPAR partition. This could also be a separate stand-alone system. However, using the dynamic resource movement introduced in V5R1 and later of OS/400, the LPAR solution is the best way of using FlashCopy when attached to an iSeries platform.

14.3.1 BRMS architecture

BRMS stores it's tailoring and management information in a library called QUSRBRM. The files in this library define both the setup of the BRMS environment and the dynamic information gathered as a result of doing BRMS operations such as save and restore tasks. This information is crucial to the recovery of the system.

When using FlashCopy, QUSRBRM is cloned from the production system to the backup system.

In Figure 14-2, we show two partitions, Production for normal day-to-day processing and Backup for taking "offline" backups. These need not be primary and secondary - they could both be secondary partitions.

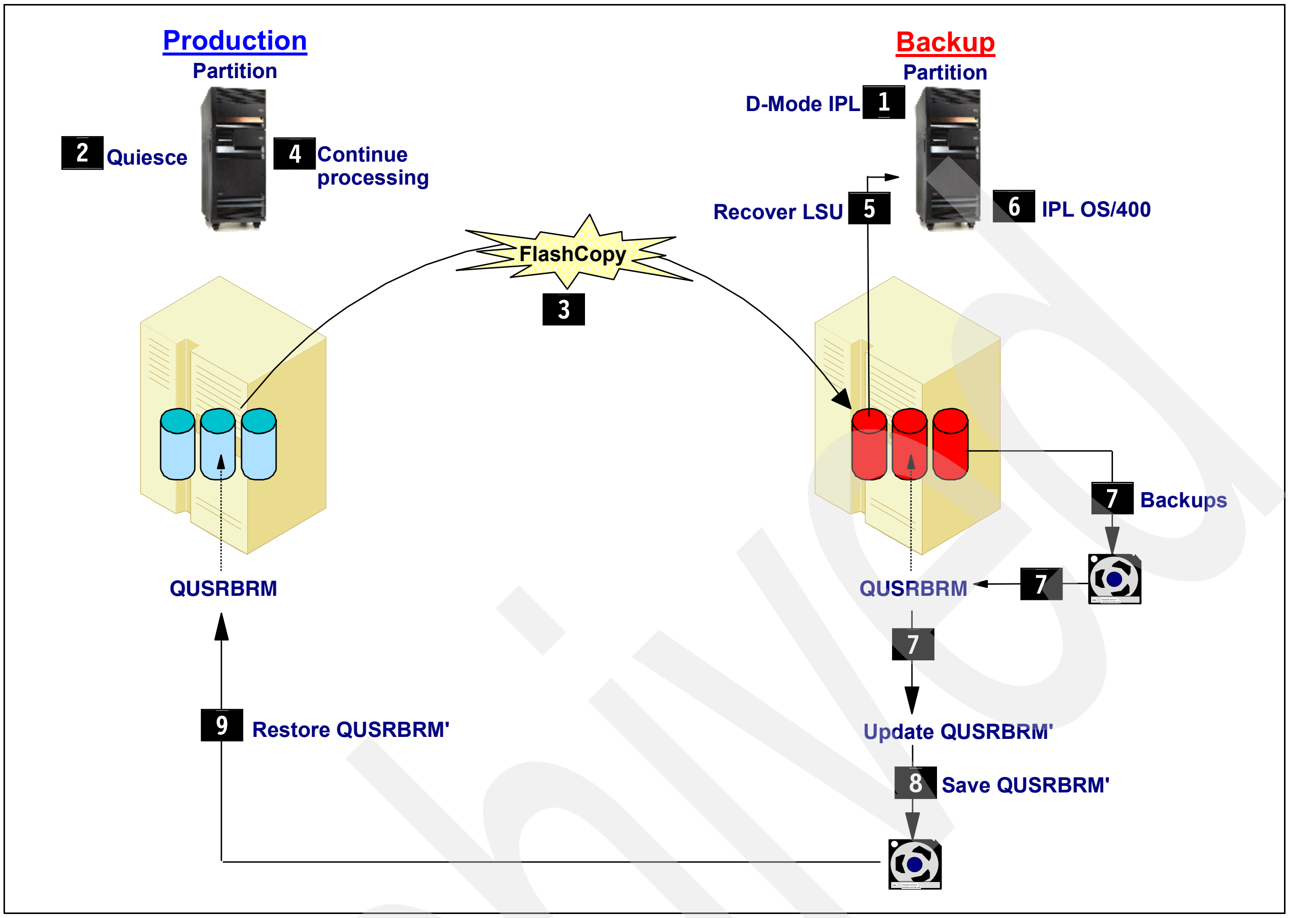

Figure 14-2 Using FlashCopy and stand-alone BRMS

14.3.2 Example of performing backups in the clone partition

In this example, we do not have a BRMS network. In other words, our production partition is standalone and does not share any BRMS information with other systems or partitions.

The steps necessary to do the BRMS backups on the second partition are shown below. They correspond to the steps shown in Figure 14-2.

1. Prepare the secondary partition for the backups. This involves performing a D-mode IPL, as described in 11.2.1, “Recovering the Remote Load Source” on page 310.

   You will most likely find it best to do this step first, although you could do it after step 2. However, the advantage of doing this step now is that it can be done well in advance of starting the backups and therefore is done when it suits you, rather than at a busier time when the backups need to be started. Once the partition is IPLed to DST, you can leave it in this state.

