Implementing the Enterprise Storage Server in Your Environment

Mark Blunden, Ivan Avsic, Ian Black, Peter Crowhurst, Carlos Sadao Miyabara

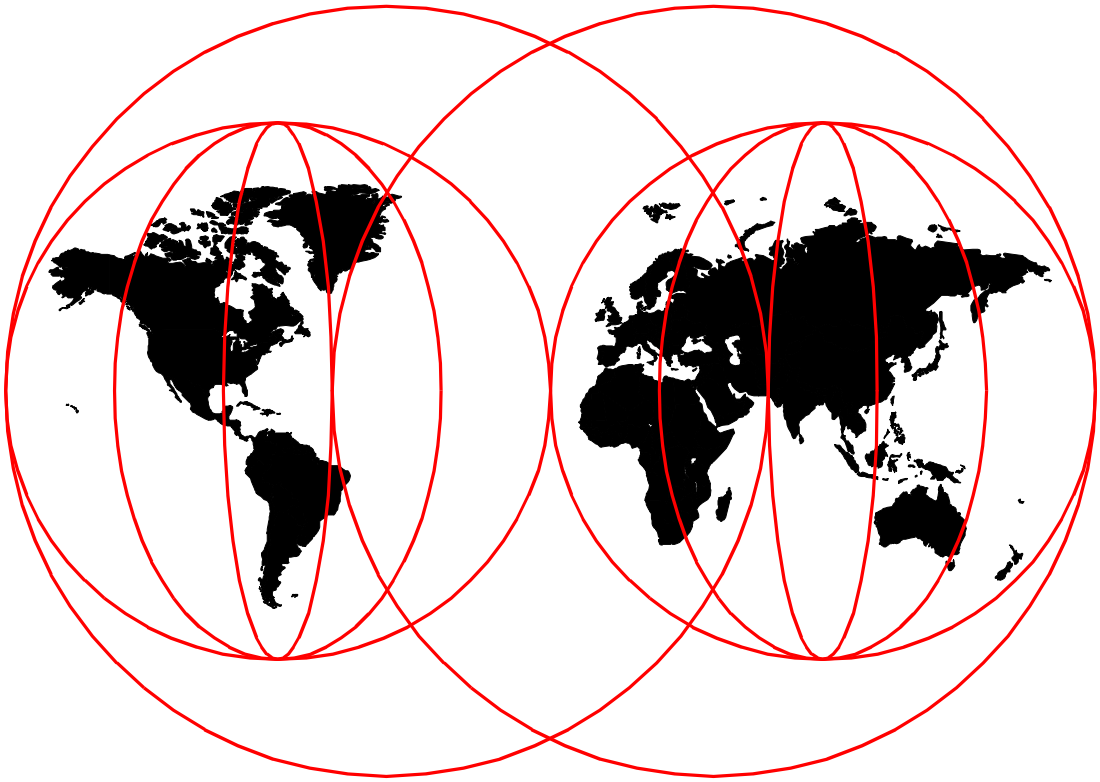

International Technical Support Organization

www.redbooks.ibm.com

SG24-5420-00

International Technical Support Organization

SG24-5420-00

Implementing the Enterprise Storage Server in Your Environment

October 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, “Special Notices” on page 255.

First Edition (October 1999)

This edition applies to the IBM Enterprise Storage Server as announced in August 1999.

Note

This book is based on a pre-General Announcement version of a product and may not apply when the product becomes generally available. We recommend that you consult the product documentation or follow-on versions of this redbook for more current information.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

Tables

Preface

This redbook will help you install, tailor, and configure the new IBM Enterprise Storage Server (ESS) in your environment.

We provide an overview of all functions available for the ESS, and explain how you can use new and existing functions to efficiently manage your storage data. With this new subsystem, you will be able to consolidate your storage requirements, and at the same time, improve your storage management strategy. In this redbook, we show how to install and implement the ESS in S/390, Open, or mixed environments.

The Team That Wrote This Redbook

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

Mark Blunden is the project leader for Open Systems Storage at the International Technical Support Organization, San Jose Center. He has coauthored six previous redbooks and comanaged three other redbooks. Mark also teaches IBM classes worldwide on all areas of storage. He has worked for IBM for 19 years in many areas of the IT business. Before joining the ITSO in 1998, Mark worked in Sydney, Australia, as an Advisory Storage Specialist.

Ivan Avsic is a Sales Support Specialist in Systems Support Center for Emerging Markets, located in Ljubljana / Slovenia. He has more than 25 years experience in Large System field. He has worked several years as an instructor, and is now working as a technical marketing support specialist for Large Systems.

Ian Black is an advisory specialist in the Enterprise Systems technical support group in Australia. Ian has 22 years of experience in S/370 and S/390 systems and hardware technical support. Ian's areas of expertise include 3990, 3380, 3390, RAMAC, and RVA. Ian joined IBM in 1976. He has co-authored one previous redbook.

Peter Crowhurst is a Senior IT Specialist in the Enterprise Systems group in Australia. He has 23 years of experience in the IT industry, with nine years working as an Applications and Systems Programmer, involved with Network planning and design. Peter joined IBM 14 years ago and has worked mainly in a Large Systems technical marketing support role for large customers.

Carlos Sadao Miyabara is the storage manager at Banco Itau, the leader in bank automation process in Brazil. Sadao has been at Itau since 1972, and has more than 17 years of experience in storage management. His area of expertise includes IBM and non-IBM storage solutions. He was project leader for storage management at the GUIDE Latin America. Since 1998, Sadao has been a member of IBM Customer Steering Committee (CSC).

Special thanks to Mike Downie from IBM Learning Services who worked with us during the writing of this redbook. While we were writing this book, Mike was creating a customer education course for implementing the ESS in a S/390, Open, or mixed environment. We were able to provide material to each other, which created a mutually beneficial situation for both parties.

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

- Alison Pate, ITSO, San Jose
- Ruth Azevedo, ATS, IBM San Jose
- John Aschoff, ATS, IBM San Jose
- Terry Quinn, IBM Atlanta
- Pat Blaney, ATS, IBM San Jose
- Brent Beardsley, IBM Tucson
- Bill Worthington, ATS, IBM San Jose
- Mike Noonan, IBM Tucson
- Helen Burton, SSD, IBM Tucson
- Jack Flynn, SSD, IBM San Jose
- Stefan Jaquet, SSD, IBM San Jose
- Richard Kirchofer, SSD, IBM San Jose
- Sam Leinart, SSD, IBM San Jose
- Jerry Coale, SSD, IBM San Jose
- Bill Micka, SSD, IBM Tucson
- Larry Perry, SSD, IBM San Jose
- John Ponder, SSD, IBM Tucson
- Joe Hyde, SSD, IBM Tucson
- Phil Norman, IBM UK
- Cay-Uwe Kulzer, IBM Germany
- Roland Wolf, IBM Germany
- Steve Van Gundy, SSD, IBM San Jose

Comments Welcome

Your comments are important to us!

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

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 269 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Send your comments in an internet note to `redbook@us.ibm.com`

Figure 1 shows the authors, with the ITSO Almaden location in the background.

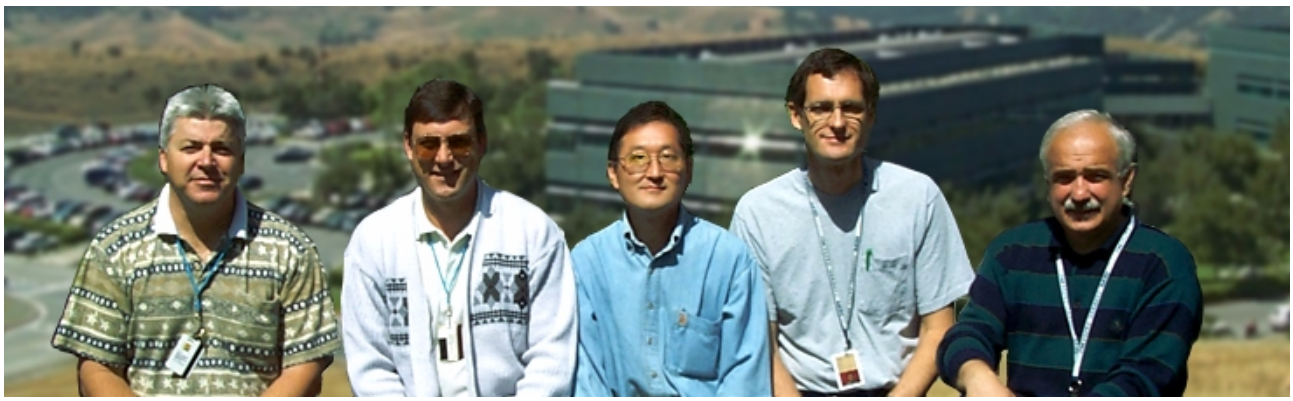

Figure 1. Authors of this Redbook: (l to r) Mark, Peter, Sadao, Ian and Ivan

Part 1. ESS Planning and installation

This Part discusses the planning and installation processes involved for the ESS.

As part of the logical configuration process, the ESS provides the capability to use configuration worksheets, which can be filled in by you and given to the IBM CE at installation time. This information will be translated to provide you with a logical configuration to match your requirements. The information in this book does not mention this process, but instead details the work necessary to implement the ESS manually.

For more information on the initial configuration process, please refer to Appendix A, “ESS Configuration Planner” on page 253.

Chapter 1. Introduction and positioning

In this chapter, we provide an overview of the Enterprise Storage Server, and the functions and features it provides. We also provide some positioning information that will give you an understanding of when the ESS should be used in your environment.

In addition, this chapter provides an overview of the terminology and definitions that will be used throughout this book.

1.1 Overview

The IBM Enterprise Storage Server (ESS) is a member of the Seascape family. It consists of a storage server and attached disk storage devices. The storage server provides integrated caching and RAID support for the attached disk devices. The disk devices are attached via a serial interface. The ESS can be configured in a variety of ways to provide scalability in capacity and performance.

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

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

The following host I/O interface attachments are supported:

- SCSI-3 Parallel Interface
- ESCON
- FC-AL [1]

On SCSI-3 interfaces, the ESS emulates a variety of fixed-block devices with either 512 or 520 byte blocks. SCSI-3 is, in general, a superset of SCSI-2. A SCSI-3 disk device can be attached to a SCSI-2 initiator, provided the cabling can be interfaced. Many SCSI-2 initiators attach directly to the cabling specified for the SCSI-3 parallel interface, but are referred to as SCSI-2 initiators because they limit their use of the command set to the SCSI-2 subset. Host systems with SCSI-2 or SCSI-3 interfaces can attach to the ESS. The ESS provides multiple SCSI I/O interfaces (busses), each with multiple SCSI targets, and each with multiple disk logical units. The storage provided by the ESS for SCSI interfaces can be configured so that it is shared among multiple SCSI interfaces if desired.

On ESCON interfaces, the ESS emulates one or more IBM 3990 control units attaching variable size IBM 3390 devices in either 3390 or 3380 track format. The ESS provides multiple ESCON interfaces that provide a set of control unit images, each with multiple disk devices. The storage provided by the ESS for ESCON interfaces is configured so that it is accessible from any ESCON interface.

1 FC-AL at GA is supported through 2108 SAN Data Gateway until native support is provided

The ESS is composed of the following components:

- The storage server is composed of two clusters that provide the facilities with advanced functions to control and manage data transfer. Should one cluster fail, the remaining cluster can take over the functions of the failing cluster. A cluster is composed of the following subcomponents:
  - Host Adapters - Each cluster has one or more host adapters (HAs). Each host adapter provides one or more host I/O interfaces. A host adapter can communicate with either cluster complex.
  - Device Adapters - Each cluster has one or more device adapters (DAs). Each device adapter provides one or more storage device interfaces. Disk drives are attached to a pair of device adapters, one in each cluster, so that the drives are accessible from either cluster. At any given time, a disk drive is managed by only one device adapter.
  - Cluster Complex - The cluster complex provides the management functions for the ESS. It consists of cluster processors, cluster memory, cache, nonvolatile storage (NVS) and related logic.
    - Cluster Processor - The cluster complex contains four cluster processors (CP) configured as symmetrical multiprocessors (SMP). The cluster processors execute the licensed internal code that controls operation of the cluster.
    - Cluster memory / cache - Is used to store instructions and data for the cluster processors. The cache memory is used to store cached data from the disk drives. The cache memory is accessible by the local cluster complex, by device adapters in the local cluster, and by host adapters in either cluster.
    - Nonvolatile storage (NVS) - Is used to store a nonvolatile copy of active written data. The NVS is accessible to either cluster-processor complex and to host adapters in either cluster. Data may also be transferred between the NVS and cache.
  - Disk Drives - Provide the primary nonvolatile storage medium for any host data stored within the ESS Storage devices. They are grouped into ranks and are managed by the clusters.

As a member of the IBM Seascape family, the ESS provides the outboard intelligence required by SAN solutions, offloading key functions from host servers, which frees up valuable processing power for applications. As a comprehensive SAN-based storage solution, the ESS provides considerable management flexibility to meet the fast-paced requirements of the next century.

Among the many factors that make the IBM ESS an ideal SAN solution are:

- Supports all major server platforms including S/390, AS/400, Windows NT, and many varieties of UNIX
- Fiber channel attachment capability
- Extensive StorWatch management capabilities through a Web interface
- Excellent scalability:
  - From 400 GBs to over 11 TBs
  - Simple selection from 16 standard configurations to meet capacity and performance needs

- Performance optimized to your heterogeneous environment needs
  - High bandwidth and advanced transaction processing capabilities provide solutions for both online and batch applications
  - Innovations such as Parallel Access Volumes to reduce resource contention and dramatically improve performance
- Availability required to support e-business applications
  - Non-disruptive access to data while making a copy using Concurrent Copy
  - Business continuity through remote copy services - PPRC and XRC
  - Rapid data duplication through FlashCopy, providing extensive capabilities to exploit, manage, and protect your information in a 7 x 24 environment
  - Storage server availability through redundancy and nondisruptive service with design for no single point of failure or repair

1.2 Functions and features

The ESS provides many functions that will help you manage your enterprise storage, and provide efficient usage of your storage resources.The functions include:

- Copy Services — These include functions such as:
  - FlashCopy—the ability to create time zero (T0) copies of logical volumes
  - Peer-to-Peer Remote Copy (PPRC)—the ability to create synchronous volume copies via ESCON channels
  - Extended Remote Copy (XRC)—the ability to create asynchronous volume copies over long distances
  - Concurrent Copy (CC)—the ability to create volume or data set copies locally
- ESS Specialist—the Web interface tool used to manage the ESS logical configuration
- High performance—provided by a suite of hardware and software features called the Enterprise Storage Server EX Performance Package. Some of these include:
  - Parallel Access Volumes (PAVs)—this feature allows you to dramatically reduce or eliminate IOSQ time for a single-host environment
  - Multiple Allegiance—this feature allows you to dramatically reduce or eliminate IOSQ time for a multiple-host environment
  - Performance Enhanced CCW commands
  - I/O Priority Queuing—allows users to define priority of application workloads
  - Custom Volumes—the ability to create your own custom-sized logical volumes
- Scalability—Disk capacity is scalable from 420 GBs to over 11 TBs

1.3 Explanation of terms

In this book, all disk capacities are in decimal form. One decimal GB = 1,000,000,000 bytes.

Cache, NVS and memory capacities are in binary. One binary GB = 1024^3 = 1,073,741,824 bytes.

Since this book will be read by planning and implementation experts with very different background and skills, the following are some of the major terms and expressions used throughout this book:

- S/390 refers to the System 390 or OS/390 environment
- Open refers to the SCSI or Fibre Channel attached systems, such as UNIX, NT, Netware, or AS/400 environments
- Device end = task complete or I/O complete
- hdisk = LUN (Logical Unit Number) = device
- ESCON = Enterprise Systems Connection Architecture. An ESA/390 computer peripheral interface. The I/O interface utilizes ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric.
- FBA (Fixed Block Architecture) = Track format used for logical devices used by open systems and accessed through SCSI (Small Computer System Interface) ports in the ESS. FBA devices in the ESS are often referred to as SCSI devices.
- CKD (Count Key Data) = Track format used for emulated 3390 logical devices used by S/390 systems and accessed through ESCON ports. CKD devices in the ESS are often referred to as S/390 devices, or ESCON devices.
- JBOD (just a bunch of disks) = Group of individual disk drives not configured as a RAID (redundant array of independent disks) array.
- Rank = A RAID array, or an individual JBOD disk drive. A rank is a unit of physical disk storage in the ESS that is allocated to either FB or CKD logical volumes.
- To those familiar with UNIX, a host adapter is an adapter on the host. Throughout this book, a host adapter is a SCSI or ESCON adapter that resides in the ESS, and is used to connect the host server bus to the ESS.
- LSS (Logical Subsystem) equates to a logical control unit (LCU) in the S/390 environment. For the open environment, LSS is a new term used by the ESS. An LSS is used internally to manage a set of logical volumes, which is associated with an individual device adapter.
- Disk group refers to a group of eight drives that can be defined as either a RAID array, or non-RAID ranks.
- CE refers to the IBM customer engineer who is responsible for the physical installation and maintenance of the ESS.

1.4 Platforms and concepts

The ESS provides support for many different platform connections. These include most forms of UNIX, AS/400, PC servers running NT or NetWare, and S/390

processors. For a detailed description of the latest server support for both hardware and software levels, please refer to the Web address:

www.storage.ibm.com/ESS

1.5 Futures

At the announcement of the ESS, IBM also provided a Statements of Direction (SODs) that told of future enhancements to the ESS, which included higher capacity devices, higher performing processors, increased connectivity options, and function enhancements.

Chapter 2. Configuration planning

This chapter details a set of planning tasks that should be considered to ensure that the ESS will be sized and configured according to your storage needs and performance metrics. The process through the sections will most likely be an iterative one until all parties involved in the planning agree that a valid configuration has been reached.

The output from the chapter will be a definition of both the physical ESS and an internal logical configuration plan.

2.1 Introduction

The Enterprise Storage Server (ESS) is designed for sharing its storage capacity between a mix of both S/390 and Open (SCSI)-based server host systems.

As a result, the people involved in the capacity requirements and the logical configuration planning will be from several different groups that are not traditionally combined, as shown in Figure 2.

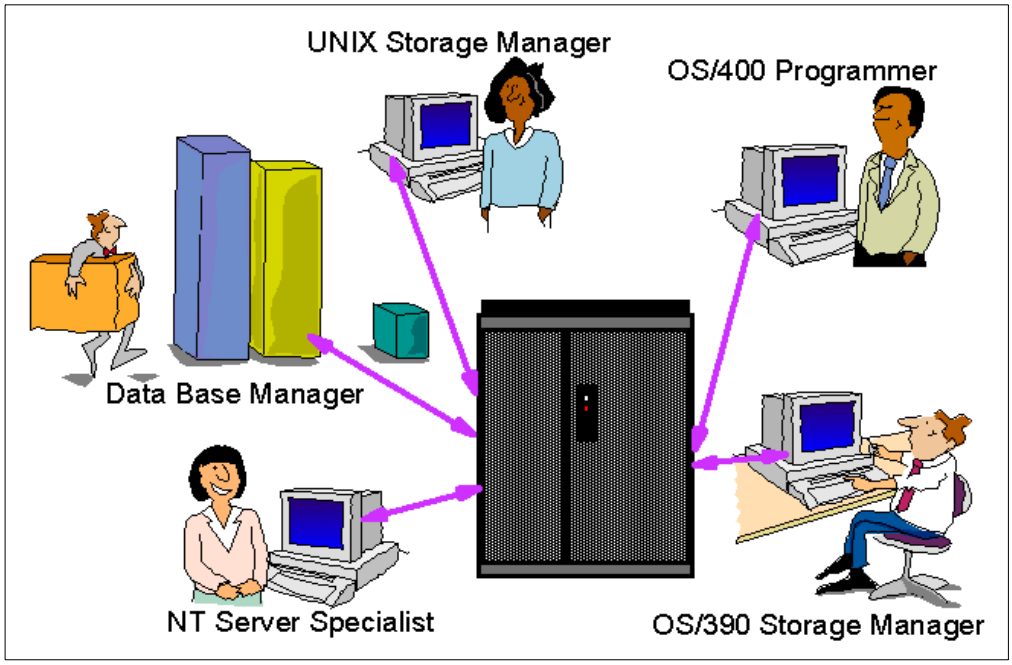

Figure 2. The expanded planning team

The disk capacity requirements from all people intending to store data in the ESS will need to be combined to decide the total ESS physical capacity. Additional details relating to the number of server hosts to connect, and what the host view of the storage allocation will be, are used to build up the ESS logical storage configuration.

This logical plan is then used to configure the ESS internally via the ESS Specialist as part of the installation process.

2.2 Capacity planning

This section is intended to aid the sizing of the physical storage capacity required for your Enterprise Storage Server. The tables included in this section will enable you to calculate the total number of 8-pack disk arrays required when ordering the ESS. The process involves:

- Calculating the number of arrays required for SCSI fixed block format data
- Calculating the number of arrays required for S/390 CKD format data
- Combining the two for a total ESS capacity requirement

It is assumed that you have a base knowledge of the ESS architecture, functions and terms used when describing the ESS. If the terms or concepts used are unfamiliar, it is recommended that you read the *Enterprise Storage Server* redbook, SG24-5465.

2.2.1 Capacity and configuration rules

This section is intended as a quick reference of rules or limitations for the different components and functions available on the ESS.

2.2.1.1 Physical features

Physical features of the ESS include:

- Maximum 16 Host Adapter cards, either ESCON or SCSI or a mix of both. Each adapter card has two ports, thus there are 32 physical host connections.
- 4 Device Adapter (DA) pairs with two loops per DA pair - total of 8 loops/ESS.
- From a minimum of two and up to six 8-pack arrays per loop.
- Each RAID-5 loop will have two spares, therefore, the first two arrays on the loop will be 6+P+S, the next four arrays will be 7+P.

2.2.1.2 Logical features

Logical features of the ESS include:

- One loop may have its arrays assigned to up to 4 LSS/LCU.
- A rank can belong to only one of the 4 LSS/LCU on a DA pair.
- You can mix RAID and JBOD arrays on the same LSS/LCU or loop.
- You can mix CKD and FB formatted arrays/ranks on the same loop.

S/390 Host

- 0/8/16 Logical Control Units (LCU) in CKD format per ESS
  - Maximum 4096 Devices
  - Maximum 256 Devices per LCU
  - Maximum 1024 Devices per ESCON port - ESCON limitation
  - Maximum 64 Logical Paths per ESCON port
  - Maximum 128 Logical Paths per LCU with up to 64 path groups
  - Maximum 2048 Logical Paths per ESS

SCSI Host

- 0/8/16 Fixed Block LSS per ESS
  - Maximum 960 LUNs (15 target x 64 LUN) per host adapter
  - Maximum 256 devices per LSS (SCSI target ID/LUN group)
  - 1 to 15 SCSI targets per SCSI bus/port
  - 1 to 64 LUNs per target

- Up to four initiators (hosts) maximum recommended per SCSI port
- LUNs may be shared between hosts, with hosts controlling contention
- Data path optimizer requires separate SCSI ports

2.2.1.3 S/390 Host view

Figure 3 shows an example of the S/390 logical path limits and considerations.

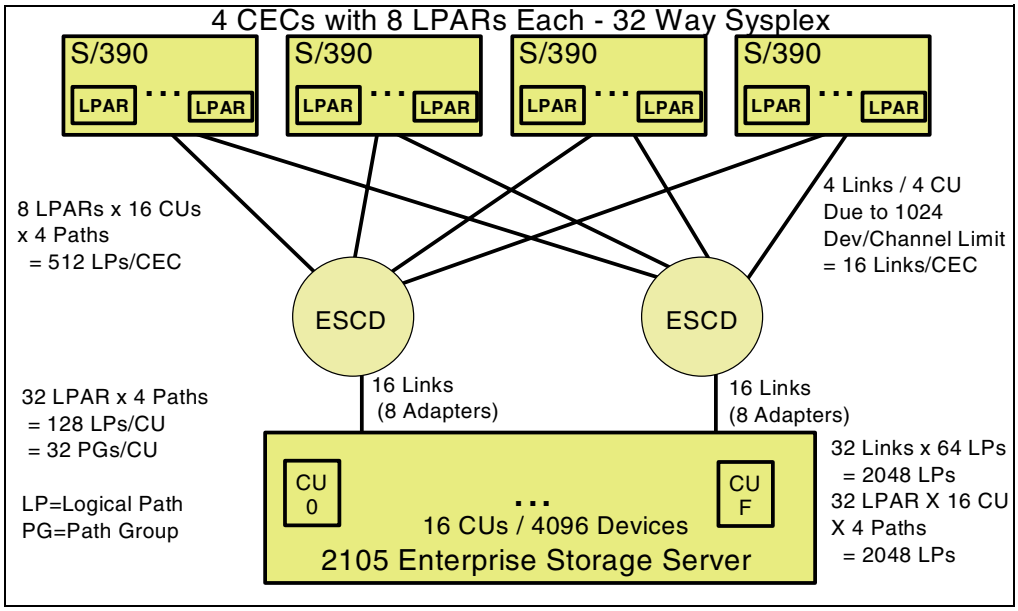

Figure 3. S/390 Logical Paths

2.2.1.4 SCSI Host view

Figure 4 shows an example of the logical SCSI bus view of the ESS.

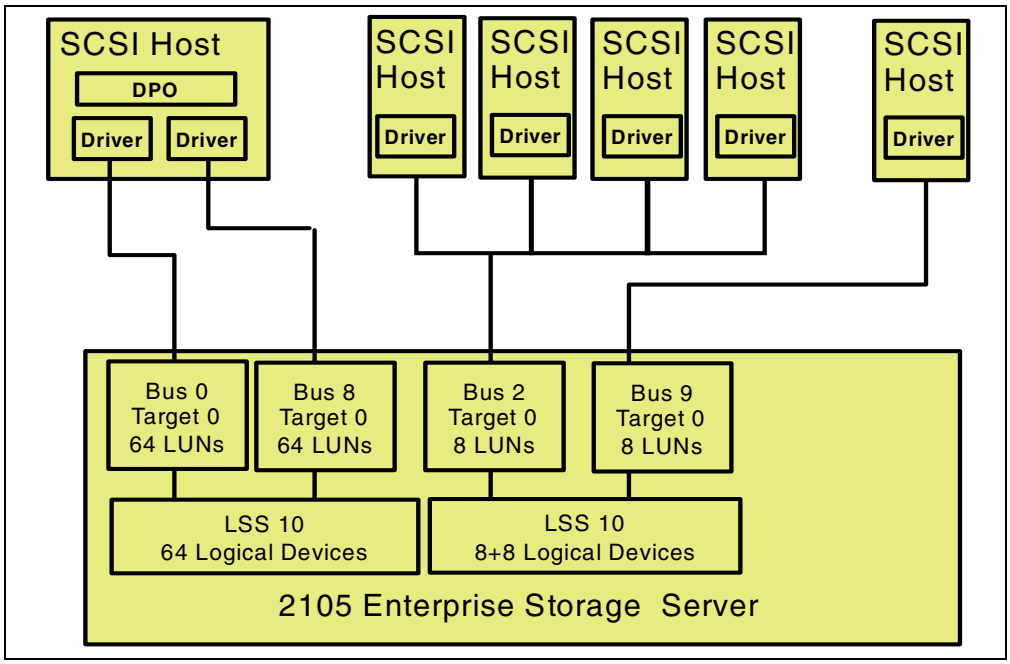

Figure 4. SCSI Logical Configuration

2.2.2 Storage capacity options

For capacity planning purposes the effective data capacity of the ESS is reduced from it's theoretical maximum, using the raw disk sizes 9 GB, 18 GB or 36 GB, due to:

- Disk sector overheads:
  - Disk sector size is 524 bytes, data area 512 bytes
  - Metadata area on each disk
- The formatting of disks into:
  - RAID-5 or non RAID support
  - SCSI FB devices or S/390 CKD devices

2.2.2.1 ESS model capacity

Table 1 details the minimum and maximum RAID-5 storage capacity available by the ESS Model and is mapped against the available physical disk drive sizes. The table is designed to assist you in the selection of an appropriately sized disk drive/ ESS model combination that will meet your current and future capacity needs.

It is suggested that you select a configuration closest to your needs, since this will simplify the ordering process and speed the installation process.

Table 1. ESS RAID-5 Disk storage capacity by model

ESS	Total Arrays	9.1GB	18.2GB	36.4GB
Min E10/E20	8	420 GB	840 GB	1680 GB
Max E20 Base	16	840 GB	1680 GB	3360 GB

ESS	Total Arrays	9.1GB	18.2GB	36.4GB
Max E20 Base plus Expansion	48	2792 GB	5600 GB	11200 GB
Max E10 Base	8	420	840	1680

2.2.3 Capacity planning by storage format

The next step in your planning requires you to separate your total storage needs into the different format types, SCSI Fixed Block (FB) or S/390 CKD. The tables in this section can be used to aid you in determining the number of 8-pack arrays required for each format and, therefore, the total for the ESS.

To aid planning, the RAID array capacity data is shown separately for ESS arrays 1 to 8 and 9 to 32 for the Model E10, and 1 to 16 and 17 to 48 for the Model E20. This is due to the fact that the first two RAID-5 arrays on a loop contain a spare disk (the exception being JBODs). You can therefore assume that all the RAID-5 arrays in the base ESS will contain a spare disk, 6+P+S, and all of the RAID-5 arrays in the expansion frame will be 7+P arrays only.

Why total the arrays? The minimum unit of storage installable or upgradeable within the ESS is two 8-packs.

2.2.3.1 Effective data capacity

Each group of eight disks can be formatted into RAID-5 or non-RAID JBOD support. In the case of RAID support, a single partition, called a Rank in ESS terminology, is allocated across the entire array of disks. For non-RAID JBOD support, the single partition is allocated on a single disk drive, one for each physical disk.

Logical volumes or logical disk devices are placed within these partitions, or ranks.

Therefore, the first decision to effect the total data capacity available is how many RAID-5 or non-RAID JBOD disk ranks are required. For example, assuming a disk group of 8 x 36GB disks the maximum available capacity for:

- JBOD is 288GB
- RAID-5 is 245GB

The effective data capacity of an array/rank is also dependent on how you configure the logical volumes; SCSI FB or S/390 CKD. As a result, you need to calculate the effective data capacity required in terms of arrays for each of the logical volume types.

When defining standard logical volume sizes or formats onto a disk or array, you may not be able to fully utilize the storage space available. This is true if the disk or array size is not a multiple of the LV size. You can, therefore, define smaller sized LVs, or S/390 custom volumes, to best utilize the remaining disk storage.

2.2.3.2 JBOD non-raid data capacity

A JBOD has one non-interleaved partition on each physical disk device in the array, thus there are 8 JBODs/ranks per group of drives. The Logical Volume (LV)

size within a JBOD is 0.5 GB to 36 GB, dependent on the physical disk drive installed. A JBOD can be defined as either SCSI FB or CKD.

A JBOD does not require a spare disk, and therefore will take over the entire 8 disks as non-Raid data disks, regardless of the position. See Table 2.

Table 2. Non-Raid Array Capacity - JBOD

Disk Capacity	All Arrays (8+No Parity or Spare)
9.1 GB	72 GB
18.2 GB	144 GB
36.4 GB	288 GB

2.2.3.3 SCSI fixed block data capacity

A SCSI FB RAID-5 array/rank has one partition that is striped across all the disks within the array and, therefore, the maximum storage size available for placing Logical Volumes is the capacity of the entire array, as shown in Table 3. UNIX, AIX or Windows NT, can allocate LVs from 0.5 to 224 GB. Multiple LVs can be placed in an array/rank.

OS/400 supports 9337 devices with SCSI FB Logical Volume sizes of 4.2 GB, 8.6 GB, 16 GB or 32 GB. Multiples of these can be placed into an array.

Table 3. Raid Array Capacity - FB

Disk Capacity	Arrays 1 to 16/8 (6+P+S)	Arrays 17/9 to 48/32 (7+P)
9.1 GB	52 GB	61 GB
18.2 GB	105 GB	122 GB
36.4 GB	210 GB	245 GB

2.2.3.4 S/390 CKD data capacity

S/390 effective data capacity can vary between SCSI FB devices due to the 3390 device formatting required. CKD has additional overheads, for example the CK itself.

The CKD track format can be placed on a RAID-5 array or non-RAID JBOD rank. The RAID-5 array/rank can be formatted with interleaved logical volumes mode or non-interleave logical volume mode.

Interleaved: Interleaved arrays/ranks have two partitions striped across the array. The first partition is in interleaved mode and accounts for the major portion of the array/rank capacity and is used to place 3390/3380 Logical Volumes. LVs are placed within the partition in groups of four until there is not enough storage space left in the array/rank to successfully store another group of four. All the remaining storage space is then allocated to the second non-interleaved mode partition.

The second non-interleaved mode partition in the array/rank will be a minimum of 5,000 3390 cylinders and can be used to define additional Logical Volumes (either standard 3390 Models or Custom Volumes).

Non-Interleaved: A non-interleaved mode array/rank has only one partition that is striped across the array. The entire partition can be used for logical volumes (either standard 3390 models or custom volume).

The maximum CKD logical volume size is dependent on the 3390 model type that has been formatted on the array/rank (3390-2 1.89 GB, 3390-3 2.83 GB, 3390-9 8.49 GB). A Custom Volume can be from 1 to 10017 cylinders (3390-9) in size.

Table 4 and Table 5 show the available data capacity for the interleaved partition (I-Mode) and for non-interleaved (NI-Mode) arrays/ranks. The non-interleaved capacity value is also equivalent to an Interleaved mode array in which custom volumes have been used to fully occupy the storage capacity of the second partition. The number and size of custom volumes allocation will affect the total number of volumes that can be created. The numbers in the Max Vols column of the table reflect the number of full-sized 3390-2/3/9 volumes.