2. When you are ready to start the backup process, you should quiesce the Production partition to ensure that any objects still in memory are written to disk so that they can be copied with FlashCopy. To do this ensure all users are signed off and all batch jobs have ended, then run the command:

```
ENDSBS SBS(*ALL) OPTION(*IMMED)
```

3. When the Production partition has reached Restricted state, you should invoke the FlashCopy from the Copy Services screen in StorWatch as shown in Chapter 13, “FlashCopy” on page 367.
4. Once the FlashCopy task has completed, you are ready to continue processing on the Production partition. To do this, issue the command:

   STRSBS SBS(*controlling_subsystem_name*)

   > **Important:** You have just created a clone of the entire single level storage for the Production partition. This includes the QUSRBRM library which contains all the BRMS definitions and management information. You should ensure that no BRMS activity which will update information in QUSRBRM takes place on the Production partition until the backup process is completed in Step 9 on page 386. Failure to ensure this could cause discrepancies in your BRMS database and cause unpredictable results. Examples of BRMS activity which would cause QUSRBRM to be updated are:

5. You should now recover the Remote Load Source Unit from the mirrored copy in the ESS. Follow the instructions described in Chapter 11, “Load Source Mirroring in the ESS” on page 303.
6. When you have successfully recovered the LSU you will see a screen similar to Figure 14-3. You are now ready to IPL OS/400.

   Before you IPL the Backup partition, you should ensure that the IPL Source is “B” as the Recover Remote Load Source function usually leaves it set to the “A” side.

```
Disk Configuration Information Report


 The following are informational messages about disk

 configuration changes started in the previous IPL.


 Information


 Recover mirrored load source completed successfully.


 Press Enter to continue.
```

Figure 14-3 Remote LSU recovered successfully

7. When the IPL has completed, the Backup partition is an exact clone of the Production partition as it was when it was quiesced in Step 2 on page 384.

You are now almost ready to perform your BRMS backups in the same way as you would have done on the Production partition.

> **Important:** It is likely that the resource names on the Backup partition do not match those on the Production partition and also that new device descriptions may be created if you have the System Value AUTOCFG=*YES. Before you start the backups, you should check that the required resource and device description names are available. If you are using an Automated Tape Library, it is particularly important that the same device names are available on the Backup partition as are on the Production partition.
>
> Further information about how to rename device descriptions and resource names can be found in Chapter 10, “Using Fibre Channel attached tape devices” on page 247.

When the backups are finished on the Backup partition, do not run BRMS Maintenance (STRMNTBRM) or any of the activities which are run as part of maintenance such as Media Movement. You should run these functions on the Production partition when you have finished Step 9.

8. When all the BRMS backups have finished, you must save the QUSRBRM library to allow the BRMS management information to be transferred to the Production partition. This is best done to a savefile.
9. The final step in the process is to restore QUSRBRM which you just saved from the Backup partition back onto the Production partition. This will provide an accurate picture of the BRMS environment on the Production partition which reflects the backups just performed on the Backup partition. To do this, take the tape used in step 8 and issue the following command on the Production system:

```
RSTLIB SAVLIB(QUSRBRM)
   DEV(tape-device-name|tape-media-library-device-name)
   VOL(XFRxxx)
   ENDOPT(*UNLOAD)
   ALWOBJDIF(*ALL)
```

> **Note:** If you do not plan to restore individual file members from your normal backups or you do not use the BRMS Archiving function, you could omit step 8 and restore the objects saved from QUSRBRM at the end of your backups. These are either saved automatically at the end of the last control group or explicitly by the SAVMEDIBRM command at the end of the backups. In this case, you should use the command
>
> ```
> RSTOBJ OBJ(*ALL)
>       SAVLIB(QUSRBRM)
>       DEV(tape-device-name|tape-media-library-device-name))
>       VOL(XFRxxx)
>       MBROPT(*ALL)
> ```
>
> However, you should note that if you subsequently start to use BRMS Archiving or require to restore individual file members, these functions will not be possible unless you follow the instructions in Steps 8 and 9.
>
> In addition, if there is a backup of library QUSRBRM on the tape containing the Media Information, you must not specify the default SEQNBR(*SEARCH) on the RSTOBJ command this would restore the QUSRBRM library save and not the latest Media Information objects.

14.3.3 Considerations when using BRMS in a "cloned" partition

Because the Backup partition is an exact clone of the Production partition, everything is copied over with FlashCopy. This introduces a number of areas of which you need to be aware. In addition to those considerations discussed in 14.2, "FlashCopy cloning considerations" on page 381 you should bear the following in mind:

*MEDCLS
: If you use *MEDCLS to select devices for a BRMS media operation, some of the device names may not exist on either the Production partition (in the case of restore operations) or the Backup partition for save operations.

Dynamic Retrieval
: If you use the BRMS Dynamic Retrieval function, you should be aware that if an object is retrieved while the "clone" partition is active, the information gathered by the retrieval function will be lost when QUSRBRM is restored in Step 9 on page 386 (unless you use the RSTOBJ option described). This information relates to length of time an object which has been dynamically retrieved will remain on the system if it is not changed. If you restore the entire QUSRBRM, this information will be lost.

Incremental saves
: As mentioned in 14.1, "Steps required to take backups" on page 380, using the "clone" method prevents the use of incremental saves. However with BRMS, some functions use the BRMS database to identify when an object was last saved and do not use object header. In these cases, BRMS would perform a *INCR save if specified. However, an incremental save is less beneficial as doing this in the Backup partition has no significant impact on the Production partition and recovery will be much simpler from a *FULL backup.

Maintenance
: As mentioned in 14.3.2, "Example of performing backups in the clone partition" on page 384, we recommend that no maintenance tasks should be run in the Backup partition.

A

Comparative discussion

In this appendix we compare the implementation of ESS Copy Service on iSeries and other Open System platforms, and the similarities and differences of these platforms.

Dispelling iSeries differences

The use of PPRC and FlashCopy with AS/400 (iSeries) is very similar to Unix and NT systems from a relational database point of view. What we mean by this, is that in Unix or NT, all database file systems need to be copied to have a working copy of a database. In the event of one LUN not being copied, the entire database is unusable, as it is with AS/400 (iSeries) Single Level Storage Architecture.

The difference's between AS/400 (iSeries) and other Open Systems revolve around the need to recover the Load Source and all volumes associated with a load source. However we can treat the Single Level Storage Architecture as a relational database such as DB2® or Informix® on other Open Systems, and the major differences begin to diminish.

In conclusion, for those out there who may be unfamiliar with AS/400 (iSeries) Single Level Storage, it may be useful to relate back to a database file systems in other Open environments.

iSeries and UNIX equivalent

We have taken some of the common iSeries terms and cross referenced them to some of the possible UNIX equivalents. Due to the number of UNIX variants, we have attempted to make the comparison as generic as possible.

QSECOFR

QSECOFR is an all rights user ID, the leading Q denotes an IBM object in OS/400. This is a standard IBM Security Officer user ID. The equivalent user ID for QSECOFR within Unix would be the super user account, commonly referred to as root.

LPAR

LPAR on iSeries is the dynamic logical partitioning of resources to run a separate copy of the Operating System. Within Unix, there are many variations of logical partitioning that occur at various levels and configuration. Refer to Table 14-1, LPAR Comparison, for platform specific information.

Table 14-1 LPAR comparisons

	iSeries	SUN	HP	Unisys
Operating Systems	OS/400	Solaris	HP-UX	Windows 2000
Supported Models	6xx, Sxx, 7xx, 8xx	E10000	L-, N-class, Superdome	ES7000
Partitioning Type	Logical	Physical (Board)	Physical now, virtual coming	Physical
Maximum # partitions	32	16	16	8
Minimum # processors / partition	0.1	1-4	4	4
Processor Increment above minimum	0.01	1-4	4	4
Memory Increment above minimum	1 MB	512 MB	512 MB	8 GB
I/O Increment above minimum	IOP	2 Buses	Bus	Direct I/O Bridge
Independent Movement of Resources (CPU, Memory, I/O)	Yes	No	CPU only	No
Dynamic Resource Movement	Yes	Partial (database restart required to use)	Planned for 2001	Yes
High-speed internal communications	Yes	Yes (via routing partition)	No	via shared memory

Single level storage

Both the main memory and physical disk units are treated as a very large virtual address space, known as Single Level Storage. This is probably the most significant differentiation in a SAN solution implementation on the iSeries when compared to other systems. It is this architecture that masks the complexities of adding additional disk units to the system. Refer to 1.2, “iSeries storage architecture” on page 4, for a more detailed discussion of Single Level Storage.

Within Unix, Logical Volume management software can be used to create volume groups to accomplish similar tasks. These volume groups are then referenced as a single address space, however unlike on iSeries, memory is not included in these virtual addressing spaces. Figure 14-4 shows a depiction of how multiple DDMS can be configured within Unix to more closely resemble single level storage.

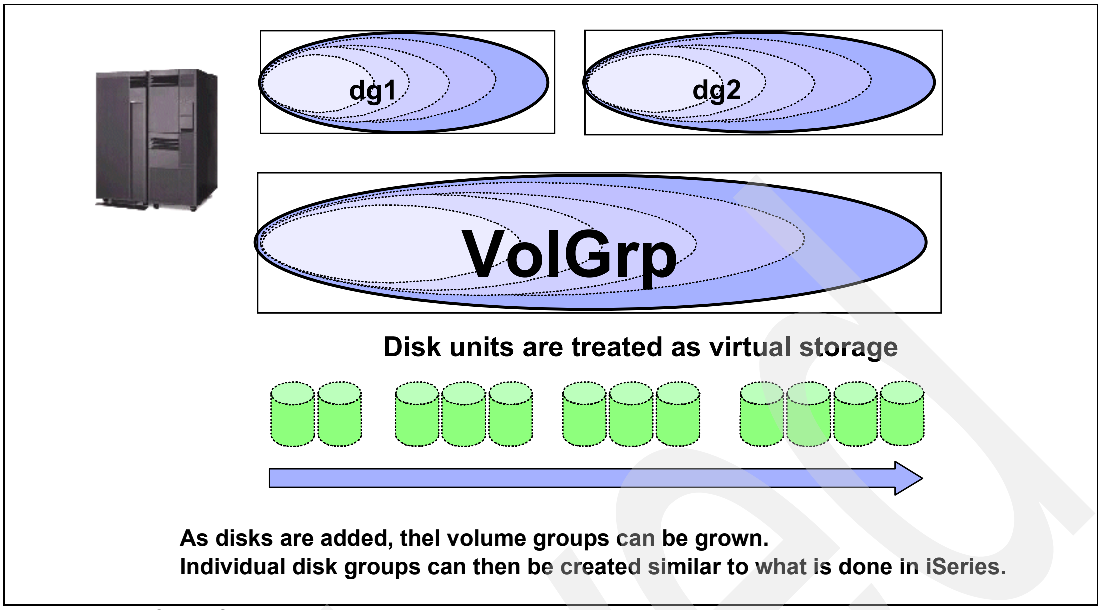

Figure 14-4 Open Systems Logical Volume Management

Auxiliary Storage Pool

Within iSeries single level storage, you have Auxiliary Storage Pools (ASPs) which are similar to disk groups, which a Logical Volume management on Unix creates. Within an ASP, OS/400 will evenly strip data across all DDMS within the ASP. As new DDMS are added to an ASP, the data can be redistributed dynamically to include the newly added DDMS. Within a Unix Logical Volume, striping not disk concatenation, can also be configured to more closely represent an ASP. However this is an operation of the Logical Volume Management software, and not the Operating System. Referring back to Figure 14-4, you can see how individual disk groups can be created within a volume group. This would be similar to the creation of ASP's within the Single Level Storage addressing space.

Dedicated Service Tool (DST)

DST is a set of tools used to re-configure various aspects of the system, and handle problems. To get to DST, you would normally need to do a manual IPL of the system, and the first panel that the system stops at is the DST panel. Examples of things you can do in DST is adjust your disk configuration, start disk mirroring, various diagnostic items, install new versions of the OS, create new LPAR partitions, continue the IPL step by step, etc. Not all task within DST will impact applications and end users, so it is possible to force DST without performing an IPL of the system or LPAR. However only a subset of DST functions are available.