Table 4. S/390 capacity for arrays - 6+P+S

Disk Drive Size	9 GB		18 GB		36 GB	
	Array GB	Max Vols	Array GB	Max Vols	Array GB	Max Vols
3390-2 I-Mode 3390-2 NI-Mode	45 52	24 27	90 105	52 54	196 210	104 109
3390-3 I-Mode 3390-3 NI-Mode	45 52	16 18	90 105	32 36	193 210	68 73
3390-9 I-Mode 3390-9 NI-Mode	34 52	4 6	68 105	8 12	170 210	20 24

Table 5. S/390 capacity for arrays - 7+P

Disk Drive Size	9 GB		18 GB		36 GB	
	Array GB	Max Vols	Array GB	Max Vols	Array GB	Max Vols
3390-2 I-Mode 3390-2 NI-Mode	53 61	28 32	113 122	60 64	227 245	124 128
3390-3 I-Mode 3390-3 NI-Mode	45 61	16 21	113 122	40 42	227 245	80 85
3390-9 I-Mode 3390-9 NI-Mode	34 61	4 7	102 122	12 14	238 245	24 28

2.2.4 Copy services

The Flashcopy function provides a time zero (T0) image copy within the same LSS or LCU. Some thought needs to go into the number of logical volumes that you may wish to use as target LVs since disk capacity must be allowed for within the same LSS or LCU. This may effect the total capacity required for the ESS.

PPRC and or XRC will also need to have additional disk storage reserved, if the target LVs are within the ESS.

PPRC requires that a primary and secondary ESS be connected by ESCON connections, even in an Open Systems environment. These paths are unidirectional, and the primary ESCON adapter ports are dedicated to PPRC and

cannot be used for normal host data traffic while in a PPRC session. These points must be considered during planning, since it may effect the number or type of Host Adapters required when ordering the ESS.

When using Copy services, you are required to nominate one ESS as the Copy services primary server, which is the central place for collection of all information. A second ESS can be nominated as the backup server. You will need to plan which ESS you would nominate in which role to ensure maximum availability. The CE will need this information during ESS installation.

2.2.5 Standard physical configurations

You may choose to order a standard physical configuration. These configurations have been designed in such a way that they meet most of the your requirements. The advantage of these configurations is that they will allow a quick setup and easier implementation into your site. You only need to know the basic capacity requirements of the user and then match those to the best available option. The standard configurations have been selected after considering:

- Performance
- Capacity
- ESCON-only environments
- Fixed block-only environments
- Mixed environments (SCSI FB and ESCON)

The proposed configurations are ordered by feature code. For details about the features codes required at ordering time, refer to your IBM representative.

2.3 Logical configuration planning

The Logical configuration of the Enterprise storage server involves defining how the ESS is seen from the attached hosts. The UNIX, AIX, NT system hosts will see the ESS as SCSI generic devices, the AS/400 will see it as a 9337 external disk, and the S/390 host will see it as 3990 subsystems.

The basic steps required are:

- Draw up a logical map/plan of the storage subsystems and devices you want to emulate within the ESS.
- Map these to the physical hardware arrays within the ESS.
- Review the ESS capacity, adjust as necessary.
- Combine all the decisions made and document them in the form of a specification or logical configuration plan, which can then be used to setup the ESS via the ESS specialist during installation.
- the logical specification is also required by the host systems or storage software specialist to enable the appropriate I/O definitions to be completed

Note: The steps shown in this section do not reflect the logical worksheet process which is available to help you pre-define your logical layout, and then help the IBM CE input these requirements at installation time. For more information on the initial configuration process, please refer to Appendix A, “ESS Configuration Planner” on page 253.

2.3.1 Logical subsystem mapping

Determine the number of SCSI bus or CKD subsystems (LSS/LCU) to be defined. Multiple targets/LUNs can share the same associated logical device.

Decisions that need to be made for SCSI Bus devices include:

- How will the hosts be connect to the ESS?
- Determine RAID or non-RAID format.
- The LV or device size.
- The number of targets and the number of LUNs per target.
- Addressing requirements. Are specific SCSI IDs required for host and target? It is recommended to set the initiator ID to 7.
- Will there be shared LUNs?

An S/390 LSS relates directly to a 3990 control unit with its associated devices or volumes. Therefore, a decision also has to made about the 3990 control unit model or models to be emulated, either Model 3, Model 3-TPF or Model 6. A four digit subsystem identifier SSID will need to be assigned to each logical control unit (LCU).

2.3.1.1 LSS/LCU to physical mapping

S/390: The CUADDR address is determined by the DA and loop that the LCU is associated with when defined using the ESS Specialist. For example, in Figure 5, the CUADDR for cluster 1, adapter 1, loop B is 08.

SCSI Bus: The mapping of the SCSI FB LSS to target/LUN is also dependent on the DA and loop association, and is automatically assigned during ESS Specialist setup. The SCSI FB LSS numbers that the ESS Specialist can be assigned are shown in Figure 5 below. These are not required to be known by the SCSI planner and have been included for informational purposes only.

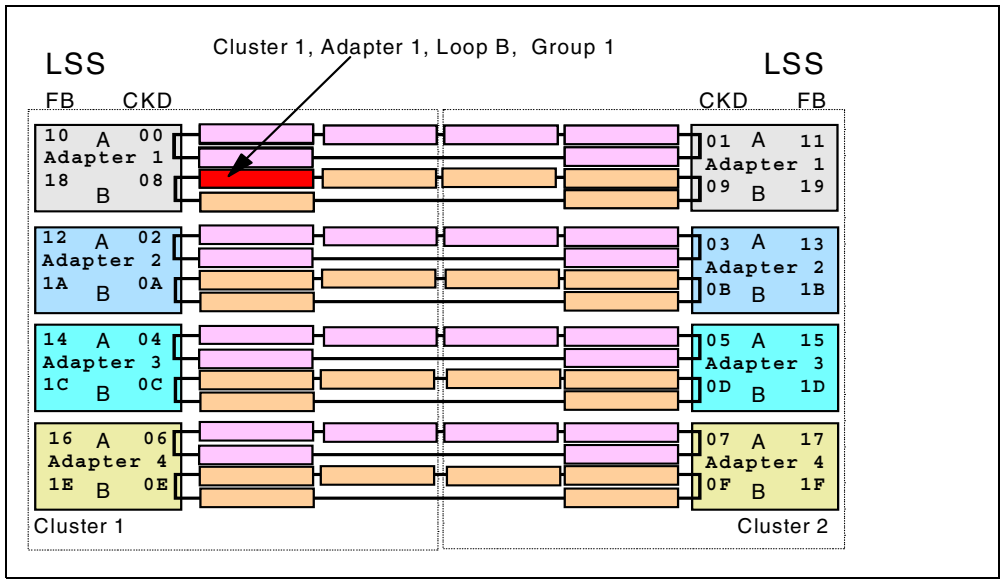

Figure 5. LSS Logical to physical mapping

2.3.1.2 Logical device mapping

RAID or JBOD: The first step requires you to define which arrays/ranks are to be assigned/mapped to which LSS/LCU. The disk arrays to be attached to the

LSS/LCU then need to be defined as RAID-5 or non-RAID JBOD disks. This will determine the rank storage capacity available. For example, a single disk rank for JBOD and a full disk array for RAID.

The drive groups on a loop are allocated evenly between the two device adapters. Therefore, on a fully populated loop, only three arrays will be available to the LCUs or SCSI target on each adapter. For example, LCU 00 or 08 will have access to disk groups 1, 3, and 5 on the loop, with LCU 01 or 09 accessing disk groups 2, 4 and 6. The ESS Specialist manages this for you.

Logical Volumes: For each array/rank, specify the number and size of the logical volumes or devices required.

SCSI: You can map the same logical devices to different SCSI ports to provide shared logical volumes. However, the SCSI host software must handle the data contention.

A JBOD can have more than one logical volume assigned to it.

The ESS maximum number of LUNs (or LVs) per target is 64, however, it is host operating system or SCSI adapter dependent. For example, a large number of servers will only support up to 8 LUNs per target.

S/390: Note whether you will be using Interleaved or non-interleaved volumes. If using interleaved volumes, will you be defining custom volumes in the second non-interleaved partition.

A decision must be made on the logical device type to emulate, 3390 Model 2, Model 3, or Model 9 or 3380 track format. If using custom volumes, decided how many volumes are required and how many cylinders are to be allocated to the volumes. Custom volumes can only be defined on non-interleaved partitions.

Parallel Access Volumes: If you need to use PAVs, the LCU needs to have the PAV feature enabled via the ESS Specialist. The volumes that require PAV Aliases need to be identified, along with the number of Aliases to be available to each volume.

2.3.1.3 Host connections

The number of host connections attached to the LSS/LCU needs to be decided.

The ESS SCSI logical volumes are specifically assigned to an HA port via the ESS Specialist at ESS setup. If using the IBM DPO product, the physical SCSI connections should be spread across different HA bays.

When planning how much capacity to attach to an SCSI interface, the following conservative guidelines are suggested:

- For Ultra-SCSI (40 MB/Sec) - Up to 120 GB per SCSI bus
- For SCSI-2 (20 MB/Sec) - Up to 80 GB per SCSI bus

For S/390 this is fairly straight forward since it is no different than normal device considerations. It is recommended that eight ESCON channels per MVS image be used. All host ESCON adapters are able to address all LCUs defined within the ESS. However, you should spread the physical ESCON connections from one Host image across the HA adapter bays. Be sure to review the logical path limitations.

PPRC and XRC also have an additional requirement for ESCON host adapter ports. PPRC primary or source volume end ESCON ports are dedicated and, therefore, cannot be shared with normal host traffic while the PPRC session is active. The ports on the secondary device can, however, be used for other host activity.

2.3.1.4 VPD settings

The following ESS global values will be set during the ESS installation:

- Nominate which ESS will be the ESS Specialist Copy services primary server and its backup.
- The hostname list must be checked and updated if all ESS cluster hostnames are not present, including its own.
- PPRC Crit=Yes Light or Heavy option - Determine how PPRC will act in the event of a problem between the pairs.
  - Light - Suspend the pair and do not accept writes to the primary if the two control units cannot communicate, or accept writes to the primary if they can still communicate but cannot access the target volume (device problem). The ESS will record changed cylinders for later synchronization.
  - Heavy - Suspend the pair and do not accept any further writes to the primary if data cannot be sent to the secondary.

You should, therefore, decide on the options required and document them so they are available to the CE at installation time.

2.3.1.5 Plan for the future

It is a good idea to consider your future capacity requirements and the effect this may have on your planned ESS configuration. This may prevent a lot of data relocation work later on. Items to consider include:

- If upgrading the ESS capacity, where will the next group of arrays/ranks be placed?
- Will this require movement or redefinition of LSS/LCU and its logical volumes due to performance, capacity or device number limitations?
- Reformatting the array from RAID to JBOD, or changing the logical volume size or the CKD device type, will require unloading and then reloading of data.

2.4 Planning for performance

This section provides some considerations on the physical and logical configuration that may contribute to better performance. Reviewing these points is recommended before finalizing the physical or logical plan for your ESS.

2.4.0.1 Disk drive capacity

There is a trade-off decision to be made when ordering your ESS. This is disk drive capacity selection versus performance. You have the option to choose 9 GB, 18 GB or 36 GB disks.

If you require high performance, select the lowest capacity disk drives that will give you the total storage capacity required. For example, the 9 GB and 18 GB disks are designed for high performance and the 36 GB disks are designed for

high capacity with good performance. The selection of 9 GB or 18 GB disks will obviously effect the maximum storage capacity available within the ESS.

If taking advantage of the ESS ability in order to install 7133 devices in the expansion frame, be sure not to mix 7133-020s with D40s. Mixed speed devices will effect the performance of the loop. For example, D40s will be reduced to 020 data transfer rate.

2.4.0.2 Host adapters

When physically connecting a host system to the ESS host adapter ports (ESCON port or SCSI bus port), multiple connections from the same host system should be spread across the host adapter bays. This not only provides better availability due to the possible loss of a bay, but also spreads the I/O load across different PCI buses in the cluster and maximizes the available data bandwidth.

S/390: For high performance workloads, attaching 16 ESCON channels per MVS image and using eight-path path groups to LCUs will spread the I/O channel activity at peak times. This may not be practical in some Sysplex environments, but it is recommended that eight ESCON channels be used. Be sure to check the logical path restrictions.

The 8 channels for a path group should be cabled to ports A and B on 4 host adapters, 1 adapter per HA Bay. See Figure 23 on page 66 for more explanation.

With 16 channels from a single image, you need to spread your data across volumes evenly split into two groups--one group associated with ESS cluster 1, one set of four device adapters, the other associated with cluster 2, the second set of four device adapters. Define as many LCUs as you can on each cluster (maximum eight per cluster) and distribute your volumes across the LCUs.

One eight-path path group should then provide access to the LCUs defined on ESS cluster 1, for example, LCU CUADDR 0, 2, 4, 6, 8, A, C, E. The second eight-path path group should provide access to the LCUs defined on the other Cluster 2, CUADDR 1, 3, 5, 7, 9, B, D, F. See Figure 5 on page 17 for details. This will maximize the bandwidth available through the host adapters.

To maximize the operations per second from the eight-way path group, configure all eight channels to the same IOP or SAP. However, this may not be achievable due to host channel availability requirements.

Data Path Optimizer: The Data Path Optimizer for NT and AIX can provide multiple SCSI bus paths to the same target/LUN group, thus spreading the I/O workload.

2.4.0.3 Cache

If you are only defining one type of LSS/LCU (S/390 or SCSI bus) within the ESS, it is recommended that you set the other LSS quantity to zero. This releases additional cache storage for data, rather than holding control information about the LSS. For example, if only using SCSI LSS, set the S/390 CU LSS value to zero. The LSS/LCU maximum values are set by the IBM CE during installation.

2.4.0.4 Device adapters

It is not advisable to mix disk devices of different speeds on the same loop. This applies when you intend to install 7133-020 drives on a loop. Use the reserve DA

pair feature (FC 9904) on the ESS base so that only 7133-020 can be installed on the loops.

2.4.0.5 LSS/LCU

You should use as many LSS/LCU as possible. This will enable you to spread all your required logical volumes across a greater number of LSS/LCU and therefore DAs. The lower number of LVs per LSS will better utilize the available loop bandwidth.

Specifying a larger number of LSS/LCU will also enable you to spread the LSS/LCUs across the two ESS clusters. Since each DA in a pair is associated with a different cluster, the I/O workload is split, and thus able to take advantage of the bandwidth and processing capacity of both clusters.

2.4.0.6 Data isolation

For performance reasons, you may wish to isolate volumes containing data with a particular activity pattern or attribute, for example cache unfriendly. Use the ESS design to your advantage by isolating the data at either the Cluster, SCSI target LSS/LCU (DA), or array/rank level.

2.4.0.7 JBOD

JBOD should be used in a situation where you require high random read performance and you do not need RAID protection. An example would be where the host application forced software mirroring to provide the availability that is lost in JBOD mode.

S/390: An example of JBOD use is TPF, high I/O rates with the software providing disk mirroring functions. Another example is the IMS Fast Path with multi-area dataset (MADS). Again, the software provides duplexing of the data.

2.4.0.8 RAID

RAID arrays will perform better than JBOD for reads and sequential operations because the data is striped across multiple disks and more I/O can be done in parallel.

S/390: The ESS provides additional performance options for S/390 RAID arrays. The arrays/ranks can be partitioned into either interleaved logical volume mode or non-interleaved logical volume mode.

Using interleave mode will give more even performance across all logical volumes, but at the expense of some additional definition work that needs to be done if you want to utilize the entire array/rank storage capacity. For example, you will need to define custom volumes in the non-interleaved partition appended after the interleaved partition.

2.4.0.9 Custom volumes - S/390

Custom volumes are variable-sized S/390 formatted volumes. They could be used to place a single critical performance dataset where we would like to have it on its own volume to minimize contention (IOSQ). Custom volumes can also assist you in circumventing hardware serialization problems at the volume level, for example reserve/release.

Large-sized custom volumes with PAV support will enable improved performance over standard sized volumes with no PAV support.

If your operating system does not support PAVs, then custom volumes could be used to reduce volume contention.

2.4.0.10 PAV - S/390

Parallel Access Volumes provide the ability to almost eliminate IOSQ time, a major contributor to response time. The decision on whether to use PAVs, and how many PAVs to allocate to a volume, should be based on a study of your application access profile for a particular system image, and then sysplex wide. The sysplex-wide view will also assist in determining if dynamic PAVs under workload manager's control would be an advantage.

PAVs can be very useful for improving the performance of large volumes that may previously have had large IOSQ values, for example, a 3390 Model 9. You should get better performance with one Base and two aliases on a 3390 Model 9 than from three 3390 Model 3 volumes with no PAV support. This would also reduce storage management cost associated with high numbers of volumes that need maintaining.

To get the maximum benefit from PAVs and the standard ESS multiple allegiance feature, Enqueue is much preferred over Reserve/Release.

For maximum PAV/multiple allegiance benefit, the define-extent coverage on the logical volume should be minimized. This is especially important for writes to the same extent which will still need to be serialized.

2.4.0.11 PPRC

The ESCON protocol for PPRC has been streamlined with less handshaking, and larger ESCON frames are transmitted between the two ESS.

2.4.0.12 Workload Manager - S/390

Workload Manager (WLM) in GOAL mode provides resources to workloads within a sysplex to ensure that predefined performance metrics are met. The WLM therefore plays an important part in the ESS performance by controlling Dynamic PAVs and enabling I/O Priority Queueing.

2.4.1 Storage performance tools

The ESS performance can be modeled using an IBM internal only tool called DiskMagic, which runs under Windows. Please refer to your IBM representative for use of this tool.

If planning a migration of an existing workload onto an ESS, contact your IBM Storage Specialist to assist you in determining the expected performance.

Chapter 3. Physical installation planning

Successful installation of an ESS requires careful planning. The main considerations when planning for the physical installation of a new ESS are:

- Interface cables for UNIX, NT, AS/400 and S/390 hosts
- Electrical power
- Cooling and airflow
- Physical placement
- Floor loading
- Call home
- Ethernet connection

3.1 Installation planning references

You need to have these books available for reference:

- *Introduction and Planning Guide 2105 Models E10 and E20,* GC26-7294
- *Host Systems Attachment Guide 2105 Models E10 and E20,* SC26-7296

For information on installation and configuration of the IBM 2105, also see the IBM Storage Web site:

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

The book *Introduction and Planning Guide 2105 Models E10 and E20,* GC26-7294, in particular, provides all the orderable feature codes for the models of 2105, as well as information on dimensions, weights, power consumption and cooling that you will need when planning for ESS installation.

Note: The ESS Specialist *help* text provides some detailed technical scenarios to help you perform some of the most common tasks, as well as hints and tips. We suggest that you print the help text to help you plan for the installation and implementation.

3.2 SCSI host attachment

If you are configuring both SCSI and S/390 logical disks in your ESS, you need to locate the ESS close to any SCSI hosts because of the 25 meter length restriction on SCSI cables. ESCON cables can extend up to 3 km to an S/390 host or ESCON switch.

3.2.1 SCSI cables

The ESS uses fast wide ultra-SCSI differential interfaces for attachment to hosts running UNIX variants or Windows NT. This allows data transfer at 40 MB/s over cables up to a maximum length of 25 meters (82 feet).

The ESS SCSI adapters support SCSI-3 protocols, and support up to 64 LUNs per SCSI ID.

The required SCSI cables are 2-byte wide, 69-pin differential, and can be ordered by feature code in standard 10 and 20 meter lengths for several types of host, using Feature Codes 97xx. See Chapter 1, "Introduction" in the *Host Systems Attachment Guide 2105 Models E10 and E20,* SC26-7296 for details of the orderable feature codes. This reference also contains information on the required SCSI adapters for the various types of host that may be attached to ESS. SCSI cables must have 2 - 2.5 meters (6 - 7 feet) of slack under the ESS tailgate. You need to allow for this extra length when planning the installation.

Straight SCSI cables allow for the direct connection between an SCSI adapter located in the host, and an SCSI port at the ESS.The SCSI port at the ESS provides necessary electrical termination of the bus. Termination needs to be provided at the host end also, either by the adapter in the host, or by a separate terminator.

3.2.2 Daisy chaining SCSI adapters

The ESS allows daisy chaining of more than one host adapter on a single SCSI bus. See Figure 6 on page 25 for details. Observe the following recommendations:

- The presence of more than one host on the SCSI bus will lower performance due to the extra bus arbitration needed.
- IBM recommends a maximum of four host initiators attached to ESS on a single SCSI bus.
- You need to obtain the appropriate Y-cable for each extra host. The Y-cables must be 2- byte differential, and must match the requirements of the SCSI adapter in the host.
- We recommend that you use SCSI IDs 7/6/5/4 for the adapters in the hosts. This ensures the most efficient arbitration on the bus.
- The SCSI differential bus must be terminated at both ends.The sum of the length of all branches on one SCSI interface cannot exceed 25 meters. This is the maximum distance between terminators on the bus.
- Ensure that no termination is provided by any intermediate adapters on the bus.
- It is possible to attach a maximum of two ESS SCSI ports to a single SCSI bus. The ESS SCSI ports may be in the same, or different ESS subsystems. Because the ESS ports provide electrical termination, they must be at the ends of the bus.
- Other SCSI targets may share a bus with ESS. The ESS supports a maximum of 16 SCSI devices (including initiators and targets) on a SCSI bus.
- AS/400 does not support daisy chaining with the SCSI adapters used to connect external 9337 disks.

See Figure 6 on page 25 for more detail.

Refer to Chapter 2, "Attaching SCSI Host Systems for Attachment to the IBM Enterprise Storage Server" in *Host Systems Attachment Guide 2105 Models E10 and E20,* SC26-7296, for information on SCSI cabling.

3.2.2.1 Unique target IDs

Care must be taken to ensure that the target ID of each adapter on the bus is unique.

A single ESS SCSI port may contain up to 15 target IDs. You need to ensure that each target ID assigned to logical devices attached to that port is unique on that SCSI bus. Use the ESS Specialist, storage allocation, and configure SCSI ports panel to display internally assigned target IDs for an SCSI adapter port in the ESS.

It is preferable to define first the hosts that will access the ESS on the SCSI bus, then the IDs of all other initiators and non-ESS devices on the bus. On the ESS specialist, configure SCSI ports panel, define your host IDs, then enter the other IDs as unrelated hosts or devices, before adding logical volumes to the port. This will ensure that the ESS assigns non-conflicting IDs for the logical volumes you add. Refer to 15.2.5, “Configure SCSI ports” on page 221 to see the process involved to do this.

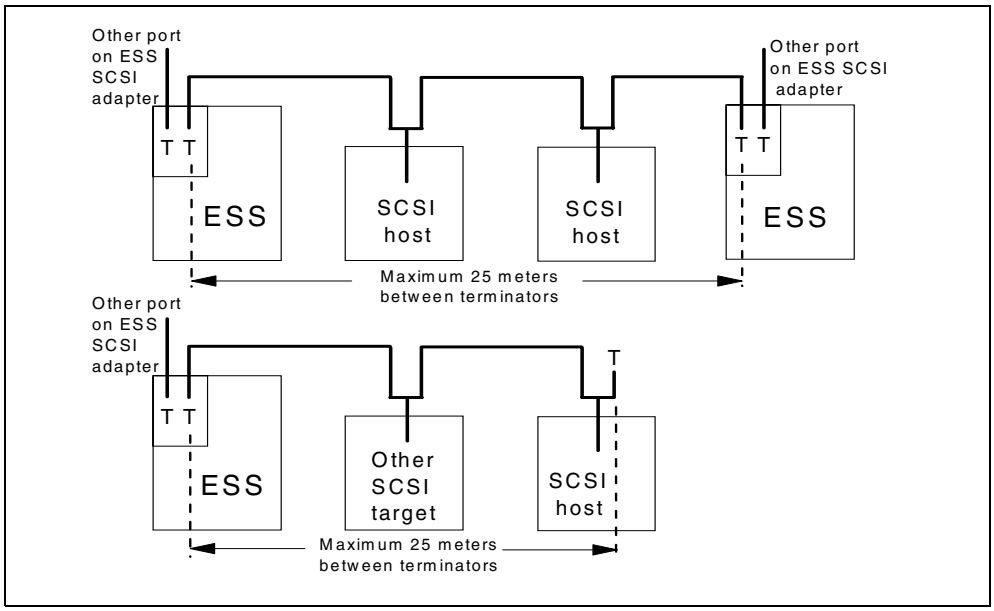

Figure 6. Examples of SCSI cabling configurations

3.2.3 SAN data gateway support

ESS supports attachment to a fiber channel arbitrated loop (FCAL) using the IBM 2108 SAN Data gateway. The 2108 converts FCAL protocols to SCSI. The 2108 must be installed within 25 meters of the ESS to allow connection using SCSI cables. The FCAL link from the 2108 to the UNIX or NT host has a maximum length of 500 meters when a 62.5 micron fiber is used, and 10 km when a nine micron fiber is used.

The ESS allows a maximum of 16 SCSI ports to be attached to the SAN Data gateway. For further information on SAN Data gateway, refer to:

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

3.3 S/390 ESCON cables

ESCON cables may be used to attach the ESS directly to an S/390 host, or to an ESCON director or channel extender. They may also be used to connect to another ESS, either directly or via ESCON Director, for Peer-to-Peer Remote Copy (PPRC). The ESS supports all models of the IBM 9032 ESCON director. It also supports the IBM 9036 channel extender to the length allowed by the 9036, and the 9729 to a maximum distance of 103 km. Customers may wish to use methods of extending ESCON channels that they already have experience with.

ESCON cables come in a standard length of 30 meters, but can be obtained in various lengths. The maximum length of an ESCON link from the ESS to the host channel port, ESCON switch, or extender is 3 km. This is using 62.5 micron fiber, or 2 km using 50 micron fiber.

3.3.1 PPRC considerations

If PPRC is to be used, ESCON connections must be provided between ESS subsystems containing primary and secondary PPRC volumes. This is the case even if there are no S/390 host connections and no S/390 volumes on either ESS. Any ESCON connections needed to support PPRC must be considered when planning the ESS configuration. ESCON host adapters must be installed to support the PPRC links.

For the rules for configuring PPRC links, refer to 12.2, “PPRC” on page 195.

For detailed information on PPRC for the ESS, refer to Chapter 12, “ESS Specialist Copy Services for S/390” on page 195, and Chapter 6, "Copy Services Functions" in *IBM Enterprise Storage Server*, SG24-5465.

TCP/IP Connection for PPRC
If the ESS Specialist Copy Services Web browser interface will be used to manage PPRC, then TCP/IP connections are needed between each participating ESS subsystem, and the PC running the Web browser. At each ESS, the LAN connection is implemented using 10BaseT ethernet.

3.4 Electrical power

You need to take note of the following considerations concerning electrical power for the ESS.

3.4.1 Dual line cords

The 2105 models E10, E20, and the expansion frames all have dual line cords. The two line cords to each ESS frame should be supplied by separate AC power distribution systems. This makes a simultaneous power interruption to both line cords less likely. The ESS continues to operate normally if power to one of the line cords is interrupted.

Because each line cord must be able to support the entire ESS frame in the event that power to the other line cord is interrupted, it is important that the AC supply circuit for each line cord has sufficient current capacity to supply the entire ESS.

See Chapter 3, "Installation Planning" in the *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294 or details on power consumption, voltage and current requirements for 2105.

3.4.2 Single phase power

A single phase line cord cannot be plugged into a normal wall outlet. It needs a circuit capable of supplying 50 or 60 amps (country and location dependent) at a nominal 208 volts or higher, as shown in Figure 7. A wall outlet typically can supply only 20 or 30 amps at 120 volts.

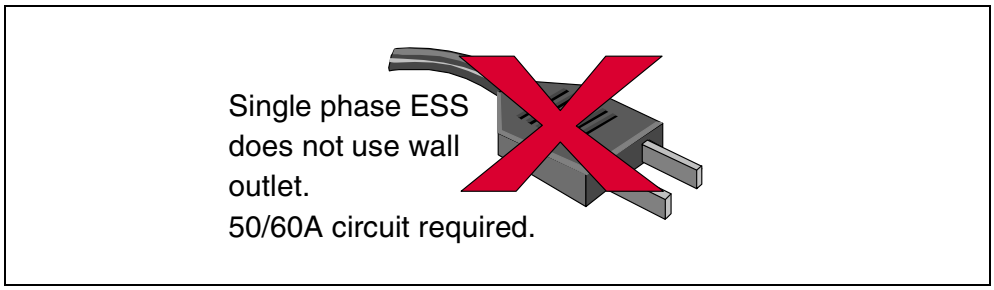

Figure 7. Single phase ESS

Note that the maximum data capacities available in the single-phase models of the ESS are smaller than those of the three phase models. Capacity is limited by the number of disk drives that can be supplied by the power available from a 50/60 amp single phase supply.

3.4.3 Three phase power

There are no special considerations for three phase power to the ESS.

Phase rotation at the line cord connection is not critical, but where a three phase line cord connector is provided, IBM recommends that the receptacle be connected to provide counter-clockwise phase rotation as you view the plug face.

The wall breaker in the circuit supplying the ESS must be capable of withstanding an inrush current of 100 amps at power on time of the ESS. Inrush current at power on, and operating current, are higher than normal if the ESS is operated with power on only one of the two line cords.

3.4.4 Other considerations

3.4.4.1 Power connectors

The cable connectors supplied with various line cords, and the required receptacles are given in Chapter 3, "Installation Planning" in the *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294. The length of the linecords is 14 feet, except in Chicago, where the length is 6 feet.

3.4.4.2 Convenience outlet power

The 2105 base frame contains a 14-foot multiple-outlet extension cord. The outlets are mounted within the rack and are provided for use by the service representative to power any required service tools. You must supply a standard

power outlet for the extension cord. Additionally, three outlets are required for the first installed ESS to supply the ESS Net PC, display, and ethernet hub.

3.4.4.3 Remote power control

If you have ordered the Remote Power Control Facility (FC1001), the 2105 can support up to eight S/360 Remote Power Control (RPC) interfaces to allow remote power control from S/390 hosts. If you plan to use remote power control, you need to obtain the RPC cables in the required lengths.

3.5 Physical location and floor loading

In this section, we discuss restrictions on where the ESS can be physically located due to cable length, floor loading, and service clearance requirements.

3.5.1 Transit

You need to consider the size and weight of the ESS rack in its shipping container, as well as how it will be moved from the loading dock to the final installation location.

3.5.1.1 Arbo crate

The weight of a fully configured 2105 frame ranges from 2160 pounds (980 kg) to 2910 pounds (1320 kg). Additionally, each frame is shipped in a wooden Arbo crate that weighs 395 pounds (179 kg) empty. Refer to Table 6 on page 28 for dimensions of the Arbo crate.

Table 6. Dimensions and weights for 2105 racks

Model	Dimensions, inches (cm)	Maximum Weight (fully configured), pounds (kg)	Dimensions of Arbo Crate, inches (cm)	Maximum Weight in Arbo Crate, pounds (kg)
2105 Models E10 and E20	Height: 70.7 (179.6) Width: 54.4 (138.3) Depth: 35.8 (91)	2590 (1175)	Height: 81.0 (206) Width: 62.0 (158) Depth: 41.7 (106)	2985 (1354)
2105 Expansion frame	Height: 70.7 (179.6) Width: 54.4 (138.3) Depth: 35.8 (91)	2910 (1320)	Height: 81.0 (206) Width: 62.0 (158) Depth: 41.7 (106)	3305 (1500)
Note: Height of installed 2105 includes casters and covers, but excludes the top-hat. With top-hat, height is 75.3 inches (191 cm).				

The 2105 should be protected from any mechanical shocks during transit. It is preferable to leave the 2105 frame inside the Arbo crate until it is at the final install location, if possible. If the ESS rack must be moved from the loading dock without the Arbo crate, on its own casters, take precautions to minimize shocks caused by passing over door gaps in elevators, ramps and so forth.

Caution: A fully configured frame in the crate can weigh up to 3305 pounds (1500 kilograms). This is the approximate weight of a medium-sized car. Ensure that the path for the 2105 from the loading dock to the final location, including any ramps and elevators, can accommodate frames of this size and weight. *This is a safety issue*.

IBM recommends that any ramp used to move a 2105 have a maximum gradient of 12.5 degrees, and be no more than 12 feet (3.6 meters) in length.

3.5.2 Physical placement

If you plan to attach the ESS to SCSI hosts, the ESS needs to be located within close proximity of the SCSI hosts to allow connection within the 25 meter SCSI cable length limit.

The 2105 requires service clearances of 32 inches (81 cm) at the front and 45 inches (114 cm) at the rear. Racks can be placed side by side if floor loading restrictions allow this. If later installation of an ESS expansion rack is planned, you should allow room for the additional rack adjacent to the base 2105 to which it will attach. The expansion frame can be installed to either side of the base frame.

When an expansion rack is attached to a 2105 Model E10 or E20 base frame, a spacer section a few inches wide is inserted between the two frames. This section serves two purposes:

- To provide a place to loop the excess length in the SSA cables connecting the expansion rack to the base frame
- To keep the floor load within the required limit when both frames are fully configured

The ESS needs a single cable access hole located under the front of the machine at center for the ESCON and SCSI cables and two line cords. See Chapter 3, "Installation Planning" in the *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294 for information on the size and location of cable access holes needed in the floor. We suggest that the ESS be located over three floor tiles, and that the center tile have approximately 8" cut off. This creates a hole 24" by 8", which is wide enough to span the cable restraints at the bottom front of the ESS without weakening the load bearing floor tiles.

3.5.3 Floor loading

The 2105 does not require a raised floor. However, if multiple cable connections will cause a problem in a non-raised floor environment, it is suggested that a raised floor be used. Ensure that the ESS is included in the raised floor equipment plan for the site.