Within Unix, the DST equivalent would be a number of applications or functions run at various system run levels or by using System Management Interface Tool (SMIT). Not all functions would necessarily require the use of the root user, some of the tasks would include volume management, adding additional DDMS, file system creation to include mounting and unmounting.

System Service Tools (SST)

SST is a subset of the DST, that can be done while the system is up and running. It includes adding new disk (but not removing old disk which requires DST), manage of LPAR partitions (but not create them, since this requires DST). Unix equivalents would be a System Management Interface Tool (SMIT) such as described in "Dedicated Service Tool (DST)" on page 392. SMIT, provides the feature functionalities of SST, but is more of a combination of DST and SST functionalities.

Initial Program Load (IPL)

Within iSeries there are a number of IPL modes, as there are boot options within UNIX. We have provided a reference between the two in an attempt to make things clearer for both iSeries and Open Systems types alike.

- A-mode

  An A-mode IPL boots from the alternate copy of the operating system that only contains the Program Temporary Fixes (PTFs) that have been made permanent.

  AIX has this equivalent funtionality. Sadly, other UNIX variants would require more creative thinking, such as creating an alternate boot disk prior to applying OS patches or updates. Then by booting off of the alternate boot disk, you would in essence have nearly similar functionality as an A-mode IPL.

- B-mode

  This is a normal boot of the system from the main copy of the operating system.

- C-mode

  Mainly used by CEs for problem isolation or recovery type tasks. A Unix equivalent would be diagnostics or recovery tasks done at a boot prom level prior to the kernel being loaded into memory.

- D-mode

  Boot from an external source such as a CD or a tape that has a *SAVSYS on it (copy of the OS). Similarly within Unix, this can be accomplished with the use of the installation media (CD), or even with a tape containing a dump of the OS.

Along with the different IPL modes, there are also different IPL types that can be used in conjunction with the modes.

- Normal

  This is a normal boot of the operating system to include the start of applications on the system.

- Manual

  A Manual IPL will boot the system to the DST screen, from here choices are made to continue the IPL or perform system level tasks not normally done while applications and users are active on the system. Booting to single user mode within Unix would be equivalent, as to restrict the use of the system to the root user account on the system console.

- Abnormal

  An abnormal IPL occurs following an improper shutdown of the applications and OS. Examples would include hardware or a power failure. Within Unix, this would be a system recover boot requiring file system checks prior the file systems being available for use.

High Speed Link (HSL)

This is a external 1 GB/s bus that is available in the 270 and 8xx ‘iSeries’ CPUs. It has made a tremendous difference to the iSeries tape save/restore rates by increasing the bus/IOP bandwidth which was a bottleneck for I/O operations in the past. It also allows the addition of more DDMS while maintaining performance, as well as the addition of new functions such as switchable disks and the use of fibre cards.

B

iSeries availability

In this appendix we discuss the availability characteristics and functions available to iSeries servers. These functions are typically built into the operating system and supplied as base functions or optional features.

High availability versus disaster recovery

Increasing demands on systems are putting enormous pressure on systems managers to maintain a thoroughly protected, stable and fully recoverable environment. With more and more customers consolidating their workloads from multiple systems onto fewer, larger systems, and with the advent of e-business, these demands are set to increase even further. Most businesses require a high, if not continuous, level of information system availability. A lengthy outage can result in a significant financial loss.

Complete and integrated system availability functions are a necessity. OS/400 provides many availability tools and options, meeting the varied requirements of different systems environments. Customers can choose to implement some, or all, depending on the demands made on their systems, and the resources available to them.

System downtime is caused by one of two types of outage:

Unplanned outages These are the result of some sort of a failure, and cannot be predicted or scheduled. Normal system operations are interrupted without warning. They can effect the entire system or only a portion thereof.

Planned outages These may be the result of a failure but are more commonly due to regular actions such as backups. Planned outages are scheduled in advance and can effect either all or just a portion of the system.

If there is an outage when the system was due to be up and running, then availability is affected. Conversely, if there is an outage but the system was not due to be up and running, that is not usually regarded as affecting system availability.

It is important to differentiate between the various levels of availability and disaster recovery (DR). Figure B-1 shows the varying levels of availability and how they are affected by planned and unplanned outages.

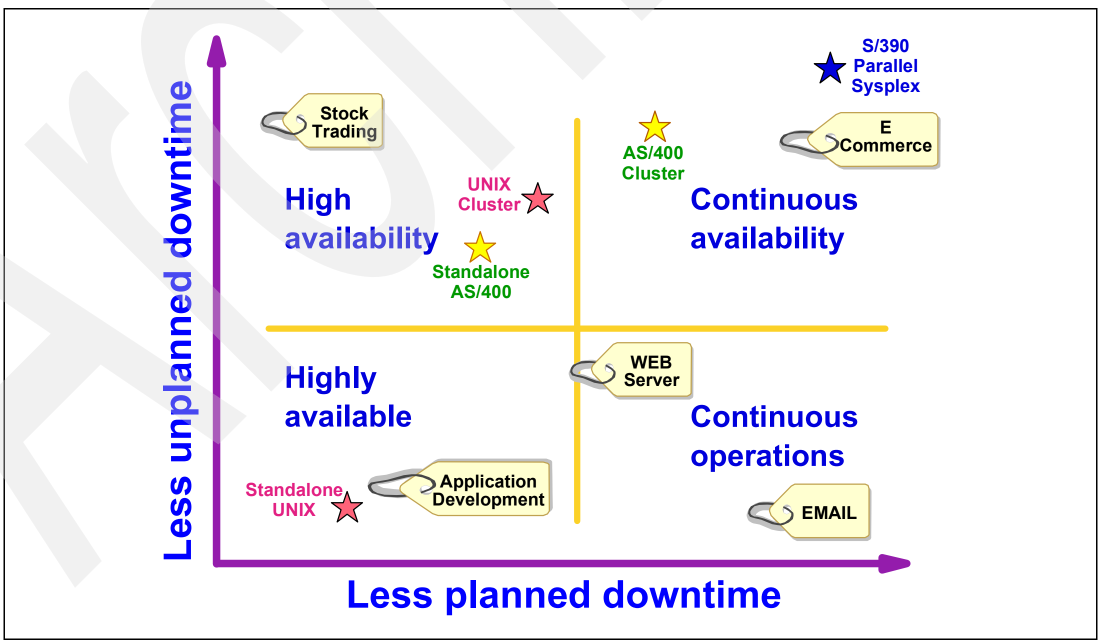

Figure B-1 Varying levels of availability

Highly available A highly available system is one that experiences both planned and unplanned outages. However, the time needed to recover from the outage is usually less than or very close to the duration of the outage.