A fully configured ESS with an expansion frame is very heavy (5500 pounds, 2495 kg). If it is to be attached to Windows NT or UNIX-based hosts, it is likely to be located on a floor designed for the smaller loads this type of computer equipment usually involves. Ensure that not only is the building floor loading capacity sufficient for the ESS configuration you plan to install, but that the raised floor is also rated to support the load. Because of the high point load on each floor tile supporting part of the weight of the ESS (727 lb, 430 kg for each of the four casters on a fully configured 2105 expansion frame), the floor tiles and support structure must be capable of supporting such loads. *This is a safety issue*. See Figure 8.

It may be necessary to install additional stanchions under the floor tiles that support the ESS.

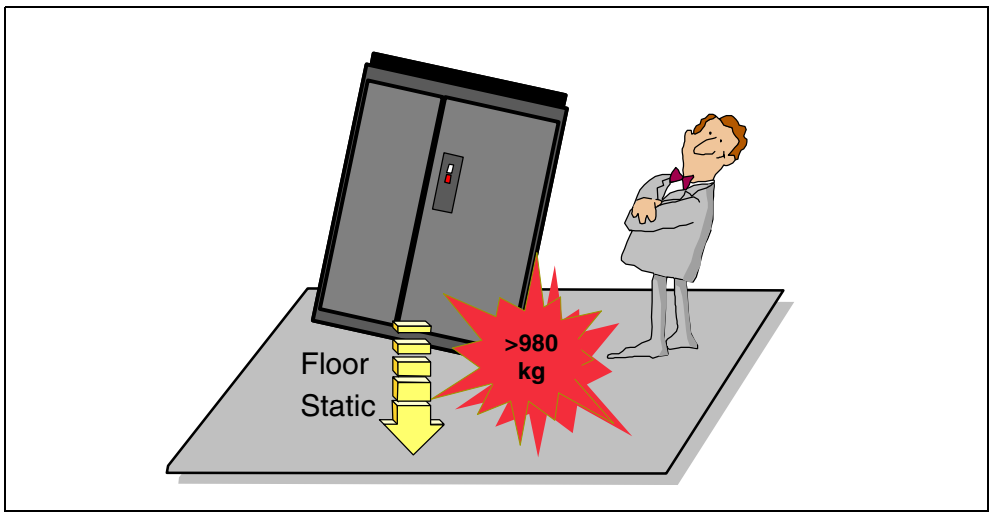

Figure 8. Caution: Floor loading

If installing more than one ESS frame, you may need to separate the frames by a specified distance to keep within the load capacity of the floor.

See Chapter 3, "Installation Planning" in the *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294 for additional weight and floor loading information.

3.6 Cooling and airflow

Adequate cooling is critical to the long term reliability of electronic equipment in general, and to hard disk drives in particular.

3.6.1 Operating temperature

The 2105 should be maintained within an operating temperature range of 20 to 25 degrees Celcius (68 to 77 degrees Fahrenheit). The optimum temperature is 22 degrees Celcius (72 degrees Fahrenheit). We strongly recommend that you avoid running the ESS, or any disk storage equipment, at temperatures outside this temperature range.

Humidity
The humidity should be maintained between 40% and 50%. The optimum operating point is 45%.

3.6.2 Airflow

Adequate airflow needs to be maintained to ensure effective cooling. We recommend that two full-vented floor tiles be located at the front of each 2105 frame, and two vented tiles at the rear. The ESS takes in cooling air through vents in the front and rear covers, and exits it through the top of the frame. The hole in the floor provided for cable access to the ESS will not pass sufficient cooling air. It may be filled by cables, or be blocked with a fire-retardant pillow.

Adjustments may need to be made to air conditioning equipment or ducting to ensure a good flow of cool air up through the floor vents. The 2105, like any disk

subsystem, works best and most reliably when the temperature is maintained near the optimum.

Avoid placing racks of 2105 in confined corners of a room where there is insufficient above floor airflow to remove the heat. See Chapter 3, "Installation Planning" in the *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294 for information on required temperature and humidity, and heat output of the 2105 models.

3.7 Call home

The setup of the modem and modem expander, and the connection to the ESS, will be performed by the IBM service representative during installation of the ESS.

Why is call home needed?

- To allow the ESS to automatically notify IBM when service is required.
- To allow remote access by IBM support specialists.
- The Call Home modem is used by the ESS to contact a pager if you elect to use the pager notification function.

3.7.1 Call home feature codes

If this is the first ESS, or other IBM device, that supports the call home support (such as a VTS, or a 3494) in this location, you need to order FC 2715. This feature code provides one Microcom Deskporte FAST+ modem, one Western Telematic APS-16 Asynchronous Switch (modem expander), and the cables to interconnect them with the ESS. Refer to Figure 9 for details of the interconnections.

FC 2715 also provides the ESSNet PC, monitor, and ethernet hub.

For each of the next six ESS installed in the same location, order FC 2716. This feature code provides just the cables to connect the ESS to the existing modem expander. Each modem and modem expander can support seven ESS, using one customer-supplied phone line.

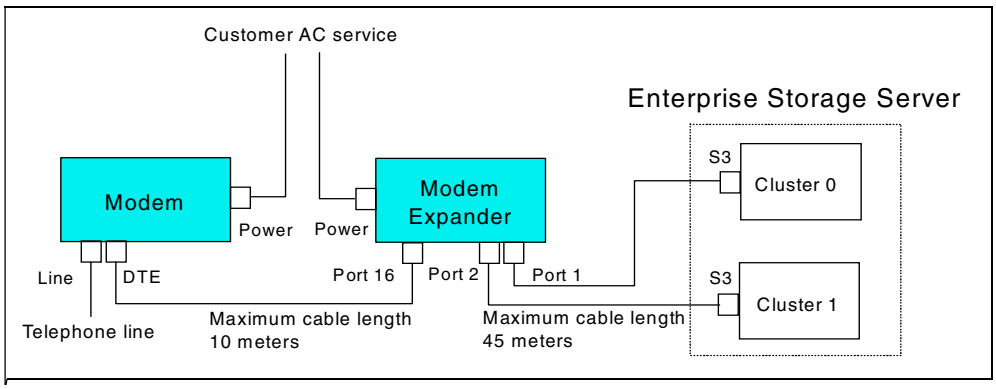

Figure 9. Call home configuration

3.7.1.1 Physical considerations for call home

The modem expander must be installed within 45 meters (150 feet) of each ESS to which it is attached, and within 3 meters (10 feet) of the attached modem. The modem expander connects to the S3 serial port of each cluster in the ESS.

A power outlet must be supplied to provide power for the modem expander and modem. You must supply the phone cord from the modem to the phone line jack.

The modem supplied with FC 2713 is fully homologated for use in your country.

Note: The customer provided phone line must be analog, not digital.

3.8 Web client and Ethernet connection

The two clusters in the ESS communicate with each other through a 10baseT Ethernet connection. With the addition of a hub, this Ethernet is used also as a link to the Web client, and can be connected to the user's LAN to provide connectivity from anywhere using TCP/IP.

3.8.1 ESS Specialist

The ESS Specialist is a software application that runs in the 2105 clusters. It is the interface provided for the user to define and maintain the configuration of the 2105 subsystem. The ESS Specialist can be accessed using a Web browser running in a Web client PC attached directly to the 2105 via ethernet, or in a remote machine if the 2105 is connected into the user's network. For detailed information on the ESS Specialist, see Chapter 5, "Enterprise Storage Server Specialist (ESS Specialist)" on page 69.

3.8.1.1 ESS Specialist Web client requirements

The speed of the Web client used to access the ESS Specialist functions in the ESS will have a direct effect on the time taken to load the ESS Specialist, and how long it takes to refresh pages in the browser. This is due to Java code being downloaded from the ESS to the Web Client. The minimum configuration for the Web client is a Pentium 166 MHz with 32 MB RAM. For best usability you should use a fast Pentium II machine. We recommend a Pentium II 233 MHz or higher, with 96 MB RAM.

Preferably, the Web client will have a 10BaseT Ethernet adapter installed. With this, only a simple Ethernet hub and three cables are needed to connect the Web client to both clusters of the 2105 in a minimum configuration. If the ESS is connected to the user's LAN, the Web client can have either token ring or Ethernet, to match the existing LAN.

3.8.1.2 Web browser requirements

The ESS Specialist Web user interface can be viewed using any browser that fully supports the Java 1.1 standard. We recommend Netscape Navigator 4.5 running on Windows NT 4.0. This browser can be downloaded free from Netscape's Web site:

`http://home.netscape.com/computing/download/index.html`

You can also use one of the following:

- Netscape Navigator or Communicator Version 4.06 or later, running on Windows 95, Windows 98 or Windows NT Version 4.0.
- Microsoft Internet Explorer (MSIE) Version 4.0 or later, running on Windows 95, Windows 98, or Windows NT Version 4.0. MSIE at this level has the necessary support for Java 1.1, but you may experience the following problem:

  When running some versions of MSIE, you may receive error message 1196: "Unencrypted update request" when you attempt to change the ESS configuration. This is related to a problem with handling secure transactions when the browser is in a different domain than the ESS. If you experience this problem, either retry the operation using another browser, or use a client machine in the same domain as the ESS.

  The proxy settings may affect the operation of MSIE in relation to the 1196 message. You may need to disable the proxy server settings in MSIE to avoid this message.

If you experience a browser hangup or crash in Netscape or MSIE, restart the browser and try the operation again. These problems are related to the Web browser, not the ESS Specialist application or Windows operating system.

3.8.2 Minimal connection to the Web client

The minimum configuration needed to enable use of the Web client uses an Ethernet hub with a minimum of three ports, and three 10BaseT cables, connected as shown in Figure 10 on page 34.

The internal 10BaseT cable connecting the two clusters together is removed by the IBM service representative during ESS installation, and a 10BaseT cable is connected between each cluster and the Ethernet hub. A third cable connects the hub to the Web client.

It is important that both clusters be connected to the Ethernet hub so that connectivity between them is maintained.

The IP addresses for the two clusters are assigned by the IBM service representative during installation, using the interface on the service terminal.

The ESS Specialist can be run by specifying the IP address of either cluster directly as the universal resource locator (URL) in the browser. For example, the IP address of the clusters might be assigned as 192.168.0.1 for cluster 1, and 192.168.0.2 for cluster 2. You would use one of these as the URL.

3.8.2.1 Multiple ESS on the Ethernet private LAN

Additional ESS clusters can be connected to the Ethernet hub, provided that they have unique TCP/IP addresses. A single Web Client can be used to configure all connected ESS subsystems.

Ethernet hubs can be daisy chained to connect extra ESS clusters beyond the capacity of a single hub.

If two or more ESS in a site participate in PPRC, and this is to be managed using the Copy services function of the ESS Specialist, they need to be connected to the same private LAN.

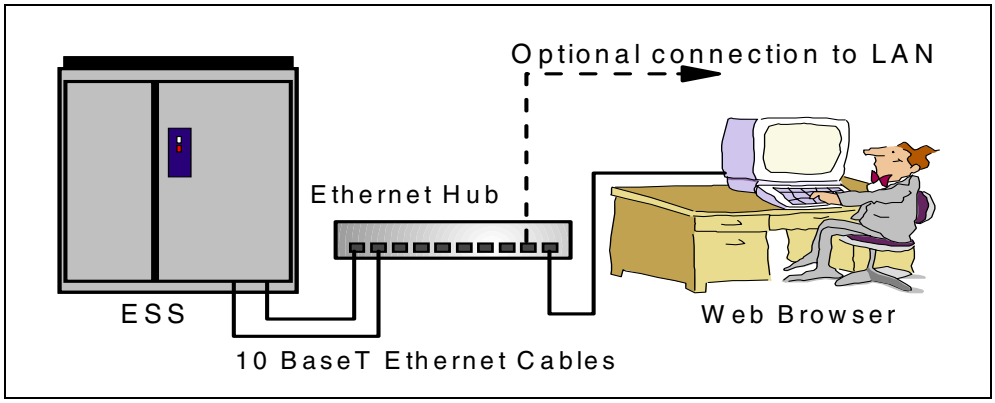

Figure 10. Ethernet private LAN configuration

3.8.3 Connection to Local TCP/IP Network

One or more ESS can be connected to the local IP network by connecting the Ethernet hub to the network. See Figure 10 on page 34 for details. The user must supply any cables, bridges or hubs necessary to make the connection.

Why connect to the TCP/IP network?

- You may need to use the ESS Specialist, or VS Expert, from a remote location.
- You may want to use e-mail and/or SNMP notification.
- Where primary or secondary volumes of PPRC pairs reside in ESS subsystems in remote locations, you may wish to manage the PPRC using the Copy services function of the ESS Specialist. This requires that the Web browser used for the ESS Specialist has TCP/IP connectivity to each local and remote ESS participating in PPRC.

3.8.3.1 Additional requirements

There are additional requirements if an ESS is to participate in the user's network:

- Two static IP addresses for the clusters must be assigned by the user to be consistent with the existing network.
- A DNS supported host name must be provided for each cluster.

Use the *Communication Resources Work Sheet* at the back of *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294, to specify the TCP/IP address, network mask, default gateway address, hostname, and nameserver TCP/IP address for the 2105 clusters.

The ESS Specialist can now be used from any TCP/IP-connected PC running a suitable browser. For example, if cluster 0 is assigned the host name essclust0, and the local IP domain is mydomain.mycompany.com, then the URL to use in the browser is `http://essclust0.mydomain.mycompany.com`.

Security

IBM recommends that any ESS connected to a user's LAN be protected from unauthorized access by use of a firewall. The ESS Specialist provides password protection and three levels of user authority to manage access by users. If

required, ESS Specialist can be configured to accept connection only from defined TCP/IP addresses.

3.8.4 ESSNet

ESSNet is provided with the first ESS installed in a site. This is a Windows NT desktop machine preconfigured with the Netscape browser. It provides the Web client function and will allow connection into the customer's LAN. The machine may connect directly to the Call Home modem, eliminating the need for a separate Modem Expander. The ESSNet PC, Ethernet hub and cables are provided with the call home modem and modem expander as part of Feature Code 2715. ESSNet supports up to seven ESS subsystems.

3.9 Defining communication functions

Before the ESS is installed, decisions should have been made about which alerting and communication functions will be used, and the information needed to configure these functions.

3.9.1 Communications worksheet

Refer to Chapter 4, "Configuration Planning" in *Introduction and Planning Guide 2105 Models E10 and E20,* GC26-7294, and use the *Communication Resources Worksheet* at the back of the book to define the needed values.

Customer and IBM personnel will use the information you enter in the work sheet during the installation process. For example, the IBM CE will enter TCP/IP values and enable the ESS Specialist during the installation using a terminal connected to one of the serial ports in the clusters. After this, the ESS Specialist can be accessed over the TCP/IP LAN.

The categories of information provided in the worksheet are:

- Time zone and date configuration
- TCP/IP configuration
- Customer information
- Customer pager information
- ESS Specialist configuration
- Call home/remote services
- Modem configuration
- E-mail configuration
- SNMP configuration

3.9.1.1 TCP/IP configuration

If a private LAN configuration is to be used, that is, the private LAN comprising the ESS clusters, Ethernet hub, and Web client only (ESSnet standard configuration), then no connection will be made to the user's LAN. See Figure 10 on page 34. The TCP/IP addresses assigned to the clusters and Web Client can have any desired values. No hostnames, nameserver, or gateway need to be specified. Standard values for the TCP/IP addresses used in ESSnet are 172.31.1.1 for cluster 1, 172.31.1.2 for cluster 2, 172.31.1.250 for the web client.

In the majority of cases, the ESS network will be connected to the customer's TCP/IP network via a hub, router or bridge. All the information in the worksheet section "TCP/IP Configuration" needs to be specified.

Note: If your LAN environment is something other than 10baseT, you need to ensure that the hub or switch into which you connect the ESS is capable of operation at 10Mb/s.

Copy services primary and backup server TCP/IP configuration
If you plan to use the Copy Services function of the ESS Specialist to manage PPRC or FlashCopy, you need to specify the host name and TCP/IP address of the Copy services primary server and the backup server. The primary server and backup server are ESS storage clusters, usually in two different ESS subsystems.

This information is used by the CE using the service terminal to configure each participating ESS for Copy services. If the Private LAN configuration is used, there will be no DNS available. The host names must also be added to the HOSTS file by the CE, using the service terminal, to allow host name resolution. This needs to be done for each participating ESS subsystem.

The Copy services primary and backup server TCP/IP configuration is not included in the *Worksheet.*

3.9.1.2 ESS Specialist configuration

Users who will be authorized to use the ESS Specialist are specified here, along with the authority level:

- Viewer
- Configurator
- Administrator

The administrator can define other users, their authority, and passwords using the ESS Specialist.

Note: The default initial administrator user ID is *storwatch*. The password is *specialist*. At the first logon to the ESS Specialist, one or more new administrator user IDs and passwords should be defined. When you define one or more administrator user IDs, the storwatch user ID is automatically deleted. To prevent unauthorized access, the ESS Specialist will not allow you to define the storwatch user ID again. This is important because the ESS Specialist can be used to reconfigure disk arrays and consequently erase customer data. However, if you delete the last administrator userid, the storwatch userid will be restored.

3.9.1.3 Call home and remote services

In this section of the worksheet you enable or disable these functions:

Modem incoming calls IBM recommends that incoming calls be enabled. This allows diagnosis of problems by a remote specialist. Refer to 4.3.1.1, "Remote support" on page 54.

Modem outgoing calls Outgoing calls must be enabled to allow the ESS to place a call to IBM when service is required.

E-mail over LAN This function would normally be enabled. It allows specified people to receive e-mail notification of events such as a modem incoming call being received.

Pager messages The ESS can send messages to your pager if you enable pager messages here.

The required information such as pager numbers and remote telephone numbers is specified in the worksheet.

3.9.1.4 E-mail configuration

Here you specify:

- **Maximum error notification count per problem**

  The ESS sends 0 to 9 (the number you specify) error notifications to each recipient for each problem. The recipients include the e-mail addresses you enter, and the pager numbers.

- **Connection number**

  The ESS will call home to IBM this number of times for each problem.

- **E-mail destinations**

  Enter here the full e-mail address for each destination. For each recipient, check which notification the destination will receive:

  - Errors
  - Information
  - All
  - None

IBM strongly recommends that you list at least one e-mail address. This allows the ESS to notify you when it receives an incoming modem call from IBM service personnel.

Note: The e-mail address must be inside the customer's intranet (behind the firewall).

Chapter 4. Availability

In this chapter, we describe ESS recovery from internal failures. Also discussed are maintenance, error notification, and configuration considerations to ensure maximum availability.

4.1 ESS internal recovery

This section describes some possible ESS internal failures, their effect on availability, and the impact of repair. In most cases of internal failure, recovery by the ESS is automatic and requires no user intervention.

The main categories of possible ESS component failure are:

- Host connection failures
- Cluster failures
- Device adapter failures
- Disk failures
- Power and cooling failures

The Enterprise Storage Server is designed to provide uncompromised data integrity in all failure situations. If the ESS is properly configured, data continues to be available in the event of any single component failure.

4.1.1 Host adapters

A failure in an ESCON or SCSI host adapter may affect one or both of the ports in the adapter.

Multiple host paths recommended
We recommend that wherever possible, you configure two or more paths from each SCSI or ESCON host to different host adapter bays in the ESS. This provides redundancy in the event of a cable failure, or an adapter failure in the host or the ESS. It enables a failing ESS host adapter or cable to be replaced nondisruptively. Additionally, performance can be significantly improved by configuring multiple physical paths to groups of heavily used logical disks.

4.1.1.1 ESCON

For ESCON, there will normally be four or eight paths in a path group between each S/390 host and a S/390 LCU in the ESS. The loss of a single path usually can be tolerated until a repair can be made, with little impact on performance if four or more paths are configured.

4.1.1.2 SCSI

For SCSI, if the Data Path Optimizer (DPO), or similar software that supports multiple paths is running on the host, and two or more paths are configured from the host to the ESS as recommended, a single host adapter failure will not interrupt access to the LUNs supported by the failing path(s).

Reconfigure spare SCSI adapter
If a failing SCSI adapter is the only path configured between a host and any logical volumes, then data access is interrupted. If a spare SCSI port is present in

the ESS, the logical volumes can be reassigned to this adapter using ESS Specialist. The following steps establish a path between the host(s) and the logical volumes using a new adapter port:

- Associate the host(s) with the new port on the Configure SCSI Ports panel. Refer to 15.2.5, “Configure SCSI ports” on page 221.
- Add the desired logical volumes to the new port. See 15.2.7, “Modify fixed block volume assignments” on page 231.

The host SCSI cable is moved to the spare adapter port to restore access until a repair is made. If only one port on the adapter is failing, this needs to be done only for the hosts and volumes normally accessed through that port. However, repair of the adapter involves interrupting access through all four adapters in the IO bay.

4.1.1.3 CPI cables

A problem in one of the cables interconnecting the storage clusters with the host adapter bays can prevent communication between a cluster and a host adapter bay. A logic problem in a HA bay or cluster can have a similar effect. A HA bay must be able to communicate with both clusters because each write operation involves a transfer of data to NVS in one cluster and to cache in the other cluster.

A logic failure on a host adapter bay planar board could prevent all four host adapter cards in that bay from operating. In these situations, the entire HA bay is quiesced and is unavailable. If the ESS is configured as recommended with multiple physical paths, access to data continues.

4.1.2 Clusters

Cluster failures can be caused by logic, or the internal disk drive used by the cluster to store microcode and configuration data. If a cluster failure makes the cluster inoperative, or prevents access from the cluster to disk arrays, or to NVS, the ESS automatically takes recovery action by transferring the functions performed by the failing cluster to the opposite cluster. The other cluster takes over all subsystem management, and supports access from the host adapters to all the logical volumes in the ESS.

Cluster failover can be manually induced by the CE to perform functions such as installation of microcode updates, or to replace components when a failure does not cause automatic failover, for example the ethernet adapter, diskette drive, or serial ports.

4.1.2.1 Normal cluster configuration

During normal operation, each cluster manages its own logical volumes, communicating with the HAs through the CPI hardware, and with the disk arrays through its own DAs. Data is stored in 3GB of cache installed in the cluster, and write data is stored in the NVS of the opposite cluster. The CPC (processors and associated logic) in a cluster, together with the associated NVS in the other cluster comprise a subsystem. Refer to Figure 11 on page 41.

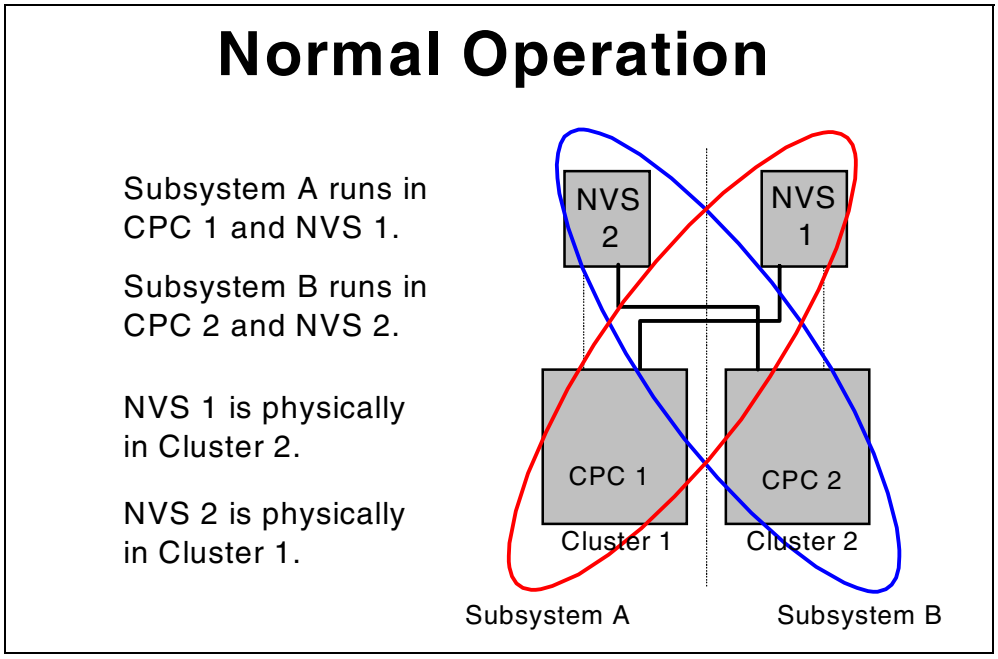

Figure 11. ESS cluster configuration for normal operation

4.1.2.2 Cluster failover

If a cluster fails in such a way that both cache and NVS in the cluster are not available, a copy of modified data for both subsystem A and subsystem B still exists in the NVS and cache of the other cluster. No data is lost.

The failover process is shown in Figure 12 on page 42. In this example, cluster 1 has failed. CPC 2 normally communicates with NVS 2 which is physically located in cluster 1. NVS 2 is no longer available. Because the only copy of modified data for subsystem B is in the volatile cache of CPC 2, this data must be destaged quickly to disk.

The sole copy of modified data for subsystem A is in NVS 1. During failover, CPC 2 establishes communication with NVS 1, and copies the modified data into cache. CPC 2 then takes over the logical volumes normally managed by CPC 1.

Subsystem A and subsystem B are now managed by CPC 2, using the cache in CPC 2, and NVS 1, both physically located in cluster 2. Access to the disk arrays is through the DAs in CPC 2.

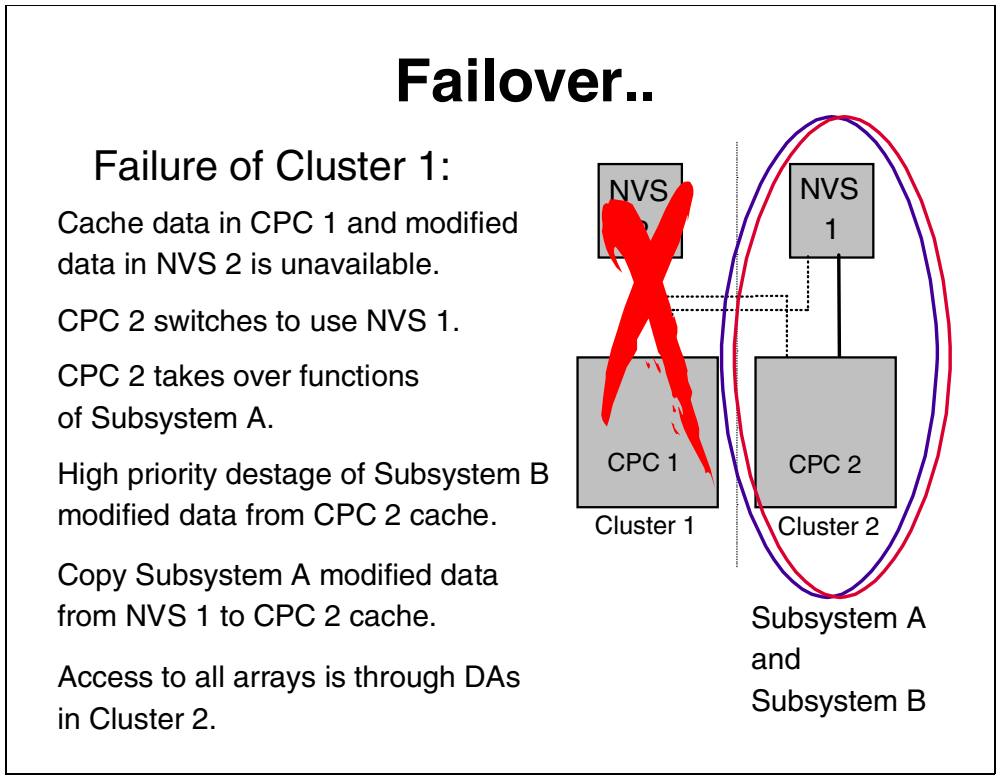

Figure 12. ESS cluster failover process

Because one cluster is now doing the work of two clusters, there may be a noticeable effect on subsystem performance, but read and write access is maintained to all logical devices in the ESS.

While cluster 1 is offline, it can be serviced while subsystem operation continues.

4.1.2.3 Cluster failback

When servicing of the failed cluster is complete, the failback process is used to return to the normal configuration. Figure 13 on page 43 shows the failback process.

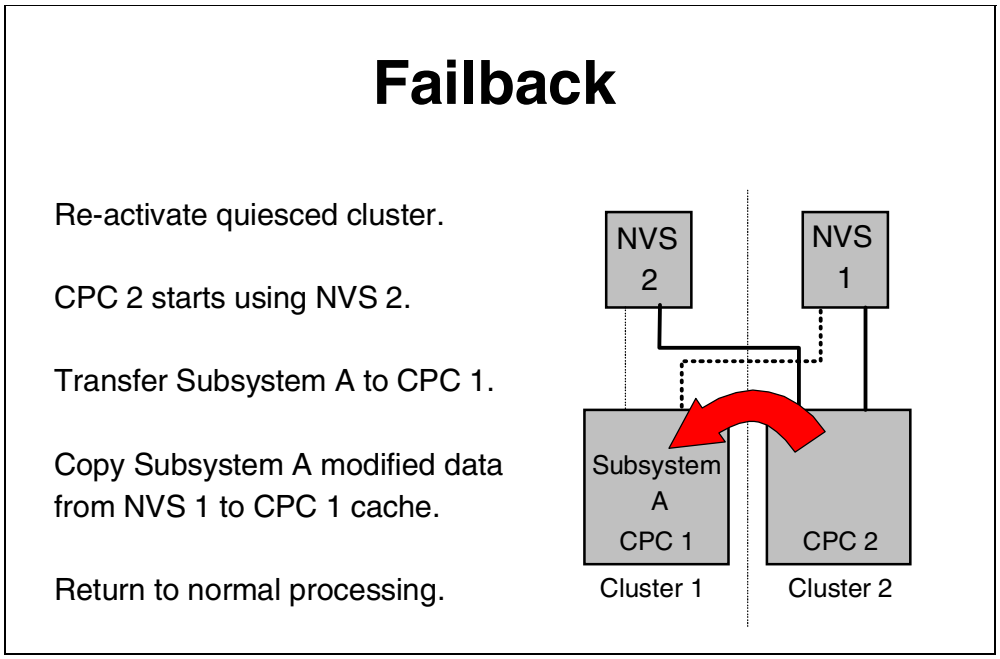

Figure 13. ESS cluster failback process

The cluster resumes control of its logical subsystems, and the associated disk drives. The host adapters are instructed to direct IO operations for those LSSs to the reactivated cluster.

4.1.3 Cache

Each storage cluster has a high performance cache consisting of 3GB of ECC Synchronous Dynamic RAM (SDRAM). The Error Correcting Code (ECC) provides single-bit, double-bit and multiple-bit error detection.

When data is written to cache, the memory controller calculates check bits and stores them along with the data. When data is read from cache, the memory controller generates the check bits again from the data, and compares these with the stored check bits. If there is a difference between the two sets of check bits, an error has occurred. The error may be:

A single data bit The error is corrected automatically by hardware using the check bit information.

Two or more data bits A check condition is generated. The data is automatically recovered from either NVS (write data), or a disk drive (read data).

One or more check bits A check condition is generated. Data recovery is not needed.

Errors occurring on cache cards are logged by the cluster. If analysis of cache errors indicates that a cache card needs to be replaced, the ESS will call home.

4.1.4 Device adapters

At any given time, a RAID 5 array or JBOD is under the control of a single device adapter; the one corresponding to the LSS or LCU to which the array or JBOD

belongs. The rank is normally accessible only by the cluster containing that DA. If the DA has a failure, management of its logical subsystems can be transferred to the opposite cluster by cluster failover.

This results in the host adapters being instructed to direct I/O for the logical devices that are normally managed by the cluster with the failed DA to the other cluster, and the other cluster takes control of all RAID arrays and JBODs on the SSA loops.

This is physically possible because each SSA loop is attached to both device adapters in a DA pair, and either device adapter can manage access to all disks on both loops if the other DA fails.

Non-disruptive repair of the DA can be done using cluster failover/failback in the normal manner for repair of cluster components.

4.1.4.1 SSA loop problems

An SSA loop can be affected by a problem in a cable, or in the serial interface chip (SIC) in one of the disk drives, which may prevent transmission of data, commands, and status past a point in the loop in one or both directions. The DAs each have one remaining path around the loop to all disks (apart from the failing disk) on the loop and use it exclusively. The failure reduces the maximum bandwidth of the loop, and may have a measurable effect on the performance of logical devices located in disks on the loop if they are heavily utilized.

4.1.5 Disk drives

A disk drive can be configured as a member of a RAID 5 array, or as a non-RAID disk (JBOD).

For a disk drive failure, continued access to data depends on whether the failed drive is configured as part of a RAID 5 array.

4.1.5.1 Disk drive failure in RAID 5

When a disk drive module (DDM) fails in a RAID 5 array, the DA starts an operation to reconstruct the data on the failed drive onto one of the hot spare drives on the loop. It does this by reading the corresponding data and parity in each stripe from the remaining drives in the array, and performing an exclusive-OR operation to recreate the data, then writing this data to the spare drive.

Effect on performance

While this data reconstruction is going on, the DA can still service read and write requests to the array from the hosts. There may be some degradation in performance while the sparing operation is in progress, because the DA and loop resources are being used to do the reconstruction. Additionally, any read requests for data on the failed drive (1/7 of the total requests to a 6+P+S array, 1/8 for a 7+P array) require data to be read from the other drives in the array to reconstruct the data. The remaining requests are satisfied by reading the drive containing the data, in the normal way.

Similarly, write operations are affected if the target for the write is the failed drive. RAID 5 recovery will be necessary for 1/7 of the total requests to a 6+P+S array, 1/8 for a 7+P array.

For sequential operations (read and write), data from the other drives needed to reconstruct data for a failing drive may exist in the DA cache, or may be read to satisfy the read or write request. Therefore few, if any extra read operations are needed for data reconstruction, and the impact of a failing drive is smaller.

Performance of the RAID 5 array returns to normal when the data reconstruction onto the spare device completes. The time taken for sparing can vary, depending on the workload on the array, SSA loop, and DA pair.

Drives are hot pluggable
Replacement of the failed drive does not affect operation of the ESS, because the drives are fully hot pluggable. In the event of a disk drive failure, replacement of the failed drive should be done in a timely manner, to ensure that a spare drive continues to be available on the loop.

Predictive Failure Analysis (PFA)
The drives used in the ESS incorporate Predictive Failure Analysis (PFA), and can anticipate certain forms of failure by keeping internal statistics of read and write errors. If the error rates exceed predetermined threshold values, the drive will be nominated for replacement. Because the drive has not yet failed, data can be copied directly to a spare drive. This avoids using RAID 5 recovery to reconstruct all the data onto the spare drive.

The ESS will call home to notify IBM of the need for a repair action.

4.1.5.2 Disk drive failure in non-RAID
For JBOD disks, the situation is as for any non-RAID protected drive. This means the data on the drive is unavailable, unless software mirroring is in use to provide an alternate copy on another device for each and every logical volume residing on the failed drive.

The broken drive can be replaced without affecting the operation of other disk drives on the SSA loop. After the replacement of the drive, data must be restored to the logical volumes on the drive from backups, using established procedures.

Note: A disk drive removed from an ESS may still be functional. It may contain confidential customer data. A drive from a RAID array contains only fragments from any one logical device. A drive that was configured as a JBOD can contain one or more entire LUNS, or several 3390 volumes. You may consider having a procedure in place to manage the secure disposal of removed disk drives.

4.1.5.3 RAID 5 recommended for data availability
We recommend that RAID 5 be used because this provides for continued data availability in the event of a drive failure without needing to provide software mirroring. Software mirroring requires twice the storage capacity for each logical volume, and uses host and ESS resources for dual write operations. Where high performance is needed for random reads, JBOD may be needed. For reads and sequential operations, any performance penalty on random write operations should be mitigated by the large cache provided in each storage cluster, the cache in the device adapters (DAs), and the data striping across multiple drives in the RAID 5 array.

4.1.6 Power and cooling

The ESS is designed so that there is no single point of failure for power that will prevent access to data.

4.1.6.1 AC power supplies

Interruption of AC power to one of the line cords is equivalent to failure of one of the two primary AC power supplies.

Because each DC power supply for the electronic bays and the disk drive cages is supplied from both AC power supplies, the failure of an AC power supply or interruption of power to one line cord leaves all DC power supplies operating.

The electronics bay fans, the disk drive cage fans, and the drive cage power planar fans all operate on DC from the DC power supplies. The fans continue to operate normally.

Either primary AC power supply can be replaced concurrently. Both clusters continue to operate normally during AC power supply replacement.

To minimize the effect of any power interruption or disturbance, the two line cords for each ESS frame should be connected to separate AC distribution systems. The loss of AC power on one of the line cords will leave the ESS operational using the remaining line cord.

4.1.6.2 Power interruptions

If power to both ESS line cords is interrupted, the two batteries integrated into the ESS rack maintain power to the clusters and disk drives for a period sufficient to allow a destage of write data from cache to the drives, and for an orderly shutdown of the ESS. Because power interruptions can occur two or more times within a short period, the batteries have sufficient capacity to support more than one controlled shutdown. The length of time that the batteries can support the rack depends primarily on the number of disk drives installed. Because the batteries conform to the 2N philosophy, either battery alone can support orderly shutdown. The batteries can be replaced concurrently by the IBM CE.

Emergency Power Off (EPO) switch

In the case of an emergency, the ESS can be powered off using the Emergency Power Off (EPO) switch on the front panel. This switch causes immediate cutoff of power to all parts of the 2105 for safety. No controlled shutdown takes place. Host systems using the ESS will be affected.

Because no destage of write data can be done, and the internal microcode is not sequenced down normally, recovery actions may be needed when power is restored to the ESS. The ESS will perform internal recovery and destage data from NVS to disk when it is brought up after power is restored. The EPO switch should be used only in an emergency.

4.1.6.3 DC power supplies

There are three groups of DC power supplies in the ESS base frame:

- Three electronic bay DC supplies for cluster 1, and HA bays 1 and 2
- Three electronic bay DC supplies for cluster 2, and HA bays 3 and 4
- Three storage cage DC supplies for one disk drive cage, or five DC supplies for two disk drive cages

The DC power supplies in the ESS are designed using an N+1 philosophy. Any single power supply failure in a group leaves sufficient power capacity remaining to continue uninterrupted operation. For example, the cages containing the disk drive bays each require two DC power supplies. For a single disk drive cage, three DC supplies are provided. For two cages, five supplies are installed. In the event of a DC power supply failure, the disk drives continue to function normally, powered by the remaining DC supplies. The broken DC power supply can be replaced with no impact to operation of the ESS.

4.1.6.4 Cooling fans

Cooling fans in the ESS base frame are categorized into several groups:

- Electronics bay fans
- Storage cage fans
- Storage cage power planar fans
- AC supply fans

Each group of electronic bay fans cools an SMP (cluster) processor and associated logic, and two host adapter bays. The storage cage fans located at the top of the ESS frame cool the disk drive bays in the front and rear of the disk drive cages. The storage cage power planar fans cool the DC power supplies for the disk drive cages. The AC supply fans cool the AC supplies in the lower part of the ESS base frame. Refer to Figure 14 on page 48.

In accordance with the N+1 philosophy of the ESS, the failure of any one fan in a group leaves sufficient cooling capacity for normal operation to continue.

If a fan fails, the remaining fans in the group increase speed to maintain normal airflow.

The cooling fans in each group can be individually replaced without the need to quiesce or failover any ESS component. Operation of the ESS is unaffected during the repair.

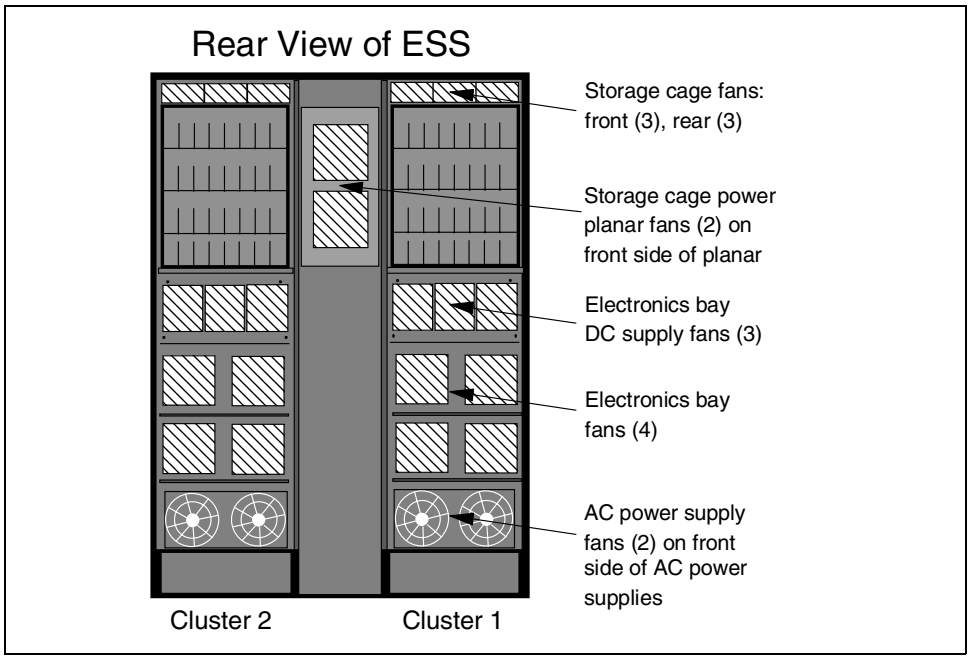

Figure 14. Cooling fan groups in the ESS base frame

In addition to these groups of fans, there is a fan in each of the three Electronics cage DC power supplies in each cluster.

The fans in the electronics cage DC supplies are not associated into a group. If any of these fans fails, the corresponding power supply is quiesced. The remaining two DC supplies in the cluster continue to function while the failing power supply is replaced.

4.1.7 Ethernet

The Ethernet adapters in the clusters are used by the clusters to communicate with each other, with a Web client to support ESS Specialist, and to send alerts using e-mail or SNMP. If the ESS participates in PPRC, and the Copy services function of ESS Specialist is used to manage PPRC, ethernet is used for communication between the Copy Services primary or secondary server and the participating ESS subsystems. An Ethernet network failure can isolate both clusters from the network or from each other. In either case, operation of both clusters continues. The connection between the Web client and either or both clusters may no longer be available. In this case, use of the ESS Specialist may not be possible.

If the network connection between the clusters and the outside network is not operational, the clusters may log errors if they attempt to use the network and find it unoperational. Use of the ESS Specialist from a remote machine will not be possible, but may be OK on a locally attached Web client. Management of PPRC may be disrupted.

In the event of a network failure that isolates the clusters from each other, some recovery actions must be performed by the IBM CE to return the subsystem to the normal state. If one cluster is isolated from the network (either by Ethernet

adapter or cable failure), it will not be able to sends an alert about the problem. At regular intervals, each cluster runs a system check procedure. When the opposite cluster runs this procedure, it detects that the first cluster is unreachable and logs the network failure. It sends notification using SNMP, e-mail or pager with the failure information.

If the problem is in the internal Ethernet adapter, a cluster failover is required to effect a repair action. After the failback process is complete, normal operation resumes.

4.2 Maintenance

As described in 4.1, “ESS internal recovery” on page 39, internal failures in the ESS are recovered automatically by the ESS without user intervention.

4.2.1 Maintenance strategy

The maintenance philosophy of the ESS is based on:

- **First time data error collection** - Data needed for analysis of the problem is collected at the time of failure and logged.
- **Non-recreate methodology** - Sufficient information is collected at the time of the failure to isolate the problem in most cases. The problem does not need to be recreated.

Log and trace data for a problem can be collected remotely by Product Engineering (PE) support for detailed analysis. The IBM CE is despatched by IBM support to repair the problem.

The CE uses a MoST (Mobile Service Terminal) service terminal connected to a cluster to access the maintenance procedures. The problem is opened for service, the required resource is quiesced (for example, a cluster could be made available for maintenance by failover). The CE is guided through the repair process by the maintenance procedures, and the repair validated by running automated diagnostic tests. If validation is successful, the resource is returned to service, and the problem is closed.

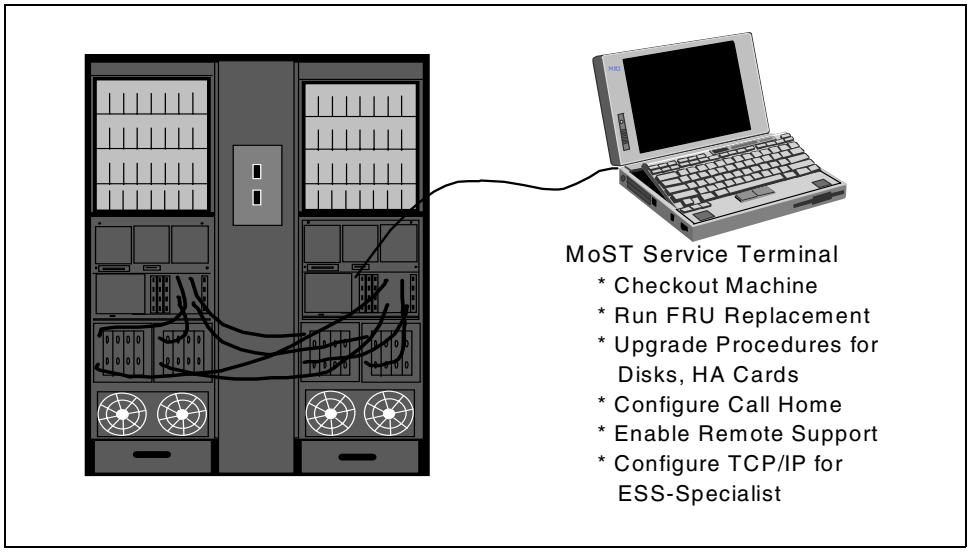

Figure 15. Service terminal functions

The MoST is an IBM Thinkpad running a terminal emulator (refer to Figure 15). It is used by the IBM CE to configure communication functions during ESS installation, to run procedures to upgrade storage or add host adapters, and to run maintenance procedures when a repair is needed. It is connected to the S2 serial port in either of the clusters.

4.2.2 Upgrades

The 2105 Models E10 and E20 are designed to allow non-disruptive upgrade of cache capacity and NVS capacity, an addition of device adapters, host ESCON and SCSI adapters, and disk drive bays.

The 2105 Models E10 and E20 come configured with the full complement of 6 GB of cache, 384 MB of NVS, and four DA pairs factory installed. Consequently, the only field upgrades are the addition of ESCON and SCSI host adapters, and the addition of disk storage capacity. These field upgrades can be performed non-disruptively.

An expansion frame can be nondisruptively attached to a base frame by a CE under guidance of the service terminal. The correct sequence must be followed when connecting SSA cables during addition of disk drive bays, or addition of an expansion frame, to ensure there is no disruption to operation of existing disk arrays. If only eight LSSs were defined at installation time, a disruptive IML may be required after addition of an expansion frame to activate new LSS/LCUs that may need to be configured to support the new storage.

4.2.3 Nondisruptive service actions

When maintenance is required, virtually all components in the ESS can be serviced concurrently with continued operation, provided that the ESS is correctly configured.

4.2.3.1 Microcode updates

Microcode in the storage clusters can be updated concurrently by using the failover procedure to make each cluster available in turn, then failing back after the new microcode has been activated. Microcode is loaded from a CD ROM drive located in each cluster. Configuration information can be copied between clusters on diskette. Additionally, microcode patches can be loaded into the ESS using the remote support facility.

4.2.3.2 Host adapters

ESCON and SCSI host adapters can be replaced *without* using cluster failover. Servicing of a host adapter card requires that the adapter bay containing up to four HA cards (eight ports) is quiesced. All the ports, ESCON or SCSI, in the bay will be unavailable during the repair.

Although a storage cluster and two host adapter bays are supplied by a common group of three DC power supplies in an N+1 configuration, DC can be independently switched off to the cluster, or either HA bay under control of the service terminal. For maintenance on a HA bay, DC power to the bay is switched off, and the bay is withdrawn to allow replacement of the host adapter or planar board.

Nondisruptive service on a host adapter bay depends on correct configuration of the ESCON and SCSI paths from each host to the ESS. There must be at least one alternative path available through the other HA bays for each ESCON or SCSI port used in the bay being repaired.

4.2.3.3 CPI cables

The common parts interconnect (CPI) cables connect the host adapter bays to the IO adapter cards in the storage clusters. Each HA bay has two CPI cables-one to each cluster. The CPI cables can be replaced nondisruptively, provided that the HA bay is quiesced.

4.2.3.4 Storage cluster components

Any component in a storage cluster can be replaced by using the failover/failback procedures to make the cluster available for service. Cluster components include the SMP processors, memory (cache), device adapters, NVS, IO adapter cards, the service processor card, and the planars on which these components are mounted. The associated CD-ROM, hard drive, and diskette drive are also components in the cluster.

While the cluster is offline, and powered off for service, all host adapters remain operational.

Cluster hard drive replacement

Each cluster is responsible for loading microcode into its SMP processors during initialization, monitoring and logging errors, and notification using the methods configured (for example, e-mail, SNMP, pager, call home). The clusters also support recovery from internal failures, and provide the maintenance facility used to make repairs to the ESS. These functions are supported by a dedicated internal hard disk drive. The data residing on the drive falls into three categories:

- ESS microcode, or licensed internal code (LIC)
- ESS customization data (for example, communications parameters)

- Configuration information (hosts, ports, and logical volumes)

A full copy of all of this data is stored in each cluster. In the case of internal disk failure in one cluster, the IBM CE replaces the failed disk, then reloads the LIC from the original CD-ROM. Configuration and customization data are saved on diskettes from the other cluster and restored to the replacement hard disk. The cluster is then returned to service using the failback process. The repair is concurrent with continued operation of the ESS.

4.2.3.5 Device adapters

Replacement of a failed DA requires that cluster failover has occurred to make the cluster available for repair. This should happen automatically if a DA fails. Replacement of the DA does not disrupt operation of the two associated SSA loops. After the DA card has been replaced, the cluster will be powered on, tested and restored to service using failback.

4.2.3.6 SSA cables

The SSA cables connect the disk drive bays to the DAs. An SSA cable can be replaced nondisruptively because the DAs reconfigure the loop to operate without it. During the repair, each DA has access to each drive on the loop via only one path around the loop instead of the normal two paths.

4.2.3.7 Hard disk drives

Any one of the SSA disk drives can be hot plugged by exchanging the failed drive for a new one. The correct maintenance procedure must be followed by the IBM CE using the mobile service terminal to change the drive, test the new drive and initiate a format operation to change the format in the new drive from the conventional 512 byte sector SCSI format to the 524 byte sector format used in the ESS. The new drive becomes a spare.

4.2.3.8 Power components

Power components fall into three main categories:

- Primary components supplying the whole rack, such as AC power supplies, line cords and rack batteries
- Secondary components supplying DC power to sections of the machine
- Power control and sensing components

In general, primary components are 2N, secondary components are N+1, power control is 2N.

Primary power supplies and rack batteries

The primary power supplies, the line cords, and the two rack batteries can be replaced while the 2105 is operating. Both clusters continue to operate normally during replacement of these components.

Electronics bay DC power supplies

These components are N+1, as described in 4.1.6.3, “DC power supplies” on page 46. Three DC supplies are provided for each storage cluster. Each of the three power supplies has an integrated cooling fan. If one of the DC supplies fails, the cluster and host adapter bays continue to operate normally, powered by the remaining two DC supplies. The failing power supply can be replaced nondisruptively without the need for cluster failover.

Disk drive cage DC power supplies
These components are N+1, and, like the electronics bay DC supplies, are hot pluggable.

Power control and sensing components
There are two Rack Power Control (RPC cards) in the ESS. These, the electronics cage sense cards, and the storage cage fan/power sense cards can be replaced with the rack power on.

4.2.3.9 Cooling fans
These components are N+1, as described in 4.1.6.4, “Cooling fans” on page 47. They can be replaced individually, except for the fans in the electronics bay DC power supplies, which require replacement of the power supply.

4.2.4 Disruptive service actions
Replacement of the planar board for the disk drive cage DC power supplies requires that the supplies be removed, and, therefore, cannot be completed nondisruptively.

Replacement of the UEPO operator panel card requires the ESS subsystem to be powered off.

The probability of failure in these components is extremely small.

4.3 Error notification
The ESS storage clusters continuously monitor the operation of the ESS hardware and microcode. If a failure is detected, or a situation occurs that requires notification to the customer or IBM, it can be reported to system administrators or to IBM in a number of ways:

- Call home
- SNMP
- E-mail
- Pager
- Service information message (SIM) - S/390 only
- Environmental recording and editing program (EREP) - S/390 only

Figure 16 on page 54 shows the methods available for communication with ESS internal functions. Error notification uses a serial port and the call home modem to notify IBM. The call home modem is also used to call a pager if this function is configured.

The IBM CE uses a second serial port on one of the storage clusters to connect the service terminal. This is used to configure the ESS during installation, and to perform maintenance procedures.

The Ethernet connection is used for SNMP and e-mail notification. It is also the means by which a local or remote Web client can be used to configure and monitor the ESS, using ESS Specialist and ESS Expert.

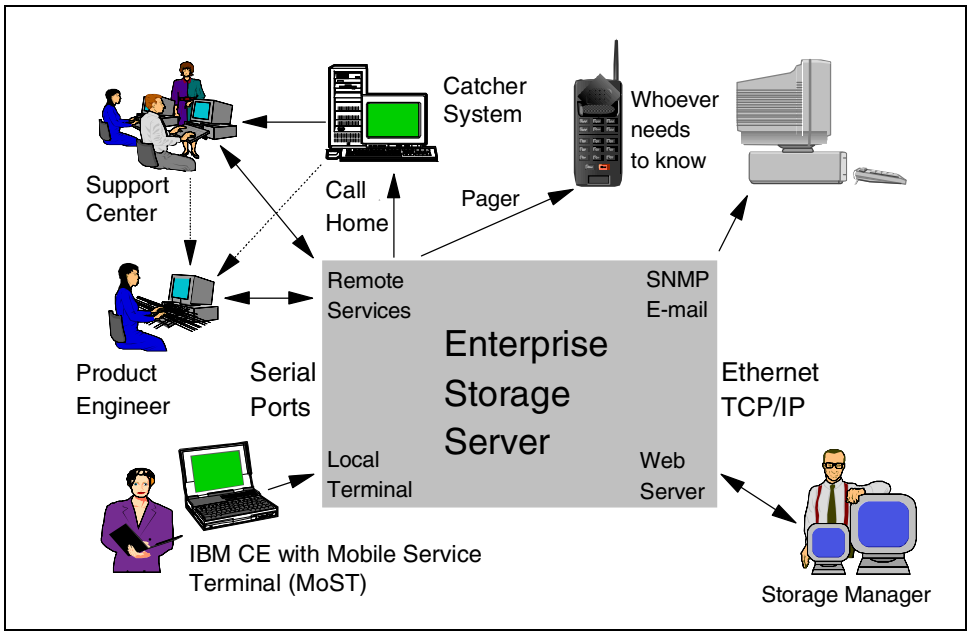

Figure 16. ESS communications

Any or all of these options can be used. Apart from SIM and EREP, they must be enabled through the service terminal used by the IBM CE during installation. For SNMP, e-mail, and pager, the necessary information such as e-mail addresses, is configured through the ESS Specialist interface.

4.3.1 Call home

The call home facility in the ESS is similar to that in the earlier VSS.

This feature of the ESS enables it to dial the IBM Support Center directly in the event that maintenance is required. The process is automatic; no customer intervention is needed. Call home must be enabled using the ESS Specialist. The call home function dials into a catcher PC in the IBM support center. The PC connects to the IBM RETAIN system to create a problem record. This record contains information on the failure and allows the support center to analyze the problem.The problem record is used by IBM to manage and resolves problems.

If required, the Support Center can dial into the ESS using remote support to collect additional information.

Call home setup is done by the CE during installation of the ESS.

4.3.1.1 Remote support

If the remote support function is enabled, a remote IBM specialist can initiate a call into the ESS and download log information, microcode traces, and other data for analysis to aid in the diagnosis of problems in the ESS. This is useful for remote diagnosis of hardware problems, but is especially important in the analysis of microcode errors. When a microcode fix is developed, it can be downloaded via the call home modem to the ESS. The microcode fix can then be non-disruptively applied by arrangement with the user.

Remote support makes use of the same modem and serial ports in the storage clusters that are used for call home.

Note: Remote access does not compromise the security of your data stored in the ESS.

To enable remote support, you must use the ESS Specialist before the ESS will accept incoming calls. Two levels of support can access the ESS through the cluster serial ports. The first level, Customer Engineering (CE) support, gives limited access to SMIT. The second level, Product Engineering (PE) support, needs root access to the operating system in the storage clusters to enable diagnosis of problems. This level of access must be specifically enabled by you, onsite, and it automatically disables after one week.

Neither support level gives access to your data. If the ESS accepts an incoming remote support call, you are notified by e-mail.

4.3.2 Simple network management protocol (SNMP)

The ESS generates SNMP traps and supports a read-only Management information base (MIB). The ESS generates both generic and product-specific SNMP traps. Product-specific traps provide information on problems that are detected by the ESS and require corrective action. The ESS sends information to the trap addresses set by the CE during the installation procedure.

The ESS supports the usual generic traps, such as cold start and warm start. Product-specific traps provide information on problems detected by the ESS that require action by the user or by IBM. Product-specific traps contain the identifier assigned by the ESS for the problem that caused the trap. The character string "Problemid=N" (where N is the problem identifier) is in the trap's description field. You can use this to find additional details on the problem using an MIB browser.

Using the MIB browser to view problem details
To use the MIB browser to view problem details:

1. Point the MIB browser to the problem that created the trap.
2. Look within the ibm2100 mib, ssProblem, ssProblemid. A display of this MIB variable is in the form "Index:Problemid".
3. Using the problem identifier (ID) from the trap, determine the index value with which it is associated.
4. Use the index value to determine which specific data is associated with a problem, and look at other MIB variables within ssProblem.

4.3.3 E-mail

The ESS uses standard TCP/IP e-mail to send error reports and notifications to the e-mail addresses you define. You can specify four levels of notification for each e-mail recipient, and the number of times that notification is sent for each problem. This parameter also defines the number of times the ESS will call home to IBM for each problem.

The ESS generates e-mail messages in two categories:

- Information
- Errors

Informational messages
Examples of informational messages are:

- A new level of Licensed Internal Code (LIC) has been installed.
- New hardware has been installed.
- The CE has run the customer-notification diagnostic test. This test verifies that e-mail messages are being received at the addresses set up by the CE or by the customer, using ESS Specialist.
- A remote support specialist has dialled into the ESS.

Error messages
The ESS sends error messages when it detects a situation that requires action by the customer or IBM. The error messages typically contain the following fields:

- Product manufacturer ID and date.
- Rack location.The Rack Location is entered by the installer during the initial installation of the product.
- Product machine type and model number (assigned by IBM).
- Product serial number (assigned by IBM).
- Customer voice phone number.
- LIC level of local storage server.
- LIC level of remote storage server.
- Report time/date stamp.
- Problem ID. The problem ID assigned to this problem by the storage facility. This problem ID can be used to access detailed problem information.
- Exception Symptom Code (ESC). A detailed error code used to define a problem. Used by the IBM CE to enter the maintenance procedures.
- SRN (Reference Number). A detailed code used by the IBM CE.
- Problem Status. See “Status” on page 57 for a list of the problem states.
- Description. A description of the problem.
- Additional message. Any additional information that is available.
- Failing cluster. The cluster on which the failure occurred. (1 or 2).
- Reporting cluster. The cluster which reported the failure. (1 or 2).
- Failing resource. This coded resource name is used during the repair process by the IBM CE.
- Failure occurred. Date and time when the failure first occurred.
- Last occurrence. Date and time of the last recorded occurrence.
- Failure count. The number of times this failure occurred.
- Presentation interval. The time between successive reports of this problem.
- Remaining presentations. The number of additional times this report will be sent.
- Isolation procedure. A pointer to a special procedure in the service guide.
- Failure actions. Actions the service provider should take.
- Probable cause. Information for the service provider.

- Failure cause. Information for the service provider.

The information above is also sent to IBM using the call home facility, if this has been configured.

Following are the fields that are most useful in identifying DDM failures:

- Description
- Failing cluster
- Reporting cluster
- Failing resource
- Last occurrence

4.3.3.1 Status

A problem can exist in various states. Each state represents a state in the reporting or repair of a problem. The problem states are defined as follows:

Pending The Initial problem state. The problem will be reported via one of the methods (SNMP trap or e-mail).

Received This state is used to indicate that the notification has been received by a host, either via e-mail or SNMP trap.

Open A repair process has begun. The required resources have been removed from use. Suspending a repair leaves a problem in the OPEN state.

Closed A repair process has been completed. All resources have been successfully returned to use.

Canceled A service representative has chosen to cancel this problem.

Expired Once a problem has been in the pending state for 30 days, it is changed to this state.

Archived Problems that have been closed, expired or canceled for more than 30 days are archived.

4.3.4 Pager

You can optionally configure the ESS to send problem information to a pager using the call home modem.

The call home, SNMP, e-mail and pager functions are configured using information you provide on the *Communication Resources Worksheet*. For information on completing the the work sheet, refer to 3.9, “Defining communication functions” on page 35.

4.3.5 S/390 notification

The two methods of error notification for S/390 systems are console messages, and by logging errors for display by EREP.

4.3.5.1 Service information message (SIM)

In the S/390 environment, the ESS uses one of the ESCON channel paths to notify a S/390 host of certain internal conditions. Sense data is logged by the host. Conditions in the ESS that need to be brought to the attention of system operators usually result in the presentation of a service information message

(SIM) on a system console. This message identifies the ESS and provides information about the error and the impact of repair. For OS/390, the SIM is presented in message IEA480E.

The ESS sends SIMs to System/390 host consoles for three types of SIMs:

DASD SIM This SIM tracks disk drive module (DDM) failures and problems

Media SIM This SIM tracks data check problems on the media.

Storage facility SIM This SIM tracks control unit, power control, and other hardware problems.

4.3.5.2 EREP

In S/390 systems, information about hardware errors and software errors and statistical information is written to a data set on a system disk. For OS/390, the data set is called LOGREC. The environmental recording and editing program (EREP) is used to select desired types of entry in the data set, edit and summarize them. In particular, the event report lists events in chronological order, and can be useful when analyzing problems such as missing interrupts.

4.3.5.3 Devserv (DS) command

The Devserv command can be used to obtain information about a disk subsystem attached to an S/390 host using the OS/390 operating system.

Devserv Paths

In OS/390, the Devserv Paths command displays useful information on the state of a device and the paths used to access it from the host. See Figure 17 on page 59 for an example of the output (IEE459I message) for the Devserv Paths command. In this example, the Devserv Paths command was for eight devices in an ESS, beginning at address 6800 (DS P, 6800,8).

The information returned includes the subsystem ID (SSID), the VOLSER for each device, and indication of the result for each channel path of an attempt to communicate with the device over that path. The plus (+) sign for each path in the example indicates successful communication.

The dynamic legend at the bottom of Figure 17 on page 59 defines the symbols that appear in the output.

```
DS P,6800,8
IEE459I 09.41.23 DEVSERV PATHS 459
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
6800,33903 ,O,000,SL6800,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   00   00        2105
6801,33903 ,O,000,SL6801,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   01   01        2105
6802,33903 ,O,000,SL6802,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   02   02        2105
6803,33903 ,O,000,SL6803,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   03   03        2105
6804,33903 ,O,000,SL6804,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   04   04        2105
6805,33903 ,A,001,SL6805,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   05   05        2105
6806,33903 ,O,000,SL6806,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   06   06        2105
6807,33903 ,A,002,SL6807,80=+ 84=+ 81=+ 85=+ C0=+ C4=+ C1=+ C5=+
     933201  22C0  Y  YY.  YY.    N   SIMPLEX   07   07        2105
************************ SYMBOL DEFINITIONS ************************
A = ALLOCATED                     O = ONLINE
```

Figure 17. OS/390 Devserv Paths output

An important attribute of the Devserv Paths command is that it causes the state of the subsystem to be obtained by performing I/O operations at the time of command execution. Apart from the VOLSER field (which is obtained from the UCB), the information is read directly from the disk subsystem. For example, if a channel path has a plus (+) sign, you know that the path was working at the time you entered the DS P command.

The Devserv Paths command can be performed to both online and offline devices. It will not display the VOLSER for an offline device (the VOLSER is not in the UCB).

Devserv Query DASD

The Devserv Query DASD command (DS QD) is used to obtain information about a device from OS/390. For example, Figure 18 shows the result of DS QD,6800,8 to an ESS subsystem.

```
DS QD,6800,8
IEE459I 09.40.05 DEVSERV QDASD 360
UNIT VOLSER SCUTYPE DEVTYPE   CYL  SSID SCU-SERIAL DEV-SERIAL EF-CHK
6800 SL6800 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6801 SL6801 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6802 SL6802 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6803 SL6803 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6804 SL6804 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6805 SL6805 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6806 SL6806 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
6807 SL6807 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
```

Figure 18. OS/390 Devserv Query DASD output

The DS QD,UCB displays the device UCB as shown in the example in Figure 19.

```
DS QD,6800,UCB
IEE459I 09.44.44 DEVSERV QDASD 183
UNIT VOLSER SCUTYPE DEVTYPE   CYL  SSID SCU-SERIAL DEV-SERIAL EF-CHK
6800 SL6800 2105E20 9332B01  3339  22C0 0113-FCA22 0113-FCA22 **OK**
  UCB AT V00ECA520
0188FFAC68000800 0000000000E4C3C2 3030200F00ECA4F8 19E60100E2D3F6F8
F0F0040400A0000B 00ECA32000000000 0000000000000000
  UCB PREFIX AT V0270EDE8
000C804000DDFD80 00D4070000010D56 289C1012FF0080FF 80848185C0C4C1C5
0108000000000001
  UCB COMMON EXTENSION AT V00ECA4F8
0000094020AA0000 0270EDE8000000A0 0000000000FCCAA4 00ECA4C000000000
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

Figure 19. OS/390 Devserv Query DASD UCB output

Examples of other areas that can be displayed using the Devserv Query Dasd command are DCE, DCPT, and SSSCB.

An enhancement to the OS/390 Display Matrix command (D M) command causes information on PAV aliases to be displayed for an ESS. Refer to Figure 20 for example output from a Display Matrix command for a channel (D M=CHP(80)), showing the devices configured on the channel, and the aliases.