High availability In this environment, it means that you have protected yourself from almost all recovery actions required due to an unplanned outage. It does not mean that unplanned outages do not occur. As designed, the iSeries is a very reliable system. However, due to requirements to take the system down for h/w and s/w maintenance as well as daily backups, there may be a fairly high proportion of planned downtime.

Continuous operations In a continuous operations environment, there are no planned outages. The system and applications are available to users when the users require them and there are no outages for things like upgrades, backups, or other systems maintenance act ivies. However, unplanned outages can still occur. Usually these outages are of short duration and recovery actions are unnecessary or minimal. Implementing database replication can avoid most planned outages. However, there may still be some unplanned downtime due to occasional hardware or software failure and a short period of time while cutover to the second system.

Continuous availability Continuous availability is very similar to continuous operations, however, there are *never* any outages of any type. When continuous availability is implemented, it becomes imperative that it is done not only at a system level, but also an application level because it means that the user will not lose any transactions or even notice that a failure or outage has occurred somewhere within the computing environment.

Disaster recovery is more difficult to define. Each organization may have a different view of what constitutes a disaster. However, it is generally accepted that it has a significant impact on the business. For example, it could be as major as a complete site loss or simply the loss of a strategic system. DR usually involves to ability to recover from a “worst case” scenario and recovery to an alternative system will take longer than any of the availability definitions described above.

To increase the level of availability towards 100%, it is necessary to employ additional techniques, as shown in Figure B-2.

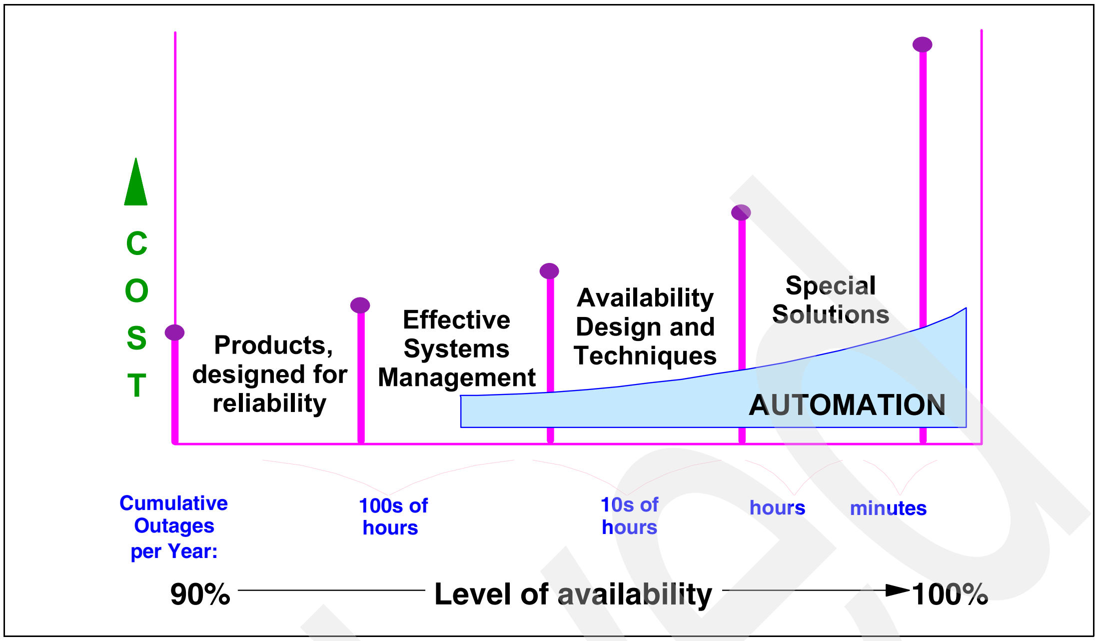

Figure B-2 Achieving the various levels of availability

The highest levels of availability can only be achieved with a second, live system. Simply duplicating data on disk drives will only provide you with a solution for disaster recovery, such as might be required to recover after a site disaster. Actually making this data copy usable usually takes longer than is available to achieve the levels of availability defined above.

IBM's strategy for high availability and beyond is based on data replication. See Figure B-3. This is where the data from the primary, production system is replicated to a second live system. In the event of downtime (planned or unplanned) users can switch to the second system and continue processing the data.

To reach the levels of continuous operations and continuous availability, even having a replica of data on another live system is not enough. These require *application resilience* and not just data resilience. Such a solution can only be provided if the application on the primary system is "in touch with" the application on the second standby system. *Clustering* provides the tools to build these solutions.

Management must determine the time frame that moves an outage from a "problem" to a "disaster" status. Most organizations accomplish this by performing a business impact analysis to determine the maximum acceptable downtime for critical business functions. Clearly, the cost of providing such solutions must be taken into account and contrasted with the potential cost of the downtime.

Hardware mirroring is implemented in the DASD subsystem (such as the Enterprise Storage Server) and is not recognized by the host system. It provides a copy of data on another set of disks. However, due to Single Level Storage in the iSeries, not all data can be guaranteed to be on disk - some will be in memory, and therefore the data on the copied disks after a system failure will be incomplete. When using these disks for recovery on a different system, an IPL is required. This will be an "abnormal" IPL and could take a long time. For this reason, "non-native" mirroring should only be considered as a solution to disaster recovery and not to high availability.