```
D M=CHP(80)
IEE174I 10.05.24 DISPLAY M 779
CHPID 80:  TYPE=05, DESC=ESCON SWITCHED POINT TO POINT
DEVICE STATUS FOR CHANNEL PATH 80
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
680  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
681  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
682  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
683  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
684 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
685 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
686 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
687 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
688  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
689  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
68A  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
68B  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
68C UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
68D UL UL UL UL UL UL AL AL AL AL AL AL AL AL AL AL
68E UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
68F UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
BX DEVICE IS BOXED           SN SUBCHANNEL NOT AVAILABLE
DN DEVICE NOT AVAILABLE      PE SUBCHANNEL IN PERMANENT ERROR
```

Figure 20. OS/390 Display Matrix Example Output

4.4 Configuration for availability

Virtually all forms of internal failure in the ESS are recovered automatically by the ESS without intervention by the user. This is achieved using mechanisms such as RAID 5, cluster failover, and multiple redundant paths between the storage clusters, disk arrays and host adapters.

Configuring for maximum availability of the ESS and access to data requires attention to these considerations:

- Maintaining connectivity between each host and the ESS in the event of a failure in an ESCON or SCSI path and during repair of the failing component
- Using RAID 5 or host software mirroring to ensure continued access to data in the event of a disk drive failure in the ESS

- Ensuring the reliability of the power supply and cooling to the ESS racks

4.4.1 Maintaining connectivity

ESCON path failures include problems with host ESCON channels, problems in ESCON cables, patch panels, extenders or switches, and faults in ESS ESCON ports. SCSI path failures can be caused by problems in host SCSI initiators, other devices on a SCSI interface, or in physical cabling. Problems may be caused by incorrect configuration such as excessive cable length or incorrect termination.

To maintain connectivity in the event of a path failure, at least one alternate physical path must be provided from the host to the logical device in the ESS. An FB logical device in the ESS can be logically assigned to two or more SCSI ports to enable sharing between hosts, or to provide multiple physical paths to a single host.

4.4.1.1 Data path optimizer

For UNIX, Windows NT, or AS/400 environments, multiple SCSI paths to a device must be supported by software. IBM Data Path Optimizer (DPO) is a software product that provides availability and I/O load balancing for the AIX and Windows NT environment. Other platforms may be supported by other non-IBM products.

DPO provides availability via automatic I/O path failover. If a failure occurs in the data path between the host and the ESS, Data Path Optimizer automatically switches the I/O to another path. DPO will also move the failed path back online after a repair is made.

DPO improves performance by sharing I/O operations to a common disk over multiple active paths to distribute and balance the I/O workload. Refer to Figure 21 on page 63.

Shared LUNS not supported by DPO

DPO cannot be run in an environment where more than one host is attached to the same LUN (a multi-host environment). This restriction includes clustered hosts such as RS/6000 servers running HACMP and NT High Availability Clusters.

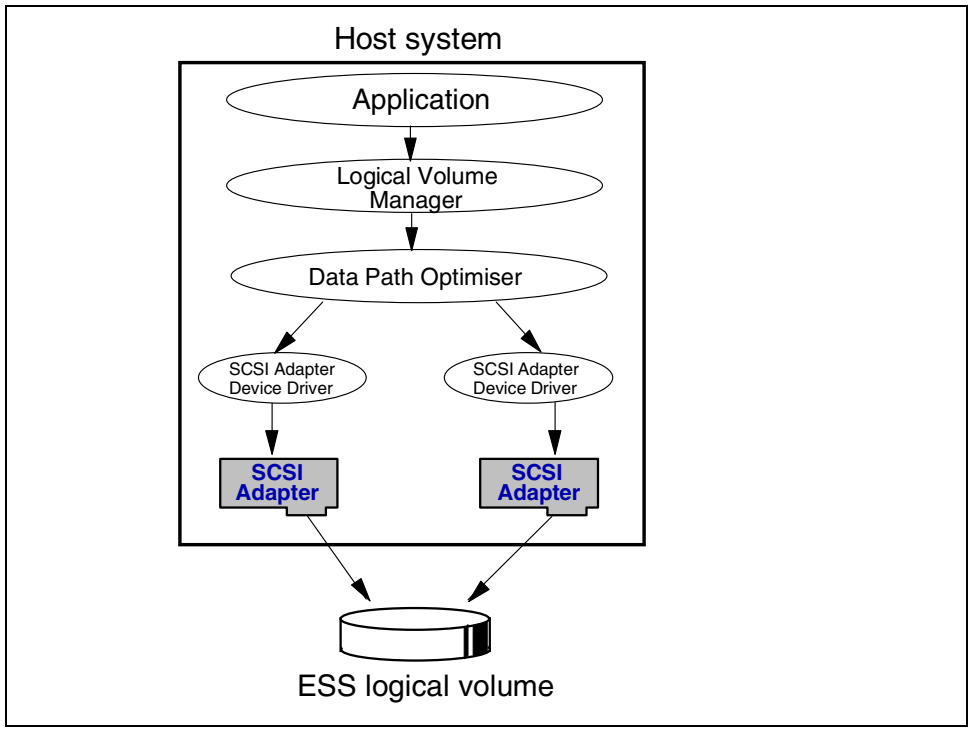

Figure 21. Function of data path optimizer

Supported Systems

Data path optimizer Release 1.0 supports Windows NT and UNIX operating systems:

- AIX 4.2.1 with PTF IX62304 and AIX 4.3.2
- Microsoft Windows NT 4.0 Server Edition with Service Pack 3 or 4

Windows NT

DPO for Windows NT is installed using the install shield. It is self configuring. The Data Path Optimizer operates as a filter mechanism. Other paths to the drive appear offline, but they are used in rotation.

AIX

SMIT is used to install and configure Data Path Optimizer.

Conversion scripts replace hdisk devices by vpath devices for volume groups. Affected file systems must be unmounted. hdisk devices are still online, but to use the DPO functions, the vpath devices must be used.

Datapath Commands

For AIX, a command line path recovery command is provided. It allows you to query devices and adapters. You can use the datapath command to vary paths online or offline.