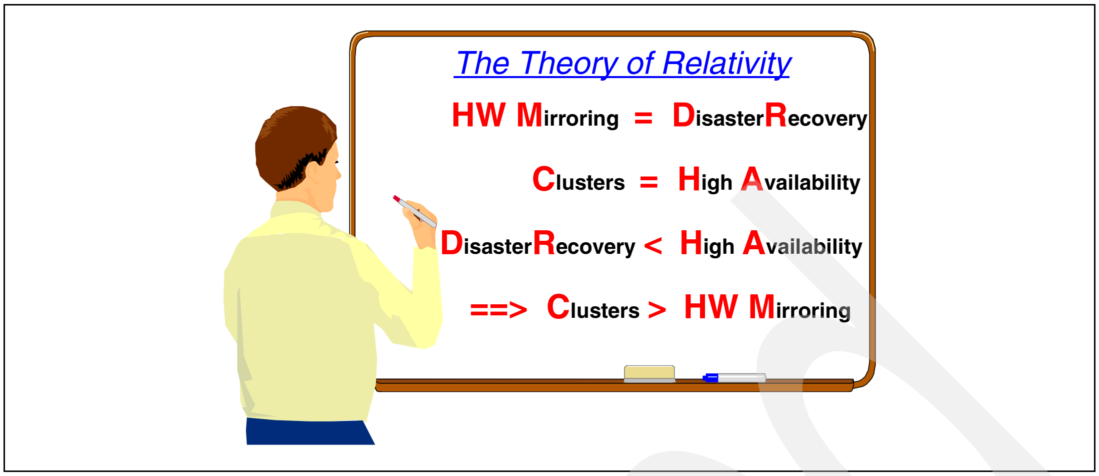

Figure B-3 Hardware mirroring versus disaster recovery

However, there area number of situations where an external hardware mirroring solution is perfectly satisfactory. Such as solution can be provided by the Enterprise Storage Server and ESS Copy Services. This is discussed in Chapter 12, “Peer-to-Peer Remote Copy” on page 327.

Using OS/400 availability techniques

There are a number of availability functions within OS/400. These provide varying degrees of availability and you should choose the most important to your organization.

Journaling

This is a database logging facility which can record after-images of any database records which are changed. It does not require changes to the application software. Optionally, you can also capture before-images of these records. If the database is restored, the after-images can be “rolled-forward” to bring it up to date. Journaling Access Paths (file indexes) can prevent significant Access Path rebuild time after a system failure. Another major advantage of journaling is that each write to the journal is automatically written out to disk, whereas an update to a database record may only be done in main storage until that page is written out to disk some time later.

Irrespective of which solution you choose for High Availability or Disaster Recovery, it is imperative that you should *use OS/400 journaling as a minimum* in order to ensure that any updates are written to disk. In practice, the actual database updates will likely still remain in main memory in the event of a system failure but the journal changes will be on disk. Any recovery will be longer if journal changes are not available at abnormal IPL time. This is particularly important when looking at recovery scenarios based on hardware mirroring such as PPRC.

Commitment Control

This extends journaling and uses before and after images to enable roll-forward or roll-back to be synchronized to a user's transaction boundary. It must be designed into the application software.

System Managed Access Path Protection

System Managed Access Path Protection (SMAPP) is a system-wide facility for minimizing the time to rebuild Access Paths after a system failure. For those customers who do not (or only partially) implement Access Path journaling, this is the minimum level of access path protection which should be used.

Auxiliary Storage Pools

Auxiliary Storage Pools (ASPs) allow you to “partition” the disks attached to an iSeries. Up to16 User ASPs can be used. If all DASD are in a single unprotected ASP (default), then a failure of one disk would require the entire ASP to be restored. Even if RAID-5 or Mirrored protection is implemented, it is advisable to partition the DASD into multiple ASPs if you have a large DASD capacity (for example, greater than 350 GB) to avoid long restore times.

With the Enterprise Storage Server, we would not recommend that you should use ASPs as an aid to availability if *all* storage units on the iSeries (apart from the Load Source unit) are in the Enterprise Storage Server. However, if you have additional internal DASD, you may consider having the internal DASD in one (or more) ASP, and the external DASD in the ESS in a different ASP. Also, if you have different sized LUNs in the ESS, you may want to use ASPs to keep equal-sized LUNs together.

Hardware features

In addition to the facilities provided by OS/400, there are facilities provided by hardware. These generally involve providing some degree of hardware redundancy.

Device Parity Protection (RAID-5)

Device Parity Protection, more commonly known as RAID-5, is a way of protecting a system from an outage caused by a single disk drive failure. Through a form of parity checking across a set of disk drives, the drives can continue to function despite a disk failure, recreating data from the failed drive as it is requested. OS/400 controls the implementation of RAID-5 across internal disk subsystems. In the case of a drive failure, the system will continue to perform normally while the operator is informed of the condition. The Enterprise Storage Server also implements RAID-5 protection in a similar way to the iSeries internal DASD. However, a failure of a disk drive in the ESS will not be notified to the iSeries operator in the way that internal drives failures would but instead will be logged by the StorWatch control system. StorWatch can also “phone home” to IBM Service in the event of a problem, or it can send an SNMP trap to the iSeries if this function is enabled.

Concurrent maintenance is provided for the latest model RAID-5 drives on the iSeries, allowing maintenance, and even drive replacement, to be completed without requiring a dedicated system.

RAID-5 is the minimum level of DASD protection recommended.

DASD Mirroring

iSeries native DASD mirroring provides protection from a system outage caused by a disk drive, disk controller, disk I/O processor, or even system bus failure, by optionally duplicating each of these items.

Once mirrored protection is started, data will be written to two disk units of a mirrored pair for each write operation. If any disk drive fails, the system will continue to use the other drive in the mirrored pair. All protection provided by DASD mirroring is visible to OS/400.