```
datapath query adapter/device [n]
datapath set adapter <n> online|offline
datapath set device <n> path <m> online|offline
```

Trace Function
In an environment with more than one path to a drive, errors could occur on a single path, while others are working error free. DPO supports driver traces to assist in resolution of single path and intermittent problems that otherwise can be difficult to isolate.

See:

`http:www.storage.ibm.com`

for further details on Data Path Optimizer.

Refer also to *IBM Data Path Optimizer,* SC26-7291.

We recommend that DPO, or a program that provides equivalent function in the host operating system (for example, Veritas Volume Manager for the SUN Solaris OS) be used to enable the configuration of two or more physical paths from each SCSI host to logical devices in the ESS. The OS/400 operating system has the needed multiple path support included.

Multiple SCSI paths from each host should be distributed to different adapter bays when possible. This minimizes the impact of failure in host SCSI adapters, cables and ESS SCSI adapters, and allows for nondisruptive repair of ESS SCSI adapters.

4.4.1.2 ESCON
In the S/390 environment, normal practice is to provide multiple paths from each host to a disk subsystem. Typically, four paths are installed. The channels in each host that can access each Logical Control Unit (LCU) in the ESS are defined in the HCD (or IOCDS) for that host. Dynamic Path Selection (DPS) allows the channel subsystem to select any available (non-busy) path to initiate an operation to the disk subsystem. Dynamic Path Reconnect (DPR) allows the ESS to select any available path to a host to reconnect and resume a disconnected operation, for example to transfer data after disconnection due to a cache miss.

These functions are part of the S/390 architecture and are managed by the channel subsystem in the host and the ESS.

A physical ESCON path is established when the ESS port sees light on the ESCON fiber (for example, a cable is plugged in to an ESS host adapter, a processor or the ESS is powered on, or a path is configured online by OS/390). At this time, logical paths are established through the ESCON port between the host and some or all of the LCUs in the ESS, controlled by the HCD definition for that host. This happens for each physical path between a host CEC and the ESS. There may be multiple system images in a CEC. Logical paths are established for each system image. The ESS then knows which ESCON paths can be used to communicate between each LCU and each host.

At the host operating system level (for example, OS/390), the paths to be available for use for each logical device are defined to the ESS at the time the device is brought online.

ESCON directors
Because a large number of hosts may be connected to the ESS, each using multiple paths, the maximum 32 ESCON adapter ports that can be installed in the

ESS may not be sufficient to accommodate all the connections. Where SCSI adapters are installed, the number of ESCON adapter ports will be fewer than 32. The solution to this problem is the use of IBM 9032 ESCON Directors to switch logical connections from multiple host channels to a single physical port connected to the disk subsystem.

A logic or power failure in an ESCON Director can interrupt communication between hosts and the ESS. We recommend that more than one ESCON Director be provided to ensure continued availability. For example, four of the eight ESCON channels in a path group could be configured to go through each of two ESCON Directors as shown in Figure 22. The complete failure of either ESCON Director leaves half of the ESCON paths still operating. In a large installation, there may be two eight-path ESCON groups connected to an ESS. Four directors might be installed, each carrying two paths in an eight path group.

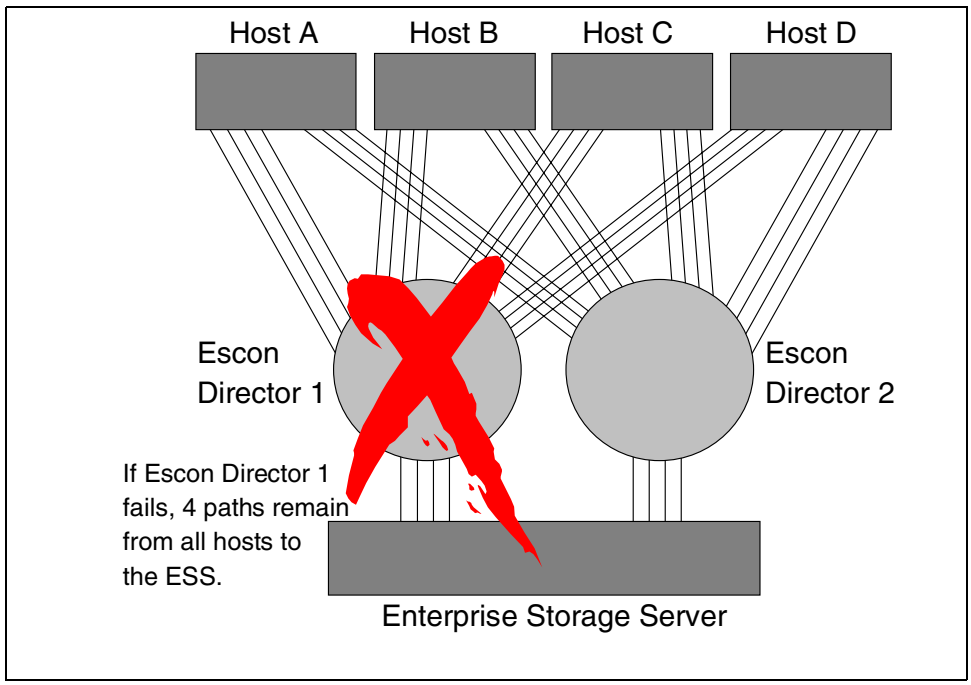

Figure 22. Multiple ESCON switch configuration

For availability, and to provide sufficient data bandwidth to provide maximum performance, we recommend that each MVS host has eight ESCON paths to all LCUs it will access.

Assuming that the ESS has ESCON adapters in each HA bay, you should configure two paths of each eight path group to the two ports in one ESCON adapter per bay. Refer to Figure 23 on page 66. The failure of one ESCON adapter card removes two of the eight paths. This is equivalent to losing one path of four, and the impact should not be significant unless channel utilization is very high.

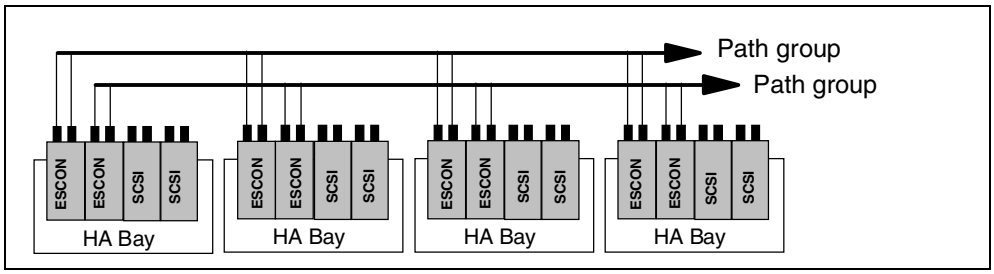

Figure 23. Physical configuration of ESCON paths

This configuration is recommended for best performance. Refer to 2.4.0.2, “Host adapters” on page 20 for guidelines on configuring ESCON channels for performance.

4.4.2 Access to data

If a disk drive in the ESS fails, data continues to be available if the drive is a member of a RAID 5 array. A spare drive is reconstructed into the array and performance returns to normal after the reconstruction is complete.

If a drive configured as a JBOD fails, data on that drive is not available. There may be more than one logical disk on the drive. None of these logical disks will be available. Where JBODs are in use, continuous availability can be provided by host software mirroring. For example, the AIX Logical Volume Manager (LVM) can be configured to mirror selected volumes. The combination of JBOD and software mirroring provides excellent random write performance with continuous data availability, at the expense of a doubling in required storage capacity per volume and increased usage of resources such as SCSI paths.

IBM recommends that RAID 5 be used in preference to JBOD, except where performance on random reads is critical.

4.4.2.1 Remote copy

Techniques such as RAID 5 and software mirroring can provide data availability in the event of a single disk drive failure in the ESS. However, if the ESS loses power, or receives physical damage (for example, fire or water damage), that prevents its continued operation, data availability depends on backups of the data having been made and safely stored, perhaps offsite. An alternate site may be needed to continue processing.

Peer-to-Peer Remote Copy (PPRC)

Peer-to-Peer Remote Copy (PPRC) provides a means to maintain a synchronous copy of selected volumes on another ESS located in the same site, or elsewhere up to 103 kilometers away. PPRC can be used for both S/390 volumes and FB volumes. PPRC for FB and CKD volumes can be managed using the Copy Services function of the StorWatch ESS Specialist. Additionally, for S/390 volumes only, PPRC can be controlled from a S/390 host using commands similar to those used for PPRC on IBM 3990 and IBM Ramac Virtual Array (RVA). A feature code must be ordered to enable PPRC on the 2105.

Extended Remote Copy (XRC)
Extended Remote Copy (XRC) can be used in the OS/390 environment to create and maintain asynchronous copies of ESS CKD volumes. The ESS can have either primary or secondary volumes of XRC pairs. The primary and secondary volumes in an XRC pair can reside in different types of disk subsystem. For example, a primary device might be in an ESS while the secondary is in a 3990-6/RAMAC subsystem up to several thousand kilometers away.

If the ESS is to act as a primary control unit, the appropriate feature code must be installed. No feature code is needed for an ESS used only for secondary volumes of XRC pairs.

Chapter 6, "Enterprise Storage Server Copy Services Overview" in *IBM Enterprise Storage Server,* SG24-5465 contains detailed information on PPRC and XRC for the ESS.

4.4.3 Power and cooling

The ESS can continue to function only if power is available on at least one of the two line cords on each rack. The internal batteries provide power for a sufficient time to override power disturbances, and to destage write data and sequence down the ESS in the event of total power loss. The batteries cannot support continued operation without AC power.

The two line cords should be connected to independent AC distribution systems to minimize the possibility of losing power to both line cords simultaneously.

Although the ESS may continue to operate at temperatures outside the extreme operating range of 16 to 32 degrees Celcius (60 to 90 degrees Fahrenheit), continued operation at these temperatures may affect the future reliability of internal components, especially disk drives. If a failure in air conditioning equipment results in an ambient temperature outside this range, the ESS should be powered off until recommended operating conditions are restored.

Chapter 5. Enterprise Storage Server Specialist (ESS Specialist)

In this chapter we describe the ESS Specialist, the Web interface for the IBM ESS. It is the configuration and administration interface of the ESS. It is the most important interface between the hardware, the software, and the storage administrators.

5.0.1 ESS Specialist

The ESS includes the ESS Specialist, a network enabled management tool that allows the storage administrator to monitor and manage storage from the desktop or remote locations. Using a secure Internet connection (LAN with a Web browser), such as Netscape Navigator, Microsoft Internet Explorer, or Sun Hot Java (the browser must support Java 1.1.6 or higher), the storage administrator can coordinate the consolidation effort and easily integrate storage capacity into the ESS.

Note: The ESS Specialist Help text provides some detailed technical scenarios to help you perform some of the most common tasks, as well as some hints and tips. It might be worthwhile printing out the help text to help you plan for the installation and implementation. To access the Help text, click on the question mark symbol from any screen.

The ESS Specialist provides the customer with the ability to:

- Monitor error logs: If a problem occurs, a description of the problem including the failed component, the problem severity, and who is to be automatically notified, is described.
- View the status of the machine: Logical schematic of the Shark environment; Host attached ports, controller and cache stores, SSA adapters, SSA storage devices, and host icons.
- View and update the configuration: View via a color coded scheme, the amount of space allocated and assigned to one or more hosts, space allocated and not yet assigned, and space not allocated to logical volumes
- Add host systems or delete host systems.
- Configure host ports.
- Add volumes, remove volumes, and reassign volumes between different servers.

  Volumes can be reassigned between hosts as follows:

  - Removing volumes (or unassigning volumes from hosts)

    Volumes can be removed by removing all attached host connections to the logical volume. However, the space can not be reused, except by reconnecting a host to the same logical volume originally defined.

  - Adding volumes

    Volumes can be added from subsystem capacity that has never been defined or after an array has been reinitialized

  - Reclaiming previously defined logical volumes
- View communication resource settings, such as TCP/IP configuration and users

- View cluster LIC levels

  You can view the active level, next level yet to be activated, and the previous (older) level.

- Select an authorization level for each user:
  - Viewer

    A viewer can view the current configuration and status information.

  - Configurator

    A configurator can view the current configuration and status information and can make changes to the configuration.

  - Administrator

    An administrator can define new user IDs, delete old IDs, and assign, change, and revoke passwords and levels of authorization.

- New functions added to the ESS Specialist from VSS include:

  The ability to change volume size without involving the IBM CE. That is, you will have the ability to reformat the array.

  Web support for copy services (such as a list of volume pairs for the PPRC function)
  - Copy services is an addition to the Specialist Web function introduced with the VSS. Its purpose is to provide MVS and OPEN customers with an easy to use Web-based interface to manage and use the Peer-to-Peer Remote Copy and FlashCopy functionality.
  - If you are running from an OS/390 host today, you must collect detailed information about your storage subsystem topology before you can use FlashCopy or Peer-to-Peer Copy. You must collect information about physical control units, logical control units, paths, switches and volumes over which you wish to establish FlashCopy and Peer-to-Peer Copy relationships. The principle aim of RAS Copy Services is to eliminate the need for your to collect this information and to make it unnecessary for you to have to deal with that information unless you choose to.
  - Copy services works by providing a user interface through the ESS Specialist which presents the topology of the storage subsystem network. You are able to see all items of interest at any level of detail you wish. You may view all control units known in the storage subsystem network, or you may focus down to the volumes of one single logical control unit. By changing focus, you can view the storage subsystem as it is known from a specific host, a specific switch, or a specific control unit. By using the mouse, you are able to click on objects representing volumes, control units, logical control units and paths to build tasks that get scheduled and executed.

- Performance enhancements for moving graphical data across the network.

5.1 StorWatch family overview

The ESS Specialist is part of the StorWatch family of products. This section has been included to introduce the StorWatch concept.

5.1.1 The StorWatch family of products

The StorWatch family is designed to work with network-enabled hardware storage systems and storage management components that are compatible with the IBM Seascape Storage Enterprise Architecture. Both IBM and non-IBM systems will be monitored and managed by StorWatch products.

Currently, the StorWatch family consists of several products, the first two, ESS Specialist and ESS Expert will be the most commonly used with the ESS.

- StorWatch Enterprise Storage Server Specialist

  A product which runs on a Web server that is running in the processor complex of the IBM Enterprise Storage server. The ESS Specialist is used to configure the ESS internal logical configuration and monitor the status of the subsystem and report problem details.

  The Copy services component allows you to manage Peer-to-Peer Remote Copy and FlashCopy on the ESS subsystem.

- StorWatch Enterprise Storage Server Expert

  A product used to report on the physical performance and capacity for the ESS. It can functionally replace the StorWatch Reporter if your only network capable storage devices are ESS.

- StorWatch Reporter

  A product that can collect capacity information from hosts on a network if they have a StorWatch Agent installed to communicate with the StorWatch Reporter Manager.

- StorWatch Versatile Storage Specialist

  A product which is built into the Versatile Storage Server subsystem, the equivalent of the ESS Specialist.

- StorWatch Serial Storage Expert

  A product to Plan or View configurations of 7133 SSA disks and other SSA disk adapters and drives.

- StorWatch DFSMS/ HSM Monitor and Tuner Version 1 Release 1

  The DFSMS Optimizer uses historical and real-time data to provide an overall picture of data usage on each system in a sysplex environment.

- StorWatch Fibre-Channel RAID Specialist

  A product which is used to configure the Fibre-Channel RAID Storage Server.

5.1.1.1 StorWatch and the Enterprise Storage server

In this document we will only discuss the Enterprise Storage Server Specialist (ESS Specialist) and the Enterprise Storage Server Expert. You may also obtain information on using these products in the *IBM Enterprise Storage Server Web Interface User's Guide,* SC26-7346.

5.2 ESS Specialist components

The ESS Specialist is the major interface between the storage administrator and the ESS itself. It is used to configure the ESS logical storage environment and

provide administrative and problem status information. This section introduces the main panels of the ESS Specialist. For detailed information on how to define SCSI devices refer to Chapter 15, “ESS configuration for fixed block storage” on page 211. For S/390 devices, see Chapter 9, “ESS configuration for S/390 storage” on page 153.

5.2.0.1 ESSNet

To aid the customer in getting the ESS set up and configured, the ESS will optionally come with a private network already set up and a PC system to access the network and the ESS Specialist. The ESS Specialist will always be available to the customer or IBM support for configuration or status. In addition, the PC can be set up to be a proxy server to the ESS Specialist from anywhere in the customer environment.

5.2.1 Prerequisites

The ESS Specialist uses a set of Java applets. Therefore, you need a Java-capable browser that supports Java 1.1.2. We used Netscape Communicator 4.5 installed on Windows NT 4.0. We suggest that you use the same or a newer version, which you can download from the Netscape Web site.

IBM has not tested all browsers. It is possible, however that you could use any browser with frames support and full Java 1.1 support.

The minimum recommended browser hardware is a 166 MHz processor and 32 MB memory, however, greatly improved performance can be achieved with larger configurations (233 MHz and 96 MB RAM).

The ESS Specialist is not enabled by default. At installation time the IBM CE will assign the customer-specified IP address and hostname alias to each cluster controller within the ESS. Once complete, other setup functions can be performed by using the ESS Specialist.

5.2.2 Components

The ESS Specialist is a software package residing inside the ESS and includes these components:

- The Storage function operation interface (SFOI)
- The Information server
- The Web server
- Client code

The SFOI, the Information server, and the Web server belong to the ESS itself and run inside the cluster controllers. The client software is a set of Java applets running inside a Web browser on the client workstation.

The SFOI is an internal application programming interface (API) to the ESS server functions required by the ESS Specialist. The ESS Specialist uses this API to perform queries of the current configuration and service setup for the facility. All functions, including the ESS Specialist, SFOI, and ESS Information server, runs on both cluster controllers of the ESS. To enable the ESS Specialist to present a consolidated view of the storage facility, the SFOI maintains a cross-cluster socket connection. This ensures that the configuration information

from both clusters is available to you even though a network connection to only one of the two cluster controllers is established.

The Information server is essentially a communication link that communicates on one side with the client applets through an Ethernet TCP/IP socket connection and on the other side with the SFOI. It provides the most up-to-date configuration information to the client applets, so that it can be displayed to the end user, or to allow the end user to change the configuration.

The ESS Specialist Web server is set up for secure connection through the Secure Sockets Layer (SSL). The data sent between the Web server and the Web browser on the client side is encrypted through the use of public and private key pairs.

The client code of the ESS Specialist consists of a set of hypertext markup language (HTML) pages and Java applets, which are executed from inside a Java-enabled Web browser running on the end-user's network-connected workstation. The user has the option of connecting to either of the two cluster controllers, by specifying a universal resource locator (URL) consistent with the controller's IP address or hostname. See 5.6.4, "Communications panel" on page 82.

The ESS Specialist is frame enabled. When running inside a Web browser, it presents two main frames on the main panel: a navigation frame, on the left with hyperlink buttons for accessing ESS Specialist functions, and the working frame on the right (see Figure 24 on page 74). At the bottom of the panel is the message area where messages are displayed while applets are processing or the client browser is waiting to receive data from the server.

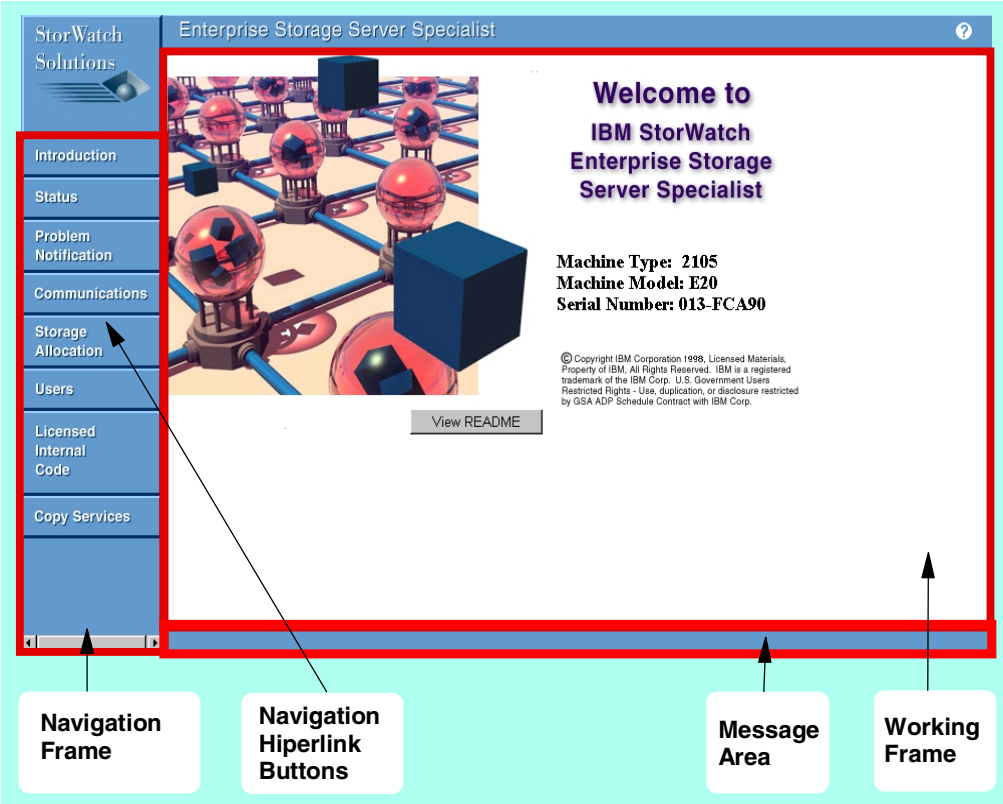

Figure 24. ESS specialist main panel

5.3 Help

The help system includes content help, task help, specific scenario assistance and a glossary. You can press the *help icon* at any time to bring up a new browser window that will display the help text. The help Icon (see Figure 25) is located on the top right-hand side of all panels. A sample of the text-based help is shown in Figure 26 on page 75.

Figure 25. ESS Specialist: Help icon

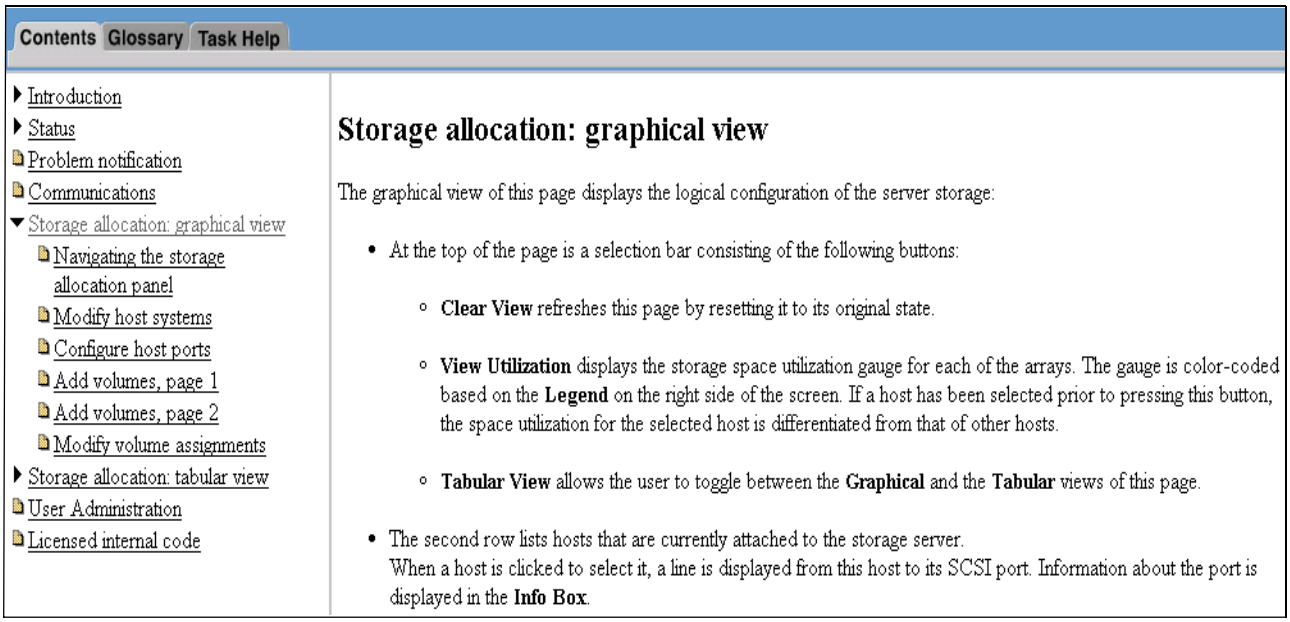

Figure 26. ESS Specialist: Sample help text

5.4 Security

Accessing the ESS through the ESS Specialist requires a valid username and password. As an option, the ESS Specialist can be configured to require that a user originate from a specific IP address or a range of IP addresses, see Figure 41 on page 88. Users can be restricted to one of three different authorization levels: viewer, configurator, and administrator.