Remote Load Source Unit (LSU) Mirroring

Remote load source mirroring support allows the two disk units of the load source to be on different IOPs or system busses, which provides IOP or bus level mirrored protection for the load source. This can allow the LSU to be on an optical bus that terminates away from the main system unit, so that the entire system, including the load source, can be protected from a site. See Chapter 11, "Load Source Mirroring in the ESS" on page 303 for further details about how this can be implemented with the Enterprise Storage Server.

‘Non-native DASD mirroring

At first look, a solution which provides non-native hardware mirroring (in other words, done by the DASD subsystem rather than the host) to another set of DASD remote from the AS/400 might seem a good choice for High Availability. The Peer-to-peer-remote-copy (PPRC) feature of the Enterprise Storage Server falls into this category. However, looking closer at such a solution soon brings out some areas for consideration:

Asynchronous versus synchronous

Depending on the delay in the network in sending the updates, the remote mirrored set of DASD will be out of sync with the primary set. If a synchronous solution is proposed, then the response time of any transaction will be delayed due to the synchronous update. PPRC only uses synchronous communication.

All or nothing

With a hardware solution, the mirroring is ‘non-selective’. Either the entire DASD contents will be copied or none at all. Contrast this with software replication (such as OMS, MIMIX or Data Mirror) where you can selectively choose which parts of the system are to be replicated and the performance of the primary system is not adversely impacted (remember, journaling is necessary, whether you are doing hardware mirroring or software replication so there is no additional overhead).

Recovery after failure

Because the non-native mirrored DASD are not directly attached to the primary iSeries system (and hence not owned by it) they must be attached to another system before they can be used. This process must involve an IPL, and due to the fact that the primary system termination was abnormal, then the subsequent IPL will be abnormal and will take much longer than a normal IPL. Contrast this with a software replication solution where the secondary system is already up and running and no IPL is necessary.

Conflict with native OS/400 functions

OS/400 and iSeries microcode do not recognize the concept of non-native hardware mirroring and so their in-built recovery functions cannot be assumed to work. Contrast this with software replication where both primary and secondary are using standard OS/400 availability facilities and are used at IPL time.

Software maintenance

After a failure with non-native DASD mirroring, the secondary system is using the mirrored DASD from the primary system. Consequently, the system software level on the secondary system will be exactly the same as the primary system. This may look like it allow you to minimize maintenance of system software as the DASD containing the system software is mirrored. Contrast this with software mirroring where both systems can be at different software levels (indeed, different versions), thus allowing more planned introduction of a new version of system software and hence help reduce planned downtime due to maintenance and allowing an easier fallback should you want to revert to the earlier software level.

Replication software

IBM's High Availability Business Partners (HABPs) Lakeview Technology (MIMIX), Vision Solutions (OMS) and Data Mirror provide solutions to replicate one iSeries system onto another. This has the advantage of being able to provide a hot-standby system, remote from the primary system, allowing switch-over in a matter of minutes. The level of replication can range from partial database to the entire system, including user profiles and other system objects.

Additional benefits can be gained by using the second for inquiry-only transactions such as data warehousing or for taking backups without significantly disrupting the service to users.

Replicating the data is the first level towards highest availability.

Clustering

A further level of resilience can be achieved by building on the dual systems approach outlined above. Splitting the applications and database onto servers with different performance characteristics (batch versus interactive) and using software replication on the database server to another database server (and possibly the application server to another traditional server) can provide further redundancy into the overall solution. Add to this the support provided by clustering solutions, allows *application resilience* rather than just data resilience.

Clustering provides the highest levels of availability for Iseries systems.

Save-while-active

When taking a backup, it is usually necessary to have no active jobs running on the system, or at least that no updates are being done against the database. This is to ensure that no objects are locked and therefore cannot be saved. However, many customers now find that the time window available for doing backups is being reduced and they do not have sufficient time to perform backups using this traditional method.

The save-while-active (SWA) function on the iSeries is an option on several of the save commands. It allows you to use the system during all or part of the backup process, thus helping reduce the planned outage time while backups are being taken. It allows you to modify objects during most of the time that they are being saved. In contrast, other save functions provided on the system allow no access to the objects as they are being saved or only allow the objects to be read as they are being saved.

There are two ways in which SWA can be used.

"True" Save-while-active (Advanced method)

This method involves actually saving the objects as they are being updated. Using this method, the time-relationship between objects is not maintained. For example, object A in library A might be saved first, then some time later, object Z in library Z would be saved. During the time between A and Z being saved, object Z is updated, losing the integrity between the two objects. In order to recover from such a save and maintain the integrity of the database, it is necessary that, as a minimum, you have journaling running when the save is taken and preferably commitment control. In the event of having to restore from the backup, the objects would initially be restored *exactly as they had been saved* (including the time differences between objects A and Z being saved) and then you must backout the updates to a known pint of integrity, using the journal. Clearly, this can be a difficult and time consuming task, and is only made simpler if commitment control is included in the application design so that changes can be backed out to a commitment boundary.

"Save-after-checkpoint" (Recommended method)

The recommended way to use SWA is to quiesce the system so that no updates are being done to the database until a checkpoint has been established. Then, once the checkpoint has been reached, you can safely allow users back on the system to continue making updates to the database while it is being saved to tape.

The system performs the save-while-active function by maintaining an image of the object being saved as it existed at a single point in time. As the object is being changed by an application during the save, the system maintains an original copy of the pages of the objects that are being changed. In virtual storage, a page is a fixed-length block that has a virtual address and is transferred between main storage and auxiliary storage. It can be viewed as though the system is maintaining two images of an object as it is being saved: one image that contains the updates to the object that normal system activity works with, and a second image that is an image of the object from a single point in time that only the save-while-active job is using to save the object to the media. The system does not maintain two complete images of the object being saved. It only maintains two images for the pages of the objects that are being changed as the save is being performed.