A viewer can view the current configuration and status information, while a configurator can view the current configuration and status information, and can make changes to the configuration. An administrator can define new user IDs, delete old IDs, and assign, change, and revoke passwords and levels of authorization, as well as configure the ESS.

All data sent between the ESS and the ESS Specialist interface in your Web browser is encrypted to avoid unauthorized modification of configuration commands during transit. As a result, the Web browser will warn the user that an encrypted site is being accessed by displaying a sequence of certificate windows.

These windows tell your browser that the Web site represented by the ESS URL should be treated as a trusted site, and that encrypted communications should be allowed between your browser and the site. Figure 27 on page 76 shows one of the several windows displayed while connecting.

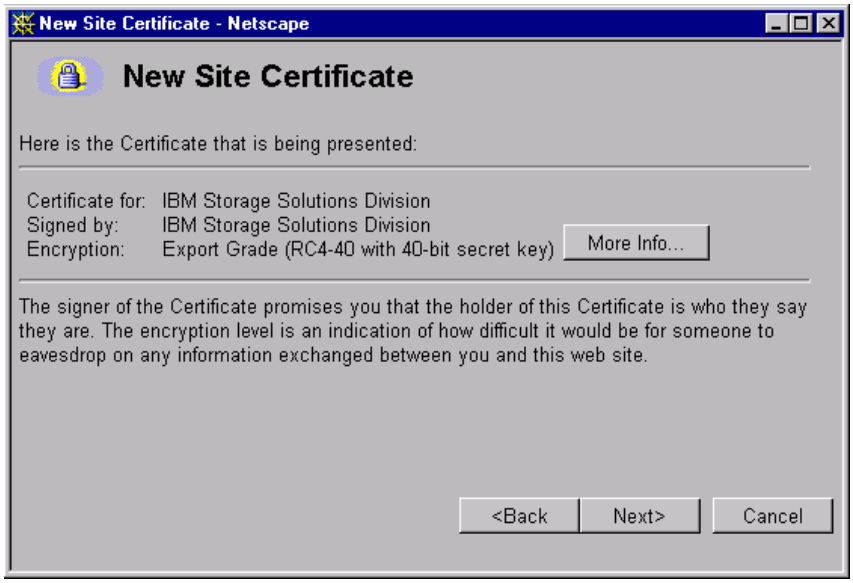

Figure 27. ESS Specialist: Certificate window

5.5 Connecting to your ESS Specialist

To connect to the ESS Specialist through the browser, enter the URL of one of the two clusters of your ESS machine (see Figure 28). You do not need to be connected to a specific cluster, even if you want to work with a cluster other than the one to which you connected.

For ease of identification, you could add a suffix such as c0 or c1 to the selected hostname for your ESS to represent the different clusters. For example, ESSa90c0 for cluster 0 as shown in Figure 28. Then bookmark it for ease of use.

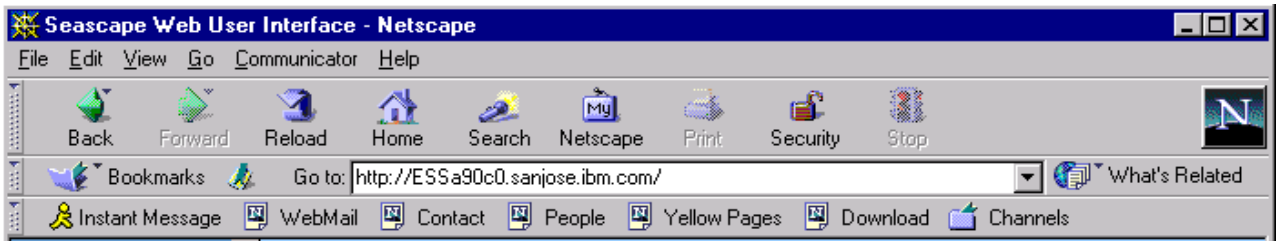

Figure 28. Entering the URL of your ESS machine

5.6 Panels

To access the ESS Specialist panels simply click on a button in the navigation frame. Figure 29 presents a logical view of all of the ESS Specialist panels, the connections among them, and the main functions.

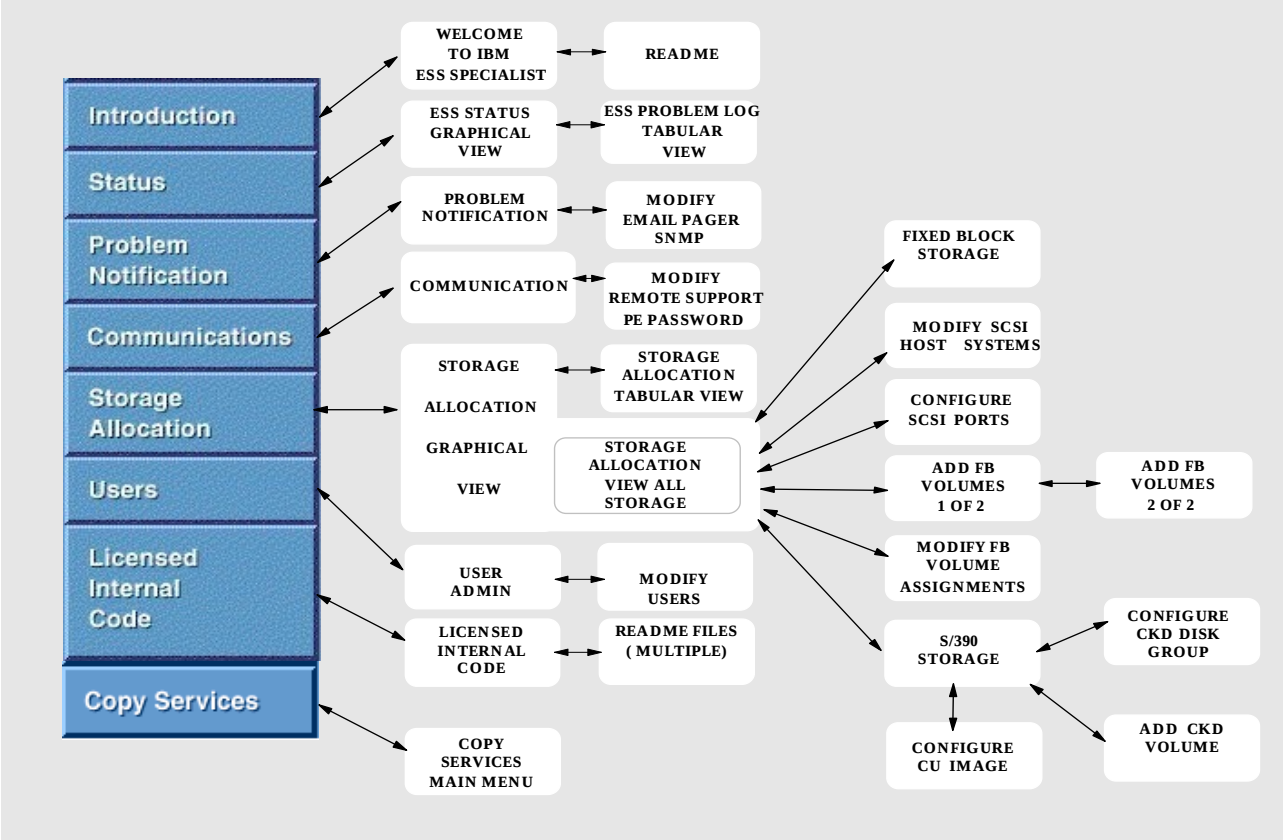

Figure 29. Logical view of the ESS Specialist panels

5.6.1 Introduction Panel

The introduction panel, also called the *information panel*, is the default panel. It shows the machine type, model, and serial number set at manufacturing time. Figure 30 on page 78 shows the introduction, or information panel.

Figure 30. ESS Specialist: Introduction panel

The title bar contains a question mark, which you can use to call the help panel. On the left side of the panel, the navigation frame contains these selectable buttons:

- Introduction
- Status
- Problem Notification
- Communications
- Storage Allocation
- Users
- Licensed internal code
- Copy Services

The Introduction panel comes up by default when you connect to the ESS Specialist.

The Introduction panel is the only ESS Specialist panel that does not require user authentication. When you enter the URL of your ESS machine, the Introduction panel appears, and you are not prompted to authenticate with the server. When you click on any other hyperlink button in the navigation frame, the Username and Password Required window appears (see Figure 31 on page 79), and you must authenticate before you can work with the panel you requested.

Figure 31. ESS Specialist: Username and Password required panel

Note: When started for the first time on a newly installed ESS, the ESS Specialist has one administrator username predefined. The user ID is *storwatch* and the password is *specialist*, all case sensitive. Logon using this username, and immediately create one or more additional administrator IDs with secure passwords. The storwatch ID will be automatically deleted by the ESS Specialist after the first new administration ID is added to the user list. Once the storwatch user ID has been deleted, it can not be re-defined as a valid ID. Attempts to do so will just be deleted by the software if another administration ID is present.

5.6.2 Status panel

When you select the Status button on the ESS Specialist navigation frame, the Status -- Graphical View panel appears as shown in Figure 32 on page 80. In this panel you can see the logical components of the ESS machine: host attachment ports and bays, clusters and storage arrays.The solid lines represent the physical data paths between the main components.

This panel should be reviewed for problem records that require service attention.

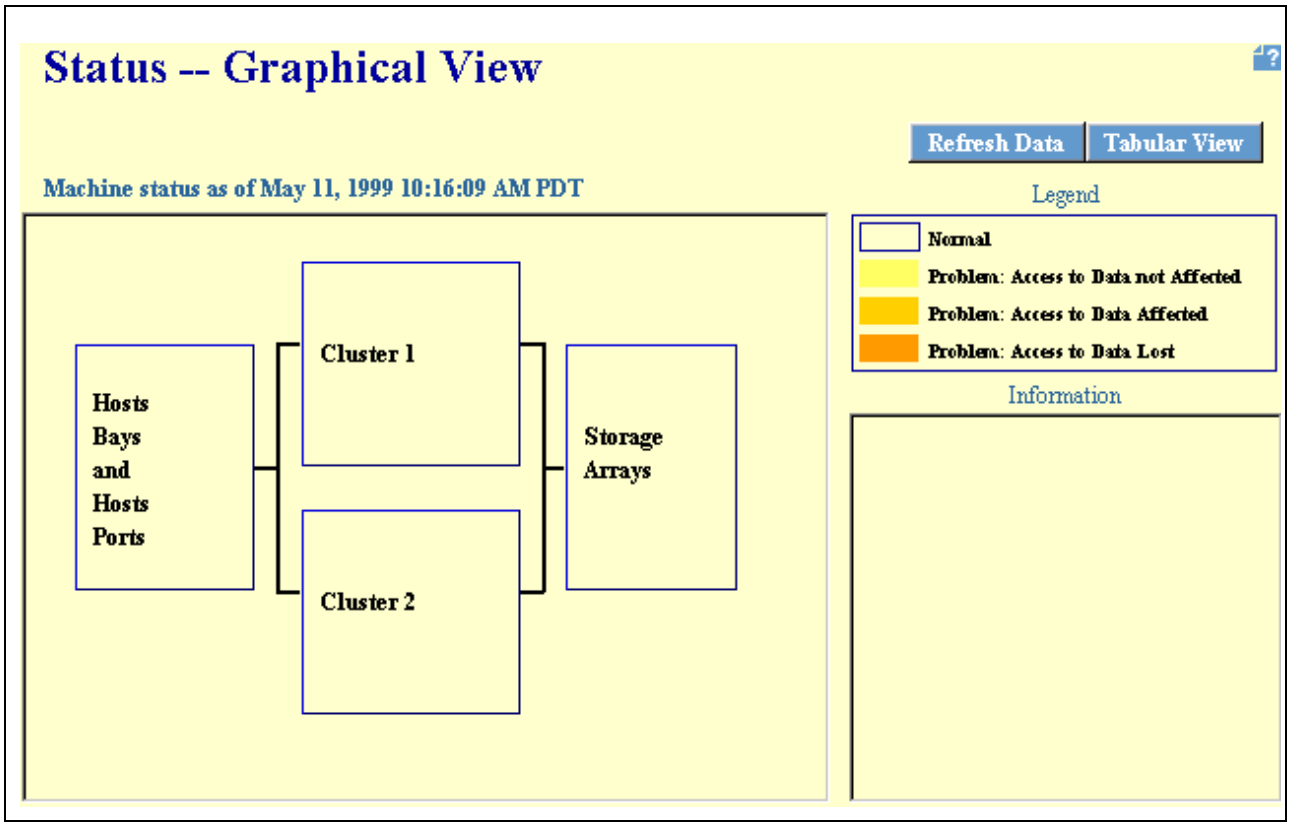

Figure 32. ESS Specialist: Status -- Graphical view panel

A failed component is highlighted in different colors. The severity of the problem determines the color used. The different colors are:

- White - Normal status. No failure has occurred.
- Yellow - Failure status. A component has failed, but access to the data is not affected. Performance may be degraded, however.
- Orange - Failure status. A component has failed, and access to the data is affected. One or more hosts have temporarily lost access to one or more logical volumes defined in the storage facility.

On the Status -- Graphical View panel, click the **Tabular View** button to access the Status -- Problem Log panel (see Figure 33 on page 81). The panel contains a list of the problems that occurred on the date recorded in the Problem Age field.

Note: Switching from a refreshed Status -- Graphical View panel to the Tabular View (Status -- Problem Log) does not mean that the list you will receive is refreshed as well. Use the Refresh Data button every time you switch between the Status panel and the Problem Log, and refresh the panel every time you want to see the current status.

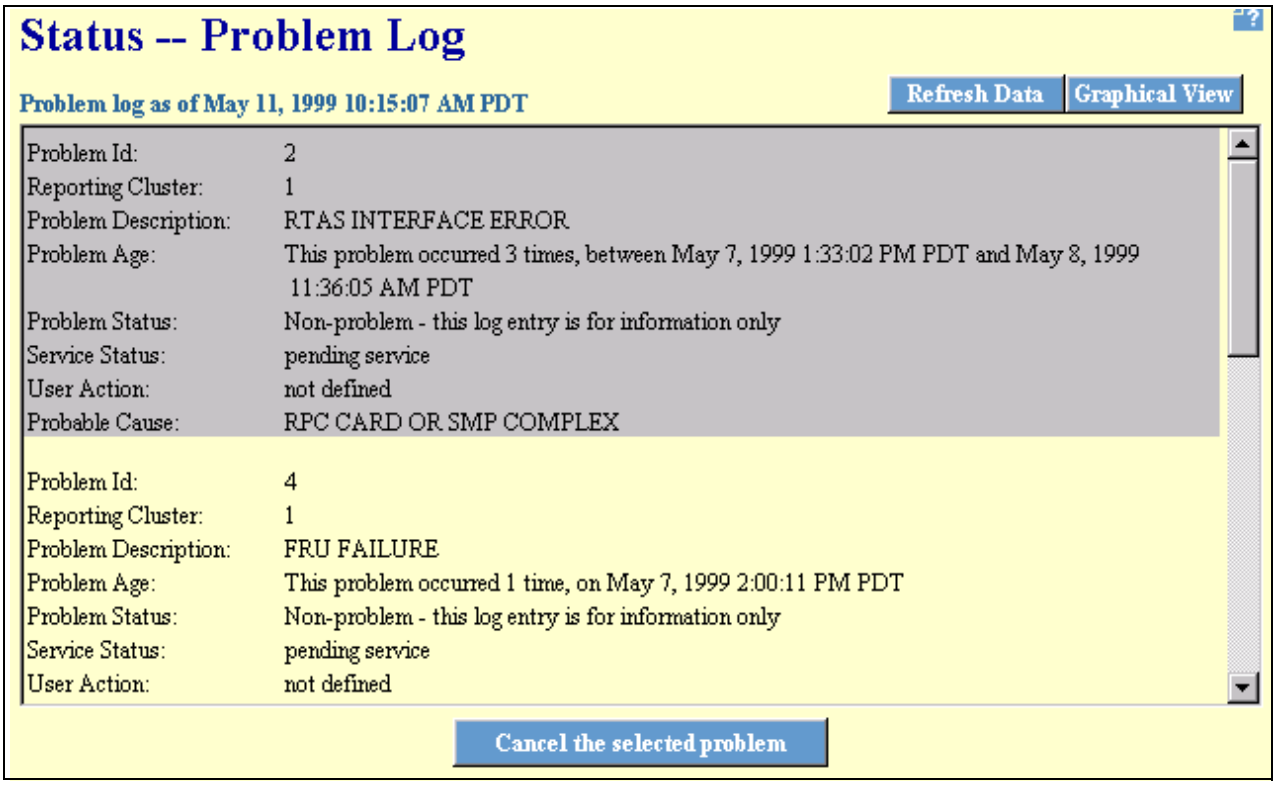

Figure 33. ESS Specialist: Status -- Problem log panel

The current status can differ, according to the successful completion or failure of subsequent recovery actions. The fixed or expired problems are not listed on the Status -- Problem Log panel.

Here is a brief description of each field:

- Problem ID: A unique ID for each problem
- Reporting cluster: List the effected Cluster ID
- Problem description: Brief description of problem
- Problem age: Date, time and frequency of problem
- Problem status: Severity of the problem
- Service status: Shows whether problem is “in service” or “pending service”
- User action: When displayed, contains suggestions of the type of service required by the IBM CE or the customer
- Probable cause: When displayed, indicates which component or machine function may have failed

You can delete a specific problem record by clicking anywhere within the problem record details area, the selected record will then be highlighted, as shown in Figure 33, then press the **Cancel the selected problem** button.

5.6.3 Problem notification panel

The ESS sends notification of a problem to a user, either through e-mail or by calling the specific phone number of a pager. The ESS Specialist Problem

Notification panel (Figure 34) shows the e-mail addresses and phone numbers, as well as the SNMP trap destinations.

The settings on the Problem Notification panel are enabled when the IBM CE installs the ESS machine. The information is entered from the service terminal and can be updated later via the ESS Specialist administrator.

The e-mail, pager and SNMP information can be modified by selecting the appropriate button.

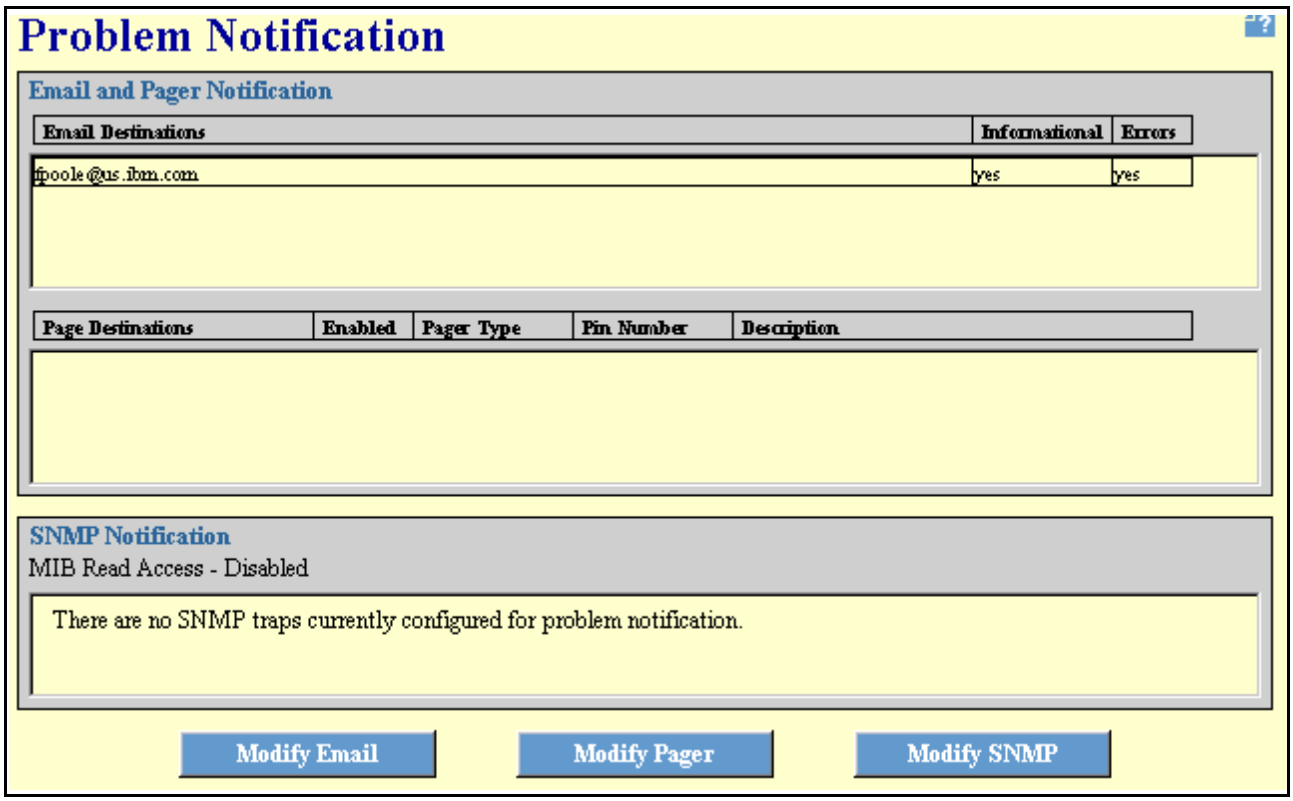

Figure 34. ESS Specialist: Problem notification panel

5.6.4 Communications panel

The Communications panel (see Figure 35 on page 83) shows the following information regarding the network and communications-related information:

- TCP/IP setup for each cluster
  - Hostname for each cluster
  - IP (dotted decimal) address
  - Subnet mask
  - Default router
  - Domain name and server address (if any)
- The ESS Remote Support functions, call home and remote service access, are enabled or disabled by selecting the appropriate button.

Note: The TCP/IP window is for informational purposes only. The IBM CE is required to modify these details.

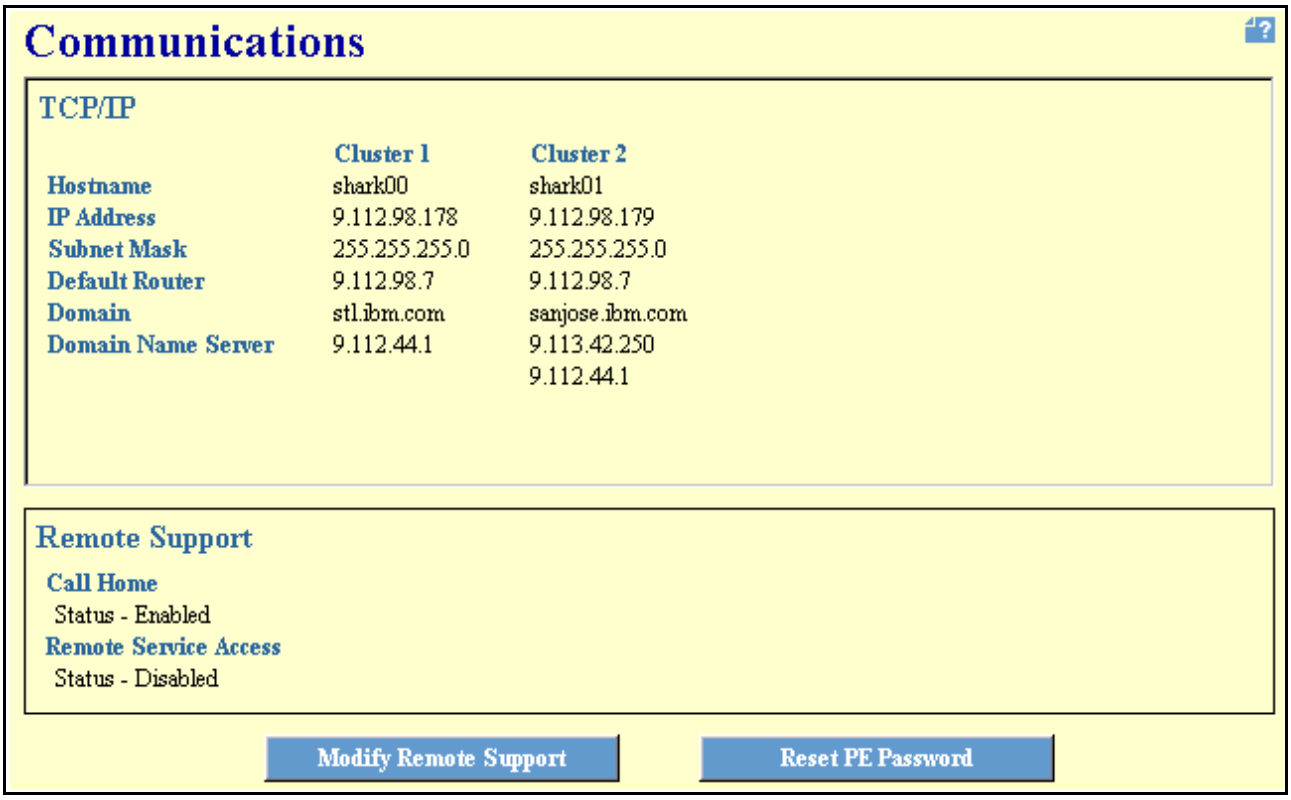

Figure 35. ESS Specialist: Communications panel

5.6.5 Storage Allocation panel

The Storage Allocation panel is the most important and detailed of the ESS Specialist panels. When you click the Storage Allocation hyperlink on the ESS Specialist main panel, the Storage Allocation -- Graphical View panel appears (see Figure 36 on page 84).

5.6.5.1 Storage Allocation -- Graphical View Panel

The Storage Allocation -- Graphical View panel provides an authorized user with a graphical representation of the hosts, host adapters (Escon or SCSI), Device adapters (SSA Adapters), Clusters, and Arrays on the ESS machine.

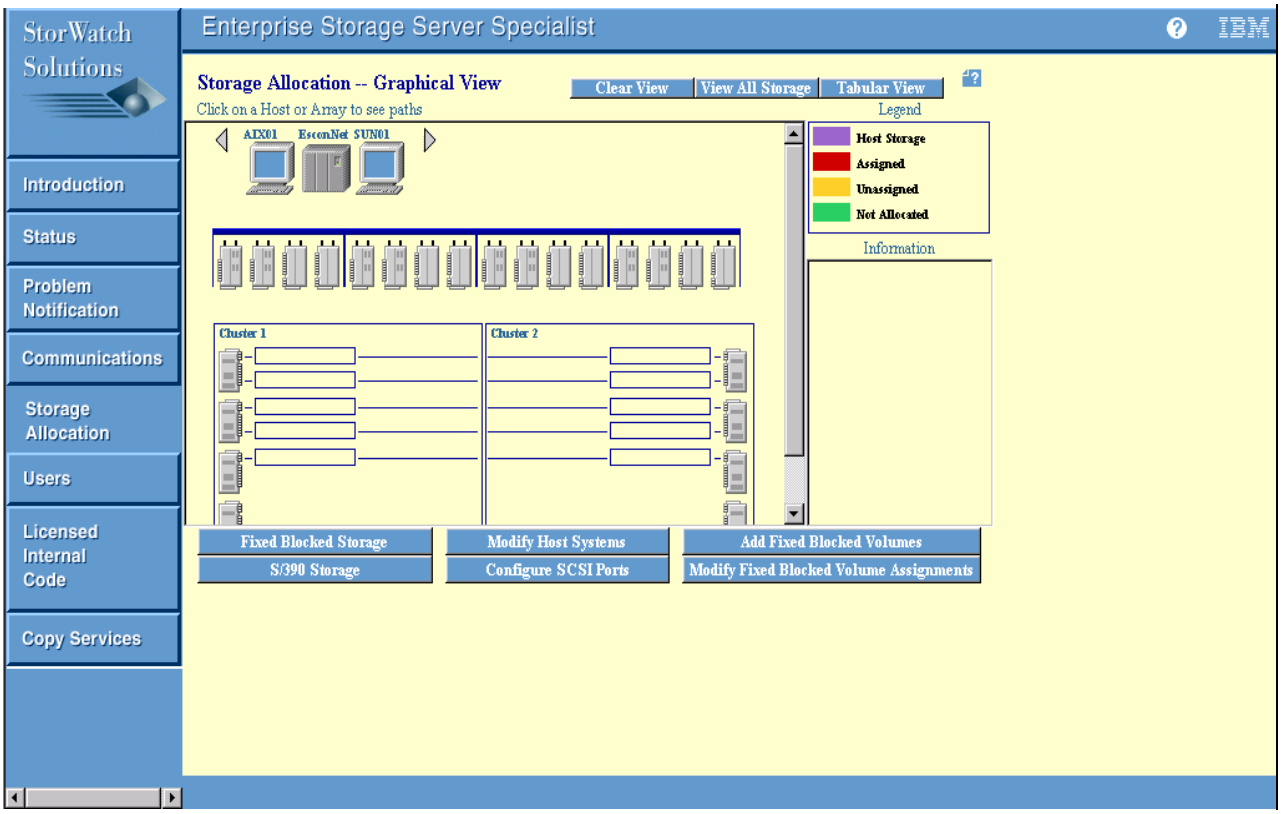

Figure 36. ESS Specialist: Storage Allocation panel; Graphical view

The first object line represents the hosts. It is a horizontally scrollable bar where the defined hosts are shown. If more hosts are defined than can be displayed in the panel, use the arrows on the left and right side of the host display to scroll to the remaining hidden host object representation.

Selecting a host will cause the panel to refresh and display connections to the host adapters and highlight the disk arrays assigned to this host; see Figure 39 on page 86.

The EsconNet host represents all escon-attached host systems. The SCSI-attached hosts are shown individually, as shown in Figure 37 on page 84.

Figure 37. ESS Specialist: Storage panel - hosts

The second object line shows the four host adapter bays with up to four host adapters in each bay. The bays are labelled left to right, 1 to 4. The adapters are also labelled left to right, 1 to 4 within each bay. The adapter ports are labeled A and B.

ESCON and SCSI host adapters are represented by slightly different icons; see Figure 38 on page 85. In the storage allocation panel example shown in Figure

36, each of the four adapter bays has two SCSI host adapters and two ESCON adapters installed.

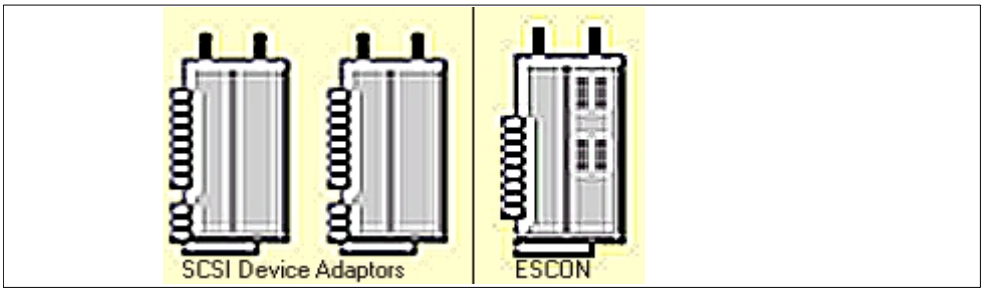

Figure 38. ESS Specialist: Storage panel - Host adapters

The third object line consists of two elements: cluster 1 and cluster 2. The clusters are two separate components internally connected to each other. You see these as separate entities. On the outer side of the two clusters are the four SSA device adapters for each cluster. The SSA adapters are numbered 1 to 4. Connected to each adapter are the two loops, A and B, with up to six disk arrays or disk groups attached to each loop.

From the Storage Allocation -- Graphical View panel, each time you select one of the elements, the information window positioned at the bottom right corner of the panel shows the details of the selected object. For example, selecting a:

- Host Adapter - Shows bay number 1 to 4, adapter number 1 to 4, and port A or B
- Device Adapter - Shows cluster 1 or 2, and DA number 1 to 4
- Array - Shows cluster, adapter, loop ID A or B, disk array group 1 to 6, and array type and storage allocation details

See Figure 39 on page 86.

The Clear View button does not have a connecting panel. You use it to clear any selection that you have made on the panel. The Storage Allocation -- Graphical View panel is the only interactive, nonstatic panel of the ESS Specialist. You can select hosts, host adapter ports, and arrays. If you want to change the selection, you can use the Clear View button to refresh the Storage Allocation -- Graphical view panel.

The View All Storage button will highlight all the storage arrays with the appropriate colors, as detailed below.

5.6.5.2 Storage allocation—Graphical view panel: Select a host

When you select a host from the top line on the Storage allocation—graphical view panel, you are not switched to another panel. The host adapter/ports assigned to the host are linked by a connecting line, and the adapter ports are colored. The disk arrays assigned in each cluster are scaled in color to signify different capacity information (see Figure 39 on page 86). The legend in the top right corner of the frame explains the colors:

- Red: Space allocated to logical volumes and accessible to one or more hosts.
- Orange: Space allocated to logical volumes, but not yet assigned to any host.

- Green: Space not yet allocated to logical volumes.
- Purple: This color has the same meaning as red, but it is used to show how much of the space is assigned to a particular SCSI host when a SCSI host is selected in the view. Purple is used to distinguish between space that is allocated to volumes accessible to that particular host, as opposed to the space that is allocated to all the other logical volumes (which is shown in red).

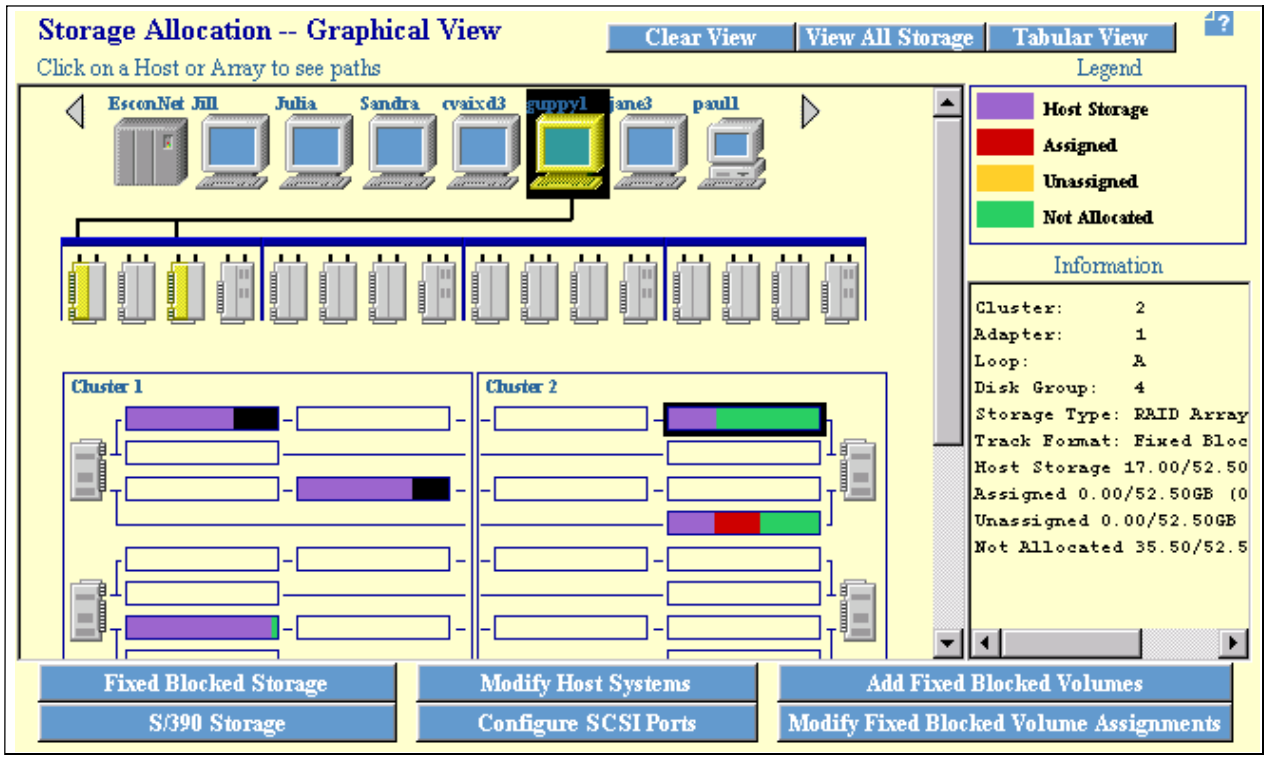

Figure 39. ESS Specialist: Storage allocation -- Graphical view: Host selection

From the Storage Allocation -- Graphical View panel, you can simply use your mouse and click on the components to see what they are and where they are allocated. The information related to the selected component appears in the information window as shown in Figure 39.

5.6.5.3 Storage Allocation -- Tabular View

Select the **Tabular View** button to see the Storage Allocation -- Tabular View List of Assigned Volumes panel (see Figure 40 on page 87).

The panel is a table with eight columns, and as many rows as necessary to describe all of the hosts and their assigned volumes.

Figure 40 shows two S/390 3390 Model 3 volumes and two SCSI logical disk devices.

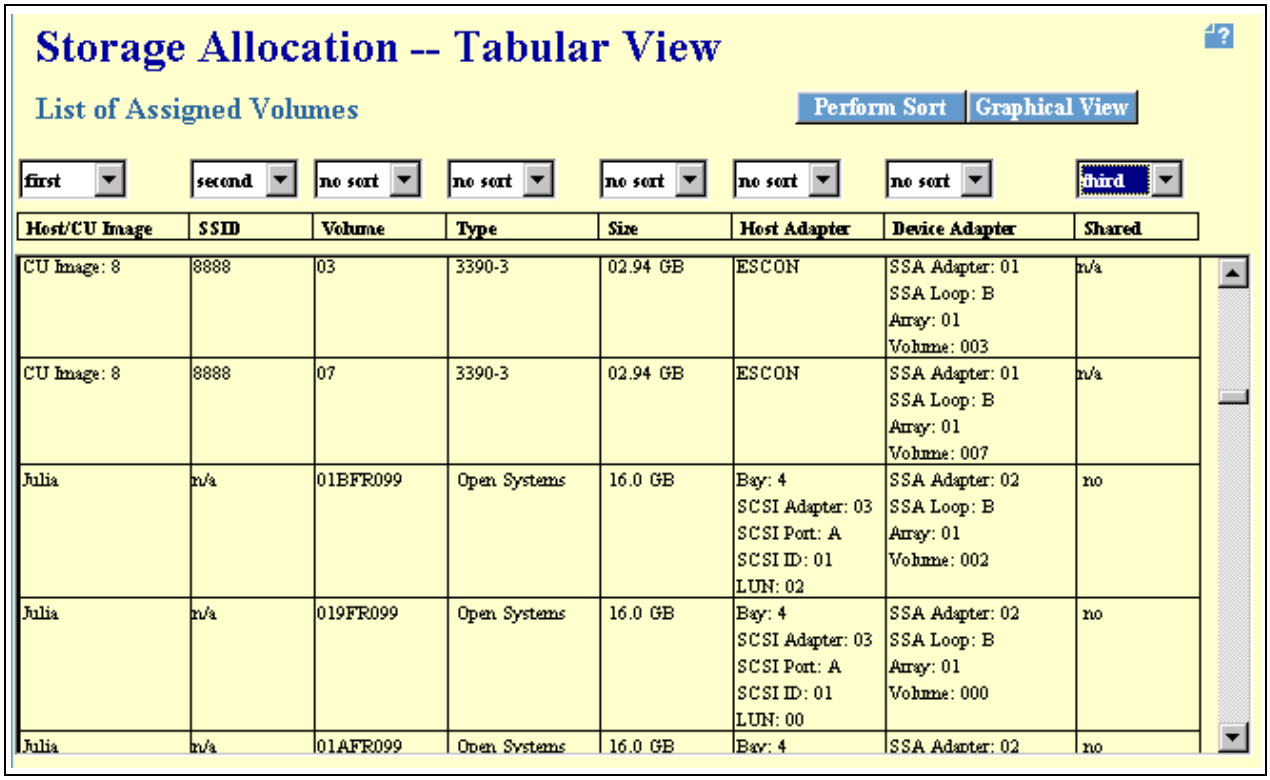

Figure 40. ESS Specialist: Storage allocation -- Tabular view panel

The host column shows the attached SCSI host name or S/390 Logical Control Unit. If the volume is shared with other SCSI hosts, there will be an entry for each host in the table, and no multiple hosts shown in a single row.

The volume column shows the unique number assigned to the SCSI volume at the time of its definition (LUN serial number), or the S/390 Logical device number. The LUN serial number is useful for identifying the host or hosts when the volume is shared between compatible hosts.

The size and type column shows the size in GB of the logical device or LUN and the emulated device type for S/390.

The host adapter column contains multiple information. Each row shows the bay number, the SCSI adapter, the port being used (A or B), the SCSI target ID, and the LUN for the volume.

The device adapter column lists the SSA Device Adapter, the SSA loop and the array number on which the volume resides. The volume logical number is also shown.

The shared column shows *Yes* when the volume is shared with another SCSI host. Sorting this column is useful for finding the other host sharing this volume.

Above each column in the table there is a sort facility pull-down list. Each list allows the selection of: No Sort, First, Second, and Third. With the exception of No Sort, which can be selected for multiple columns, the First, Second, and Third values can be selected only once in a column. If one of the values is selected, it will not be available in the other pull-down lists, and the Java applet automatically

removes the selection from the other lists. This permits a maximum of three sortable columns. The column with the First value set in the sort list field will be sorted first, then the second and the third, if defined. To run the sort, click on the **Perform Sort** button, and the columns will be ordered in the selected sequence.

5.6.6 Users panel

The Users hypertext link on the ESS Specialist main panel takes you to the User Administration Panel. With administration authority, you can grant access to the ESS, via the ESS Specialist, to other users in your organization.

The User panel is an information-only panel with four columns:

1. Username
2. Access level
3. IP address range (optional)
4. User comment (optional)

The User Administration panel allows sorting of all fields, and a maximum of three fields can be sorted concurrently. The same method as described in 5.6.5.3, “Storage Allocation -- Tabular View” on page 86.

The username is the name given to the user in the User Account Name field during user definition, as shown in the Modify Users panel, Figure 41 on page 88.

Pressing the Modify Users button will bring up the Modify panel, Figure 41 on page 88.

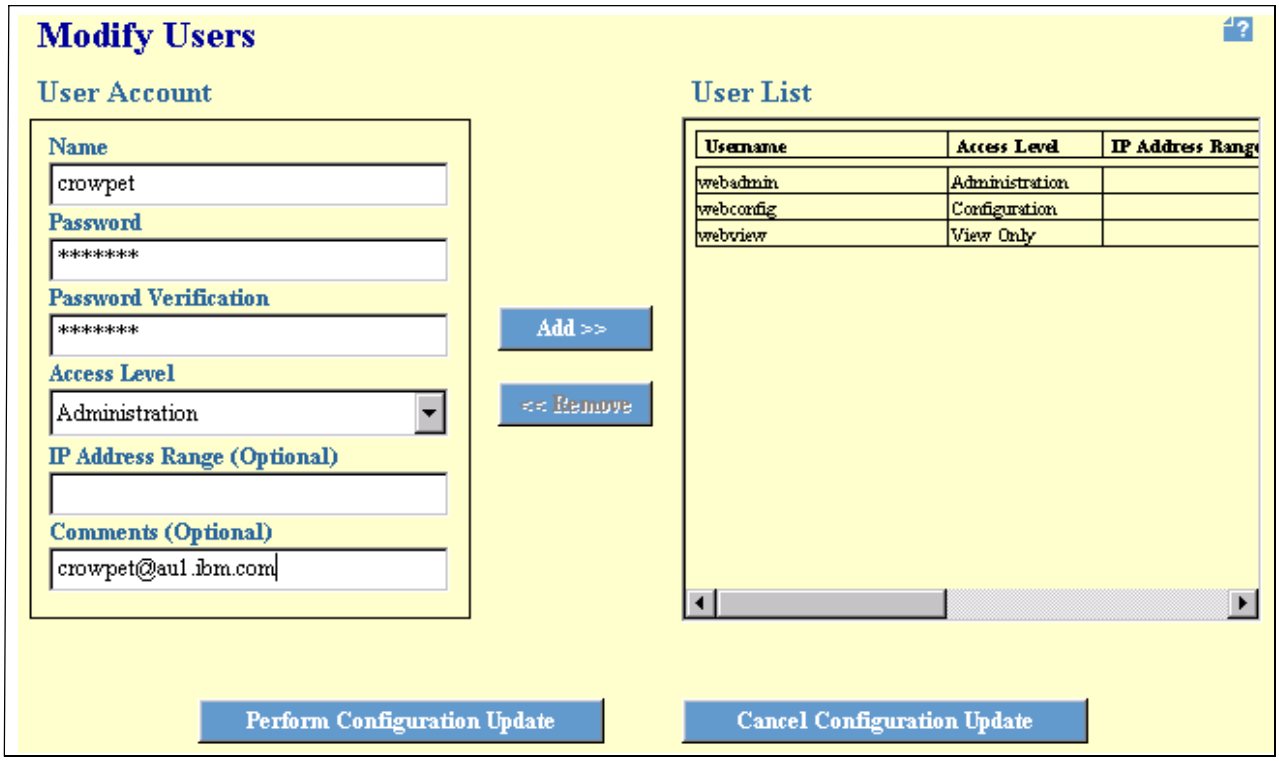

Figure 41. ESS Specialist: User administration - Modify users panel

Users can have three different levels of authorization and can be restricted to accessing the ESS Specialist from a specific IP address or a range of IP addresses. If you choose to restrict users to specific IP addresses, you may have to create alternative user IDs from other IP addresses in case part of the IP network is unavailable. With this security restriction, you may not be able to connect to your ESS Specialist if you are connecting through a provider that assigns you to a temporary randomly selected IP address (DHCP - Dynamic Host Configuration Protocol).

Figure 41 on page 88 shows three existing users, in the User List window, of different authority levels. A new user definition is in progress in the User Account window.

User comments are useful when the given usernames are not self-explanatory and do not allow fast recognition of the person accessing the machine. We recommend using the User Comments field to record a contact name for and telephone number of the username.

5.6.7 Licensed internal code panel

From the ESS Specialist main panel, click on **Licensed Internal Code** to view a list of the levels of LICs in cluster controller 1 and cluster controller 2. This panel is for informational purposes only.

The table on the LIC panel has four columns:

1. LIC Source - One of the three LIC types listed above.
2. Version - The version of the LIC in a unique standard format.
3. Activation Date - The date of activation (only for LIC source types of Active and Previous).
4. Additional Info - A hypertext link to the README file associated with that LIC (see Figure 43 on page 90).The README file typically contains more detail about the specific level of LIC, release date, and other useful information.

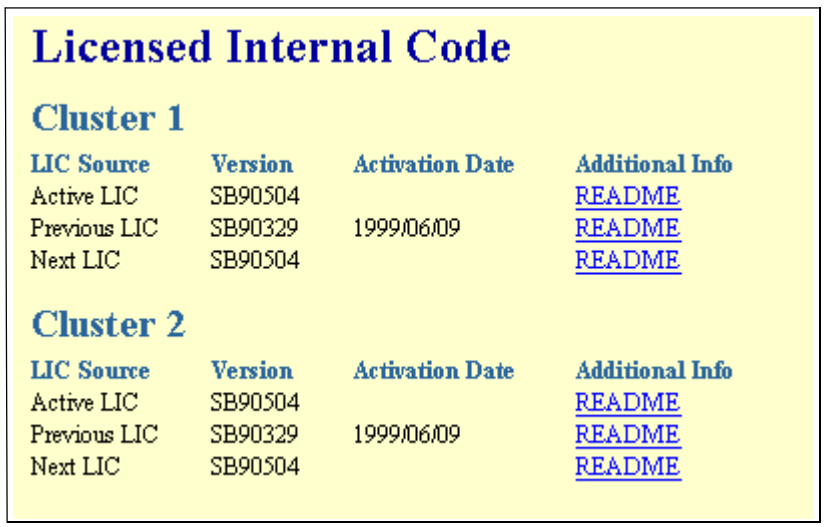

Figure 42. ESS Specialist: Licensed internal code panel

In each cluster controller section (as shown in Figure 42 on page 89) the LIC sources are:

- Active LIC - The current and active version of LIC that is controlling the operation of the appropriate cluster.
- Previous LIC - An older level of LIC, which is no longer active and has been replaced by the active version. We recommend keeping at least one back-level version of the code. If a newly activated LIC causes problems, it is then possible to revert back to the old code and avoid any outages.
- Next LIC - A newer level of LIC, which has been copied to the cluster processor but has not yet been activated.

Figure 43 shows the readme file associated with a LIC level within the ESS.

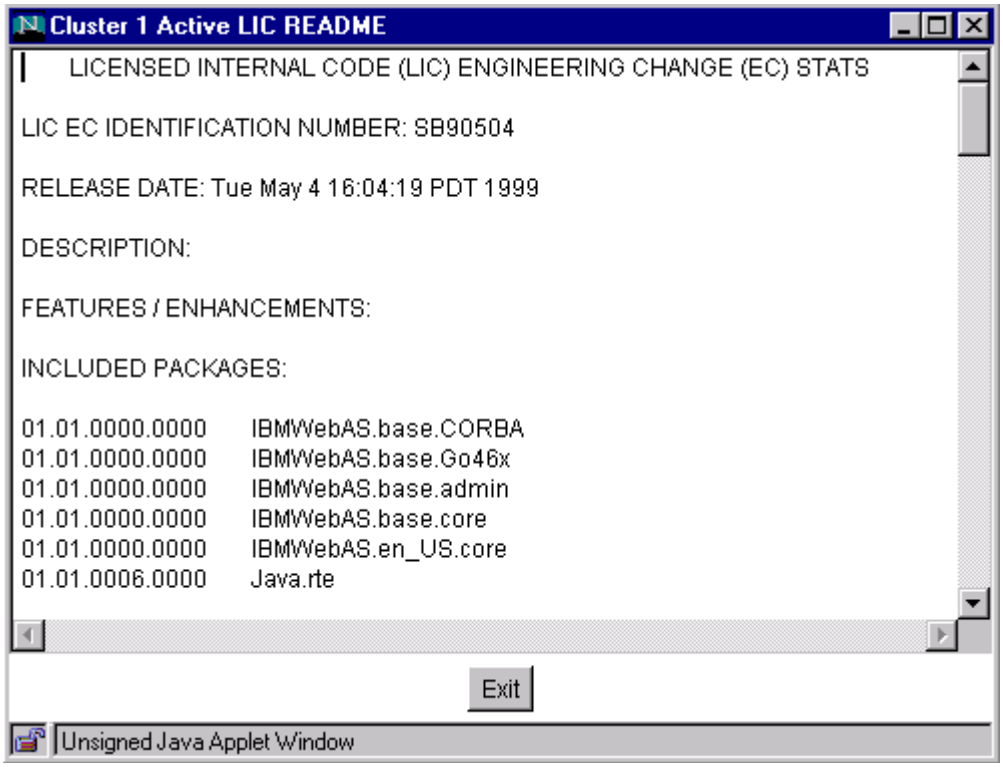

Figure 43. ESS Specialist: README file for an LIC source

5.6.8 Copy services

The Copy services panel provides the main menu for setting up and managing the PPRC and FlashCopy volume copy functions. Select the menu button to accomplish the required task.

- Volumes - Display volumes and define them as source or targets and initiate the copy
- Control Unit - Display physical ESS and its logical control units information
- Paths - Define and display the status of PPRC/FlashCopy paths between control units
- Tasks - Manage tasks you have defined, for example run tasks, remove, export or import tasks

- Configuration - Nominate Copy Services primary and backup server, also displays a Copy Services error log
- Exit Copy Services - Returns you to the ESS Specialist menu panel

Figure 44 shows the Copy Services main menu panel with the selection frame on the left.

Figure 44. ESS Specialist: Copy services main menu

5.7 ESS Expert

Another valuable tool for use with the ESS is the StorWatch ESS Expert. The ESS Expert is accessed via your browser and is used to provide storage resource management functions for IBM storage servers like the ESS. The Expert prepares reports and graphs from data it has collected.

For more information and sample reports see 6.1, “Using ESS Expert” on page 93.

Chapter 6. Enterprise Storage Server Expert (ESS Expert)

In this chapter, we discuss the Enterprise Storage Server Expert.

The IBM Enterprise Storage Server Expert (Expert) provides storage resource management functions for the IBM storage servers.

6.1 Using ESS Expert

Expert, a Web user interface, is an optional feature that you can purchase to use with IBM storage servers.

Expert supports the StorWatch management disciplines:

- Asset management
- Capacity management
- Performance management
- Configuration management

Expert prepares reports and graphical charts from the data it has collected. It produces performance reports on disk utilization and cache holding time.

6.1.1 The user interface components

The Expert interface consists of these major components:

- The information frame.

  You can press the **Help** icon at any time to bring up a new browser window that has the help system.

- The navigation frame.
- The work frame is 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 tasks. Every screen displayed in the work frame contains a page-help icon.

  When you click the help icon, it links to the help system and displays the help information specific to the current work frame screen. The help system is always contained in a separate browser window, so you may have to switch from the main browser window to the help browser window after clicking the help icon to actually see the displayed help information.

6.1.2 Previews