The checkpoint image of an object is the image of the object at a single point in time. The checkpoint for an object is the instant in time that the checkpoint image of the object is taken. Checkpoint processing is part of the work that is performed by the system during the save preprocessing phase of a save-while-active operation. Other work that is performed by the system during save preprocessing includes determining which objects are to be saved and locking those objects to be saved. The save post-processing phase of a save operation includes unlocking the objects that were saved.

Figure B-4 shows how SWA works using the recommended "quiesce" method. The save process is split into three phases - pre-processing when a checkpoint is taken, the actual save to tape and the post-processing tidying up.

Let us assume for simplicity that there is only one object to be saved using SWA. In practice, there will be many objects in many libraries but the principle remains.

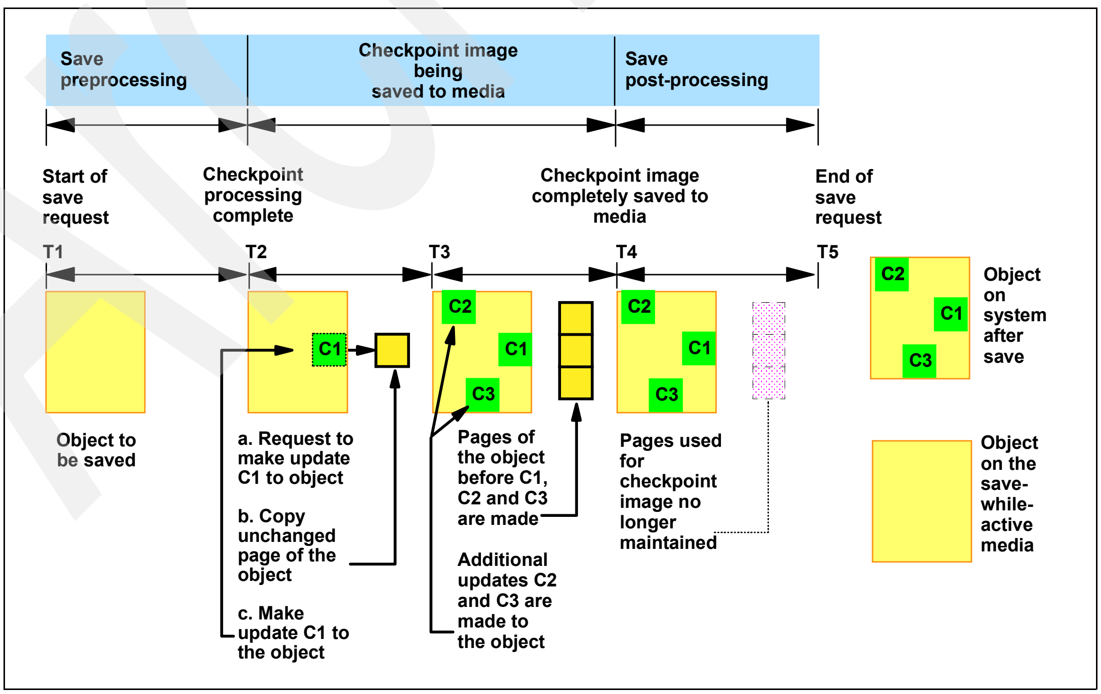

Figure B-4 Save-while-active

1. At time T_1, freeze the object and establish a checkpoint. There should be no jobs performing updates at this time. At time T_2, the object has reached a checkpoint and users can be allowed to begin making updates because the state of the object is known (checkpointed).
2. Any changes to the object during period T_2-T_4 activate a “shadow” or “side” file. Figure B-4 shows an update shown as C1 being made to the object before the object has been saved to the media. Pages which are changed are “marked” so that the side file can be referenced when those pages are saved.
   a. A request is made to update C1.
   a. A copy of the original pages are sent to the side file first.
   a. The change is made to the object. The original page copied is part of the checkpoint image for the object.
3. The save process reads the unchanged pages from the actual object and from the side file for any pages which were changed. Time T_3-T_4 shows two additional changes, C2 and C3, made to the object. Additional change requests that are made to the pages of the object already changed for C1, C2, or C3 do not require any additional processing as those pages have already been marked and copied to the side file. At time T_4, the object has been completely saved to the media.
4. After time T_4, the copied pages for the checkpoint image of the object are no longer maintained and are deleted automatically because they are no longer needed.
5. At time T_5, the save process has been completed. Figure B-4 shows that the object on the system has the C1, C2, and C3 changes. The copy or image of the object saved to the media does not contain those changes.

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 405. Note that some of the documents referenced here may be available in softcopy only.

- *IBM eServer i5 and iSeries System Handbook: IBM i5/OS Version 5 Release 3, October 2004*, GA19-5486 (or latest version)
- *IBM eServer i5, iSeries, and AS/400e System Builder: IBM i5/OS Version 5 Release 3, October 2004*, SG24-2155 (or latest version)
- *IBM Enterprise Storage Server*, SG24-5465
- *IBM TotalStorage Enterprise Storage Server Model 800*, SG24-6424
- *IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments*, SG24-5757

Other publications

This publication is also relevant as a further information source:

- *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7444

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

E

N

O

P

iSeries in Storage Area Networks: A Guide to Implementing FC Disk and Tape with iSeries

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

iSeries in Storage Area Networks

A Guide to Implementing FC Disk and Tape with iSeries

Learn how iSeries fits in a SAN environment

Plan how to migrate from SCSI to fiber channel disk and tape

Understand the new Multipath function

This IBM Redbook gives a broad understanding of the new Power5 architecture as it applies to logical partitioned @server i5 servers. This functionality is delivered through a new configuration and management interface called the Hardware Management Console (HMC).

This redbook will help you design your server partition scheme from scratch. We also discuss the requirements to migrate from existing iSeries servers with and without logical partitions.

This book will help you install, tailor, and configure logical partitions (LPARs) on @server i5 servers. You will understand how to setup the server and HMC via the guided setup wizard, and also cover the creation of multiple i5/OS partitions.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6220-01 ISBN 073849190X