Some of the reports have a preview section. These columns are located at the right side of the report. They give a summary or preview of what you will see in the detailed reports that Expert displays when you click any of the highlighted links. These previews can save you time. If there are no performance statistics in the preview that warrant investigation; you need not display the detailed report.

6.1.3 Accessing Expert

The Expert can be accessed using a Web browser running in a Web Server PC attached directly to the ESS via Ethernet, or in a remote machine if the ESS is connected into the user's network.

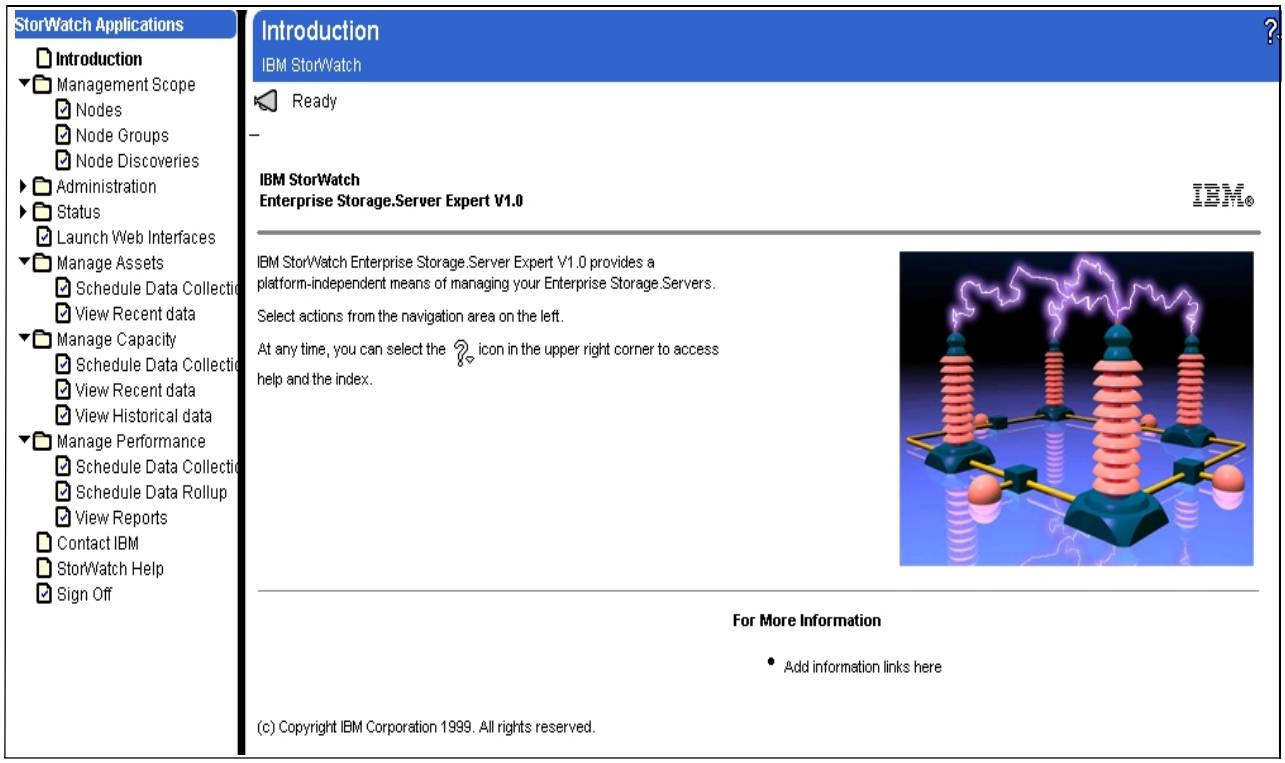

Figure 45. StorWatch introduction panel and menu

You can open the folders to display Expert options by clicking on the arrows to the left of the folders. Opening these folders displays

- Management Scope
  - Nodes
  - Node groups
  - Node discoveries
- Manage assets
  - Schedule data collection
  - View recent data
- Manage capacity
  - Schedule data collection
  - View recent data
  - View historical data
- Manage performance
  - Schedule data collection
  - Schedule data rollup
  - View reports

6.1.4 Managing scope

The first task you need to perform is Node Discoveries. During the performance of this task, you make the IBM storage servers known to Expert. Without this information, Expert would not know where to collect the data for the asset, capacity, and performance reports it creates.

1. Open the Management Scope folder to display its options by clicking on the arrow to the left of the folder (see Figure 45 on page 94)
2. Click the **Node Discoveries** option. Expert displays the following panel shown in Figure 46.

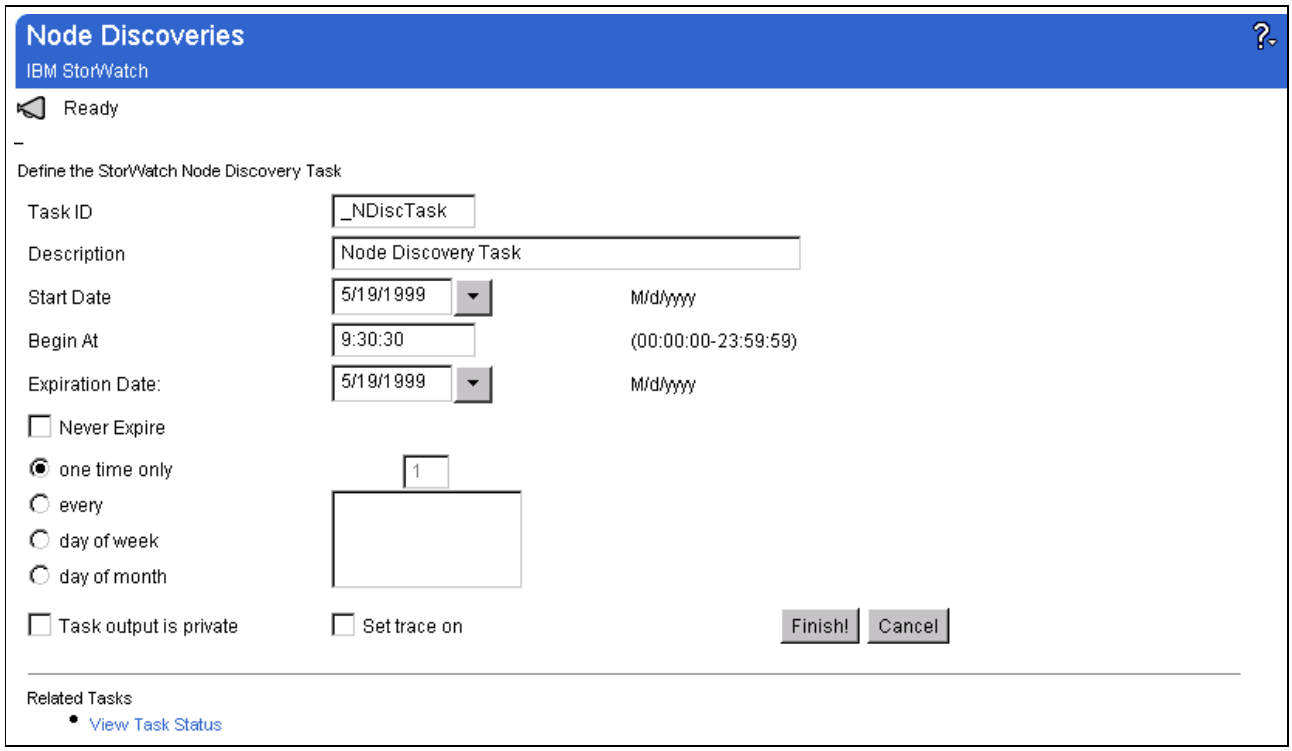

Figure 46. Node discoveries panel

Use this panel to define:

1. Task ID

   Enter some unique identifier.

2. Description

   Enter a description of the task.

3. Start date (m/d/yyyy)

   This is the date you want Expert to begin collection data.

4. Begin at (00:00:00-23:59:59)

   Enter the time that you want Expert to start collection data.

5. Expiration date (m/d/yyyy)

   This is the date you want Expert to stop collecting data.

   You may click the check box Never Expire instead.

6. To schedule interval, click the check box:
   - One time only
   - Every:
     - Day of week
     - Day of month

   If you selected Every day of week, or Every day of month, go to the box to the right of the panel and click a day of the week (Monday, Tuesday, and so forth) or a day of the month (1, 2...31).
7. Click the check box of **Task output is private** to prevent unauthorized view of the reports.
8. Click the check box **Set trace on** to generate error-tracking reports.
9. If you are satisfied with the settings, click **Finish.**
10. If you don't want changes to be applied, click **Cancel.**

6.1.5 Managing assets

The Manage asset function for Expert collects data on the IBM servers it *discovered* during the node discovery task. You can request that Expert provide this information as a summary by server type, or detailed information on a specific server. Expert can also provide information on trends in your storage server growth.

The Manage assets options include:

- Schedule data collection

  This panel allows you to set the period over which you want Expert to collect the data, and to select the servers for which Expert will collect the data.
- View recent data

  This panel allows you view asset information that Expert just collected.

6.1.5.1 Schedule data collection

Use this panel to define, modify, or view a data collection task.

- Click **Schedule Data Collection** from the selections under Manage Assets (see Figure 45 on page 94). From the next panel (Figure 47 on page 97):
  1. Click **Create** to define a data collection task.
  2. Click **View** to view a data collection task that already exists.

     Note: This selection only becomes available when you check the box to the left of a task.
  3. Click **Edit** to change a data collection task that already exists. This function allows you to change the parameters for the task you select.

     Note: This selection only becomes available when you check the box to the left of a task.
  4. Click **Delete** to delete a data collection task that already exists.

     Note: This selection only becomes available when you check the box to the left of a task.

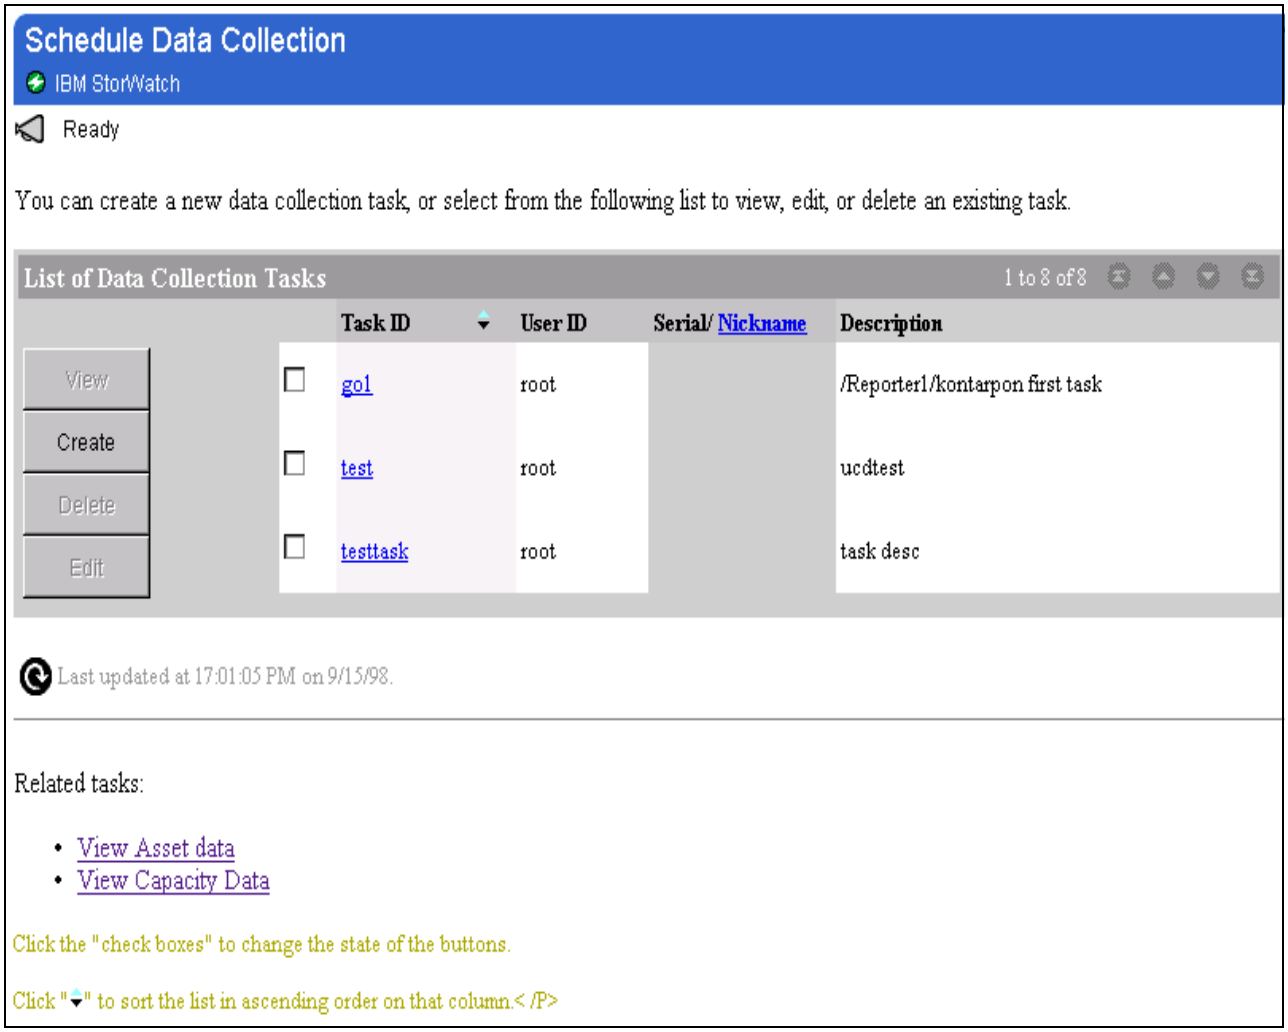

Figure 47. Schedule data collection panel

Creating a New Data Collection Schedule

When you click **Create** on the Schedule Data Collection panel (see Figure 47), you are given the option to schedule data collection for specific IBM storage servers.

To define a new data collection task, enter task and scheduling information in the fields.

1. The first panel for Schedule Data Collection displays a List of Storage Servers by serial number. Click the box to the right of the server you want to collect data on, and enter a nickname in the Nickname field (optional).
2. The password fields are filled in for you. Click the circle next to Yes if you want to change the password, and fill in the panels that Expert displays for setting new passwords.
3. Click **next** to move to the next panel (see Figure 48 on page 99).
4. Enter a task identifier in the Task ID field.

   Enter a unique identifier for this task. For example, *data on server1.*

5. Description.

   Enter a description of the task. For example *data collection for August*. This is a required field. Expert will return an error code when you try to move to the next panel without filling in this field.

6. Enter the Start Date (mm/dd/yyyy PST).

   This is the date you want Expert to begin collection data. For example, *07/29/1999.*

7. Enter the Expiration Date (mm/dd/yyyy PST).

   This is the date you want Expert to stop collecting data. For example, *08/29/1999.*

8. Enter the Begin at time as hh:mm:ss PST.

   This is an optional field. Enter the time that you want Expert to start collection data. For example, *12:00:00*. The default is midnight on the Start Date **???.**

9. Click the check box to the left of Set trace on to generate error-tracking reports.

10. Click the check box to the left of Public read access to task/output to allow IBM or your authorized service provider to view the reports.

11. For the Choose Schedule Interval field, click one of the following options:

    a. One time only

    b. Every

       1. Day of week

       2. Day of month

12. If you selected Every Day of week, or Every Day of month, go to the box to the right of the panel and click a day of the week (Monday, Tuesday, and so forth) or a day of the month (1, 2...31).

13. Click **Next** to display the task settings you have defined.

14. If you are satisfied with the settings, click **Commit.**

15. If you want to change the settings, click **Back to List.**

16. On the last panel, click **Done.**

To view, edit, or delete a data collection task you can (see panel on Figure 47 on page 97):

1. Click the Task ID for the task in the Task List.

2. Click **View** to view a data collection task that already exists.

3. Click **Edit** to change a data collection task that already exists. This displays the fields for scheduling the task. Enter your changes, then click **Commit**.

4. Click **Delete** to delete a data collection task that already exists. Verify that you want to delete this task by clicking **Delete** on the verification panel that Expert displays.

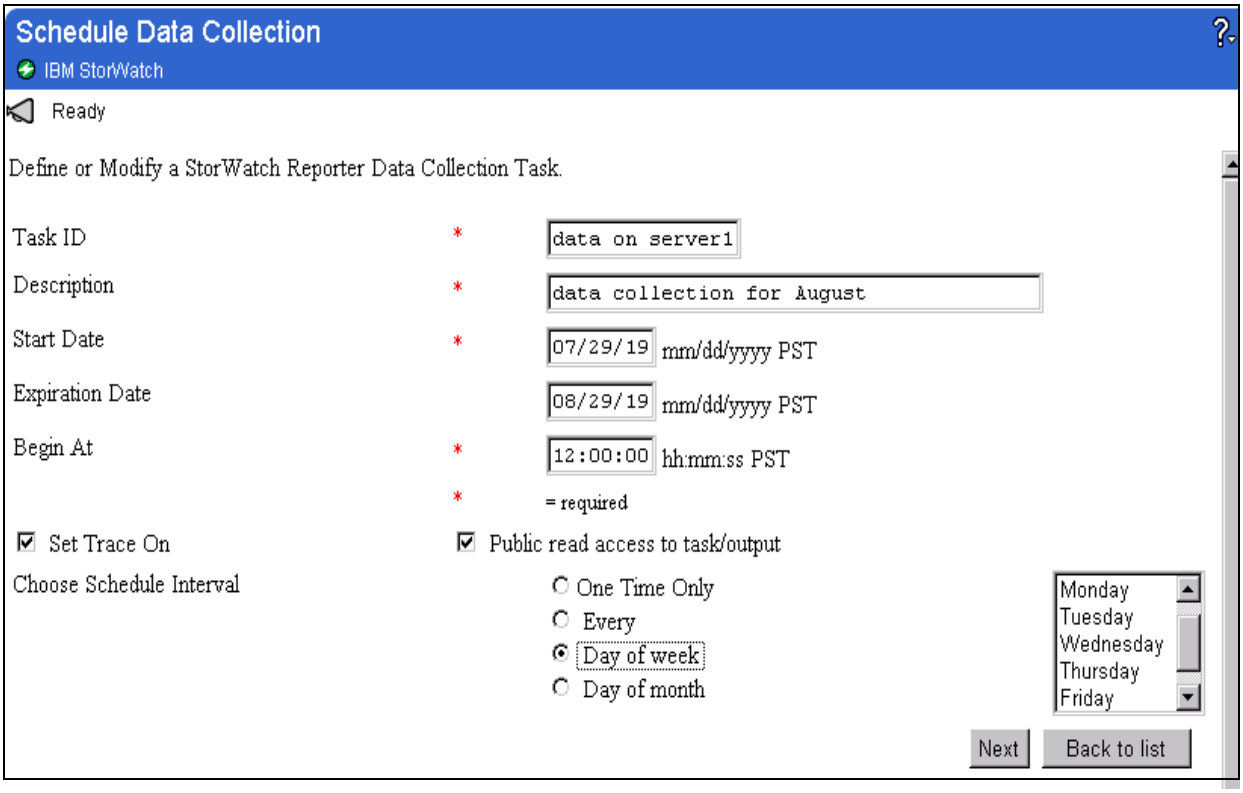

Figure 48. Creating a new data collection task, second panel

6.1.5.2 View recent data report

To view a recent data report:

- Click **View Recent data** from the selection under Manage Assets (see Figure 45 on page 94).

For example of this report, see Figure 49 on page 100.

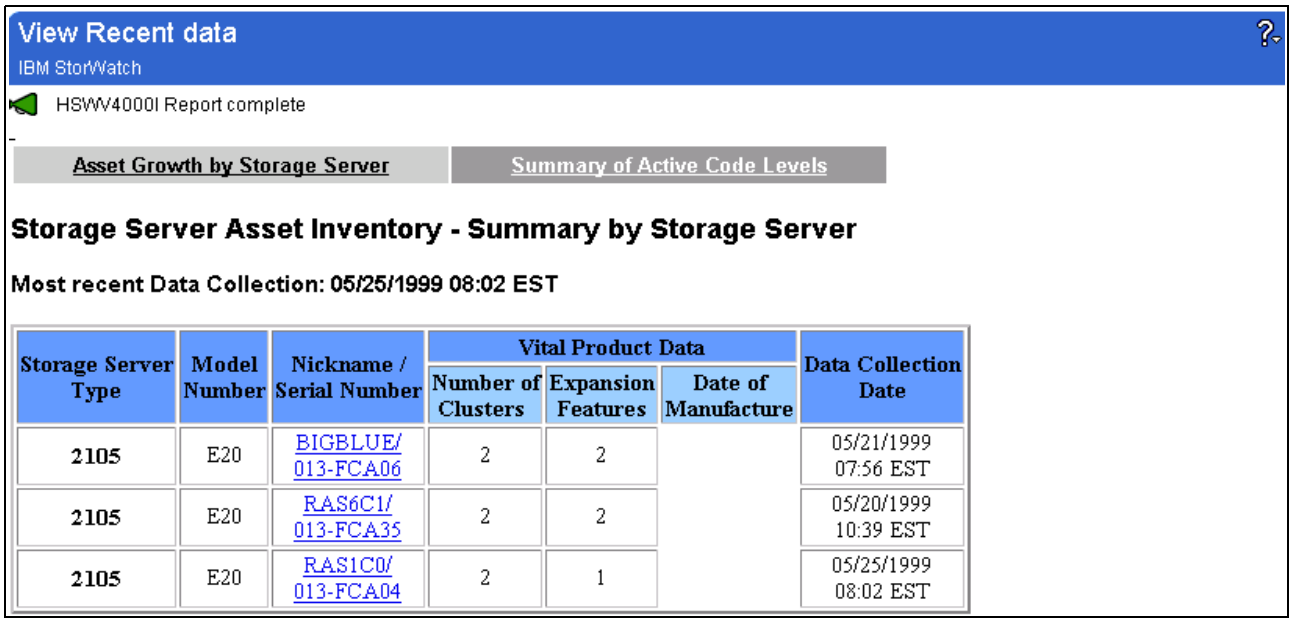

Figure 49. Storage server asset inventory - Summary report

Table 7 provides an explanation of the fields used in the previous figure.

Table 7. Storage server asset inventory - Storage server summary

Field	Description
Storage Server Type	The storage server machine type, such as IBM ESS 2105.
Model Number	The model number of the machine, such as E20.
Nick Name / Serial Number	This is the serial number for the machine, such as 013-FCA06. If you assigned a nickname to this machine (an optional task) Expert displays the name here.
Vital Product Data	The three fields under Vital Product Data (VPD) display information that Expert has gathered on each machine.
Number of Clusters	For IBM storage servers that are part of the Seascape family of products, the clusters are always in two. The clusters provide failover protection to the IBM storage servers.
Expansion Features	Expert lists the number of expansion racks with additional DDM that you have attached to your IBM storage server.
Date of Manufacture	This is the date and time of manufacture for your machine. You need this information for IBM or your service provider.
Data Collection Date	This is the date and time that Expert captured recent data.

6.1.6 Managing capacity

The Manage Capacity function for Expert collects data on the IBM servers it "discovered" during the node discovery task. You can request that Expert provide this information as a summary of capacity, in gigabytes (GB) by server type, or detailed information on a specific server. Expert can also provide this summary by server model type, or provide a summary of the capacity of each host system to which you have assigned space on a storage server.

The Manage Capacity options include:

- Schedule data collection

  This panel allows you to set the period over which you want Expert to collect the data, and to select the servers for which Expert will collect the data.

- View recent data

  This panel allows you view capacity information that Expert just collected.

- View historical data

  This panel allows you to select how Expert displays historical data to you in a report, over a specified time period.

  - Growth by storage server capacity in GB
  - Growth by host system capacity in GB

The manage capacity panels give you:

- A list of serial number and capacity, in gigabytes, for each IBM storage server
- A summary of servers organized by model
- A list of hosts that summarizes the capacity for each host having assigned space on one or more servers

6.1.6.1 Schedule data collection

Procedures for Schedule Data Collection, Creating New Data Collection and modifying a data collection task are the same as described in 6.1.5.1 on page 96. The only difference is that you have to:

- Click **Schedule Data Collection** from the selections under **Manage Capacity** (see Figure 45 on page 94).

6.1.6.2 Viewing the Recent Data Report

Figure 50 and Figure 51 show examples of the logical view and the physical view of the storage server capacity.

View Recent data

IBM StorWatch

HSWV4000I Report complete

Logical View: Storage Server Capacity - Summary by Storage Server

Most recent Data Collection: 05/25/1999 08:02 EST

Nickname/ Serial Number	Type of Storage (GB)			Total Raw Capacity (GB)	Formatted Storage (GB)			Control Unit Images	SCSI Attached Hosts	
	Raid	Non-Raid	Undefined		Fixed Block	CKD	Unformatted		Number	GB
BIGBLUE/ 013-FCA06	255	0	0	255	53	54	0	1	0	0
RAS6C1/ 013-FCA35	0	0	0	0	0	0	0	0	0	0
RAS1C0/ 013-FCA04	764	0	146	910	0	324	146	6	2	0
TOTAL	1019	0	146	1165	53	378	146	7	2	0

Figure 50. Example of the logical view: Storage server capacity - Summary report

Physical View: Storage Server Capacity - Summary by Storage Server

Nickname/ Serial Number	Storage Server Type	Host Adapters		Device Adapters	Disk Groups			Total DDMs	Cluster Info	Data Collection Date
		SCSI	ESCON		Raid Arrays	Non-Raid Groups	Free Ranks			
BIGBLUE/ 013-FCA06	2105	4	2	4	4	0	0	28	VIEW	05/21/1999 07:56 EST
RAS6C1/ 013-FCA35	2105	0	4	8	0	0	16	0	VIEW	05/20/1999 10:39 EST
RAS1C0/ 013-FCA04	2105	4	3	4	12	0	2	100	VIEW	05/25/1999 08:02 EST
TOTAL	3	8	9	16	16	0	18	128	--	--

Figure 51. Example of the physical view: Storage server capacity - Summary report

Table 8 explains the fields in the Logical View: Storage Server Capacity - Summary by Storage Server report

Table 8. Logical view: Storage server capacity - Summary by storage server

Field	Description
Nickname / Serial Number	This is the serial number for the machine, such as 013-FCA06. If you assigned a nickname to this machine (an optional task) Expert displays the name here - BIGBLUE, for example.
Type of Storage (GB)	The three fields under Type of Storage identifies the storage type allocated to storage server.
Raid	Amount of RAID-5 protected storage.
Non-Raid	Amount of Non-RAID protected storage - JBOD.
Undefined	Amount of allocated, but not defined, storage.
Total Raw Capacity	The total storage capacity of unformatted disks.
Formatted Storage (GB)	The three fields under Formatted Storage show the allocated formatted storage capacity.
Fixed Block	Amount of allocated FB formatted storage.
CKD	Amount of allocated CKD formatted storage.
Unformatted	Amount of allocated unformatted storage.
Control Unit Images	The number of defined Control Units (LCUs).
SCSI Attached Hosts	The two fields under SCSI Attached Hosts identify the number of SCSI Hosts and allocated storage.
Number	The number of SCSI-attached hosts.
GB	Amount of allocated storage to SCSI attached hosts.

Table 9 explains the fields in the Physical View: Storage Server Capacity - Summary by Storage Server report.

Table 9. Storage server capacity summary - physical view

Field	Description
Nickname / Serial Number	This is the serial number for the machine, such as 013-FR099. If you assigned a nickname to this machine (an optional task) Expert displays the name here.
Storage Server Type	This field identifies the storage server machine type, such as 2105.
Host Adapters	The information in the two fields under this umbrella shows the number of adapters.
SCSI	The number of host adapters for SCSI-attached host systems.
ESCON	The number of host adapters for ESCON-attached host systems.
Device Adapters	The number of device adapters for attached SSA devices (DDM).
Disk Groups	These fields describe the DDM groups.
RAID Arrays	The number of disk groups that are in RAID arrays.
Non-RAID Groups	The number of disk groups that are non-RAID. For example JBOD.
Free Ranks (Arrays)	The number of SSA loops that are free for disk groups.

Field	Description
Total DDMs	The number of disk drive modules (DDM) on the storage server.
Cluster Info	Click **View** in this field to see detailed information on the two clusters.
Data Collection Date	This is the date that Expert collected the capacity information on each server.

6.1.6.3 Viewing the historical report

To see historical data on capacities, click **View Historical Data**. The panel is then displayed (See Figure 52 on page 105). Enter information about the data you want to see in the fields:

1. Select the type of growth report you want to see. Use the spin button in the *Capacity Growth Report type* field to select **Growth by Storage Servers**, or **Growth by SCSI Attached Host**.
2. Enter the starting date, as mm/dd/yyyy in the Show capacity growth since field.
3. Enter the ending date, as mm/dd/yyyy in the Stop reporting growth on field.
4. Use the spin button to select a period interval (week, month, quarter) in the Show capacity growth by field.
5. After you have filled in the fields, click **Show Report**. The historical report is displayed. See Figure 53 on page 106 and Figure 54 on page 107 for examples of these two types of reports.

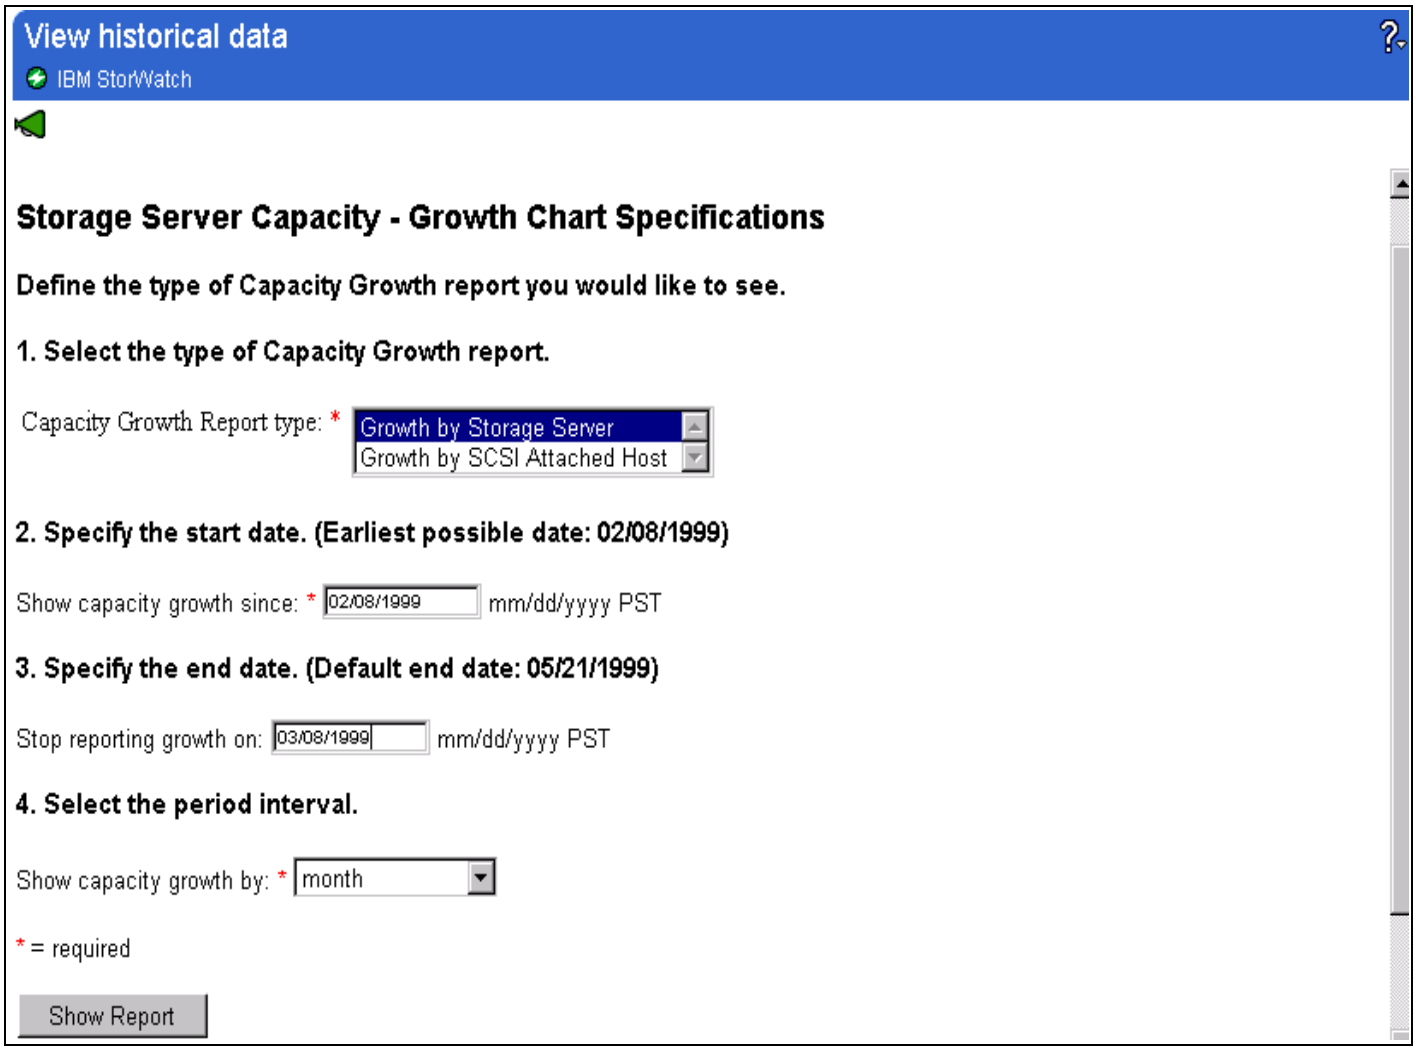

Figure 52. Setting up a historical data report for capacity growth

Capacity Growth by Storage Server

Figure 53 shows the Capacity Growth by Storage Server screen, and Table 10 provides an explanation of the fields used.

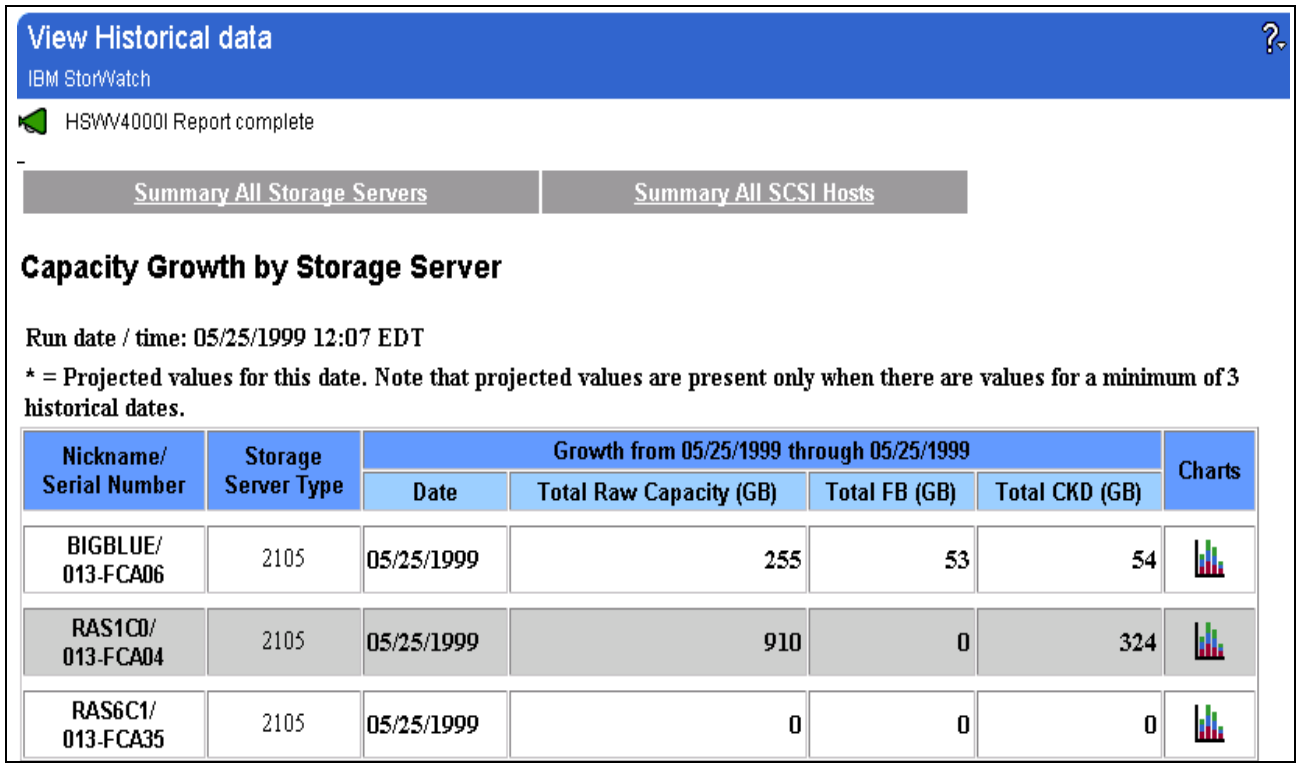

Figure 53. Example of the capacity growth by storage server

Table 10. Capacity growth by storage server

Field	Description
Nickname / Serial Number	This is the serial number for the machine, such as 013-FCA04. If you assigned a nickname to this machine (an optional task) Expert displays the name here.
Storage Server Type	The storage server machine type, such as 2105.
Growth from beginning date to ending date	The capacity growth information for each server during the time interval.
Date	The date that data was first detected, and consequently collected, by Expert.
Total Raw Capacity	The total storage capacity of unformatted disks.
Total FB (GB)	The storage that is formatted for fixed block (FB), in gigabytes (GB).
Total CKD (GB)	The storage that is formatted for count-key-data (CKD), in gigabytes (GB).
Charts	Click icon for graphical presentation.

Capacity Growth by Host

Figure 54 shows the Capacity Growth by Host screen, and Table 11 provides an explanation of the fields used.

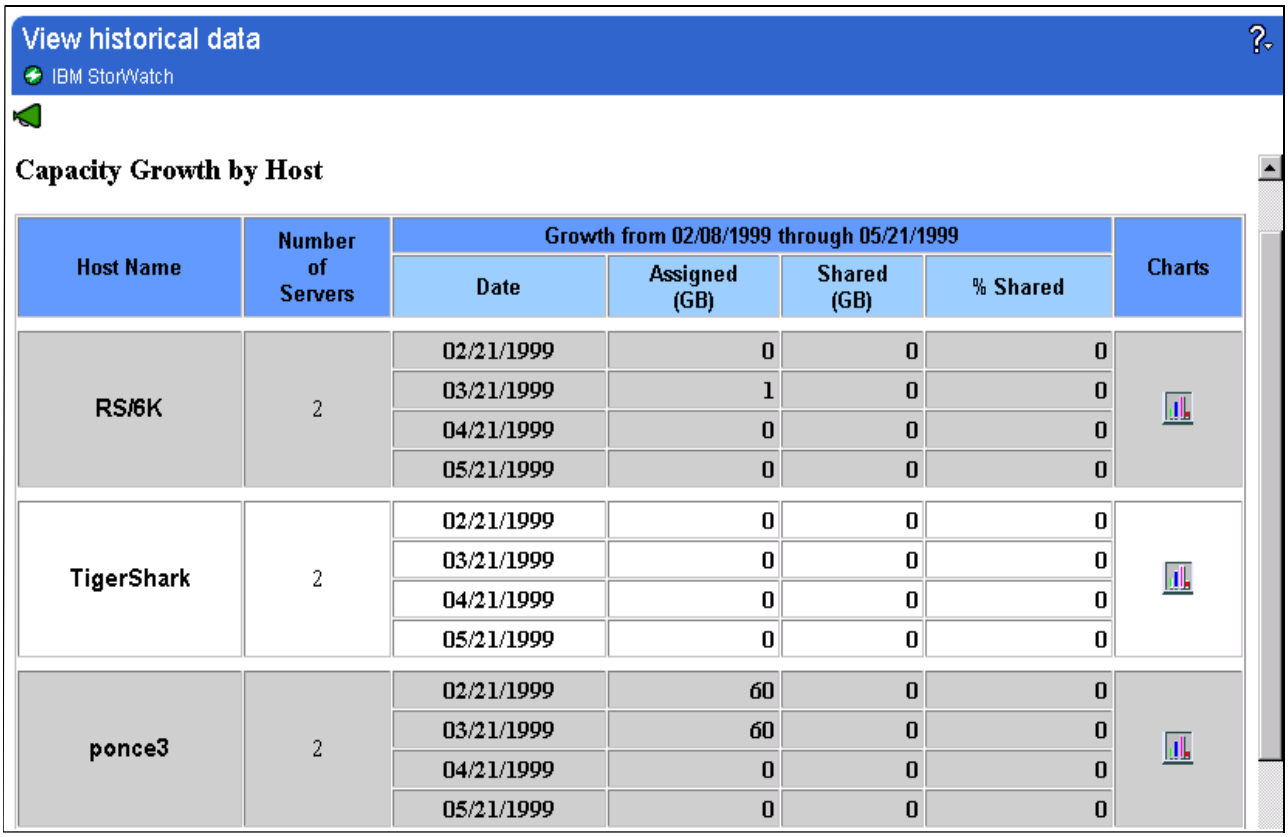

Figure 54. Example of the capacity growth by host

Table 11. Capacity growth by host

Field	Description
Host name	The name you have given to your host system that is attached to an IBM storage server
Number of Servers	The number of storage servers to which the host system is attached
Growth from beginning date to ending date	The capacity growth information for each host during the time interval
Date	The date that data was first detected, and consequently collected by Expert
Assigned (GB)	The total storage capacity on the server that is assigned to a host system
Shared (GB)	The storage that is shared with other host systems, in gigabytes (GB)
% Shared	The percent of the storage that is shared with other host systems
Charts	Click icon to see this information in chart (graph) form

6.1.7 Managing performance

With the Manage Performance function, you can collect and display performance information on the storage server

- Number of I/O requests
- Number of bytes transferred
- Read and write response time
- Cache use statistics

The information you collect can help you make informed decisions about volume placement and capacity planning, and isolate I/O performance bottlenecks.

6.1.7.1 Schedule data collection

Procedures for Schedule Data Collection, Creating New Data Collection and modifying a data collection task are the same as described in 6.1.5.1 on page 96. The only difference is that you have to

- Click **Schedule Data Collection** from the selections under **Manage Performance** (see Figure 45 on page 94).

View Reports

After a scheduled data collection date, you can view the following types of reports for performance:

- Disk Utilization
  - Disk Utilization Summary
  - Disk Utilization Report: Device Adapter Level
  - Disk Utilization Report: Disk Group Level
- Disk/Cache Transfer
  - Disk/Cache Transfer Summary Report
  - Disk/Cache Transfer Report: Cluster Level
  - Disk/Cache Transfer Report: Device Adapter Level
  - Disk/Cache Transfer Report: Disk Group Level
- Cache Report
  - Cache Report Summary
  - Cache Report: Cluster Level
  - Cache Report: Device Adapter Level
  - Cache Report: Disk Group Level

Fill in the fields in the Performance Reporting Specifications panel (Figure 55 on page 109) to give the specifications for the types of reports you want to see

1. Identify the storage server for which you are collecting performance data. Use the spin button in the Storage Server (Available date range) field. This field includes the available date range.
2. Enter the beginning date for the performance report in the Show performance report beginning field. The date should be in the mm/dd/yyyy format. For example, *03/03/1999*.
3. Identify the time interval for which you are collecting performance data. Use the spin button in the Show date range by field. You can set the time interval to a day, week, or a quarter.

4. With the spin button, select the type of performance report in the Begin with performance report type field. Your options for report type are:
   - Disk Utilization
   - Cache/Disk Transfer
   - Cache Report
5. Click the desired report type, then click **Show Report.**

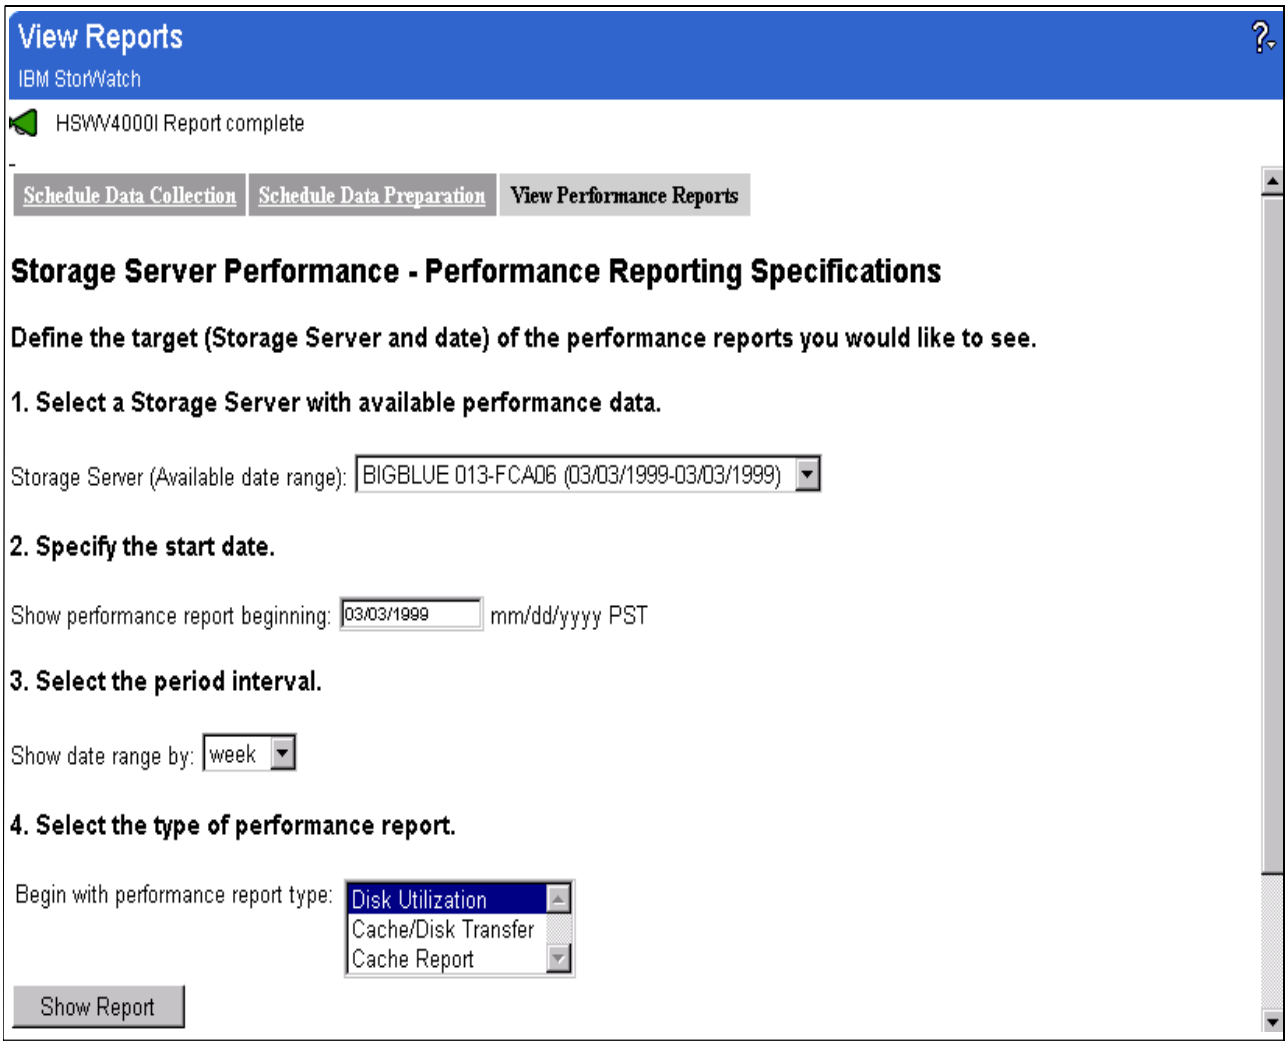

Figure 55. Performance reporting specifications

6.1.7.2 Disk utilization summary

To display the Disk utilization summary reportas shown in Table 12:

1. Click **Disk Utilization** in the Begin with performance report type field on the Performance Reporting Specifications panel (see Figure 55 on page 109).
2. Click **Show Report.**

The Disk Utilization Summary report (see Figure 56 on page 111) provides information on the following areas.

Table 12. Disk utilization summary

Field	Description
Date	The date that data was first detected, and consequently collected by Expert
Cluster	This field identifies the cluster
Disk Utilization (Percent)	The information in the two fields under this umbrella shows hourly disk utilization
Hourly Average	The percentage of hourly average disk utilization
Hourly Maximum	The percentage of hourly maximum disk utilization
Num Device Adapters	The number of device adapters for attached devices (DDMs)
Num Disk Groups	The number of disk (DDM) groups
Num Logical Volumes	The number of logical disk volumes
Config Change	Has the logical configuration changed during period of data collection

You may want more details on Disk Utilization for your analysis:

1. On the Disk utilization summary report (see Figure 56), click the highlighted Date field.

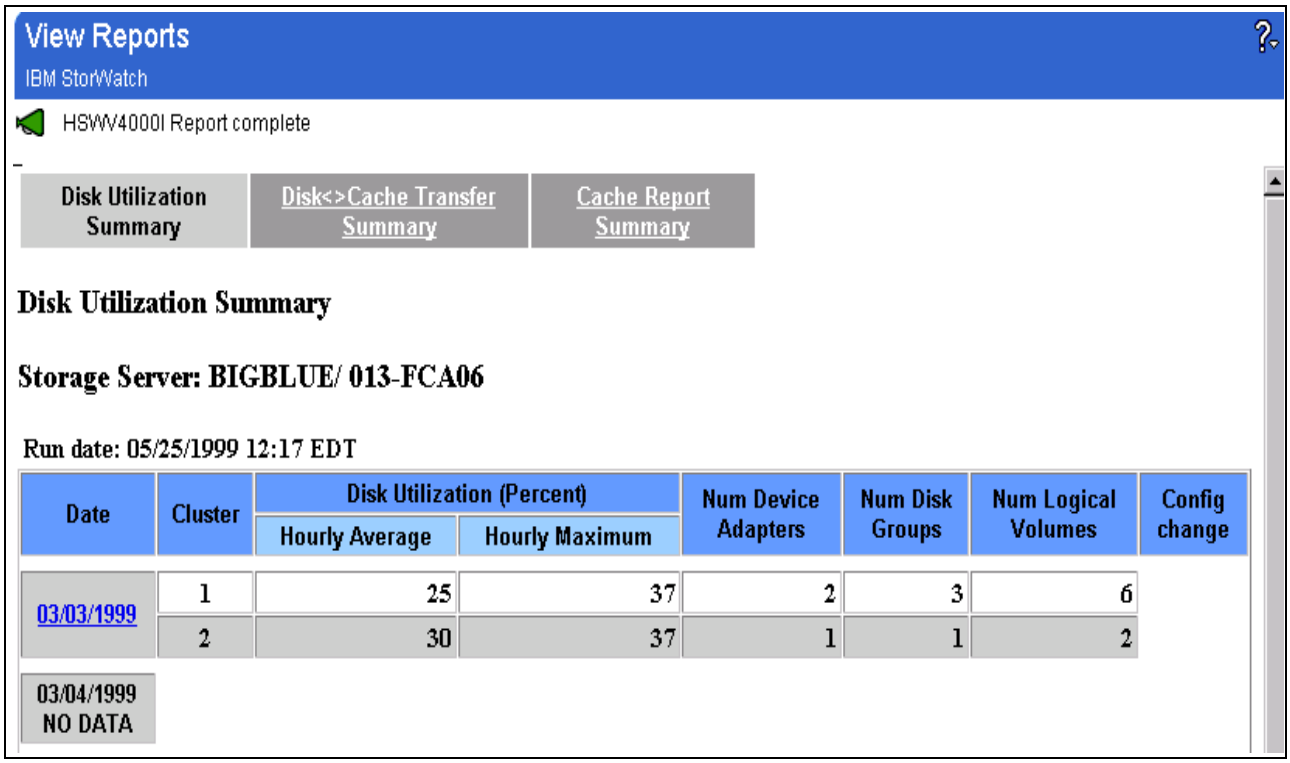

Figure 56. Example of the disk utilization summary report

This links you to the report on Device Adapter Level as shown in Figure 57 on page 112.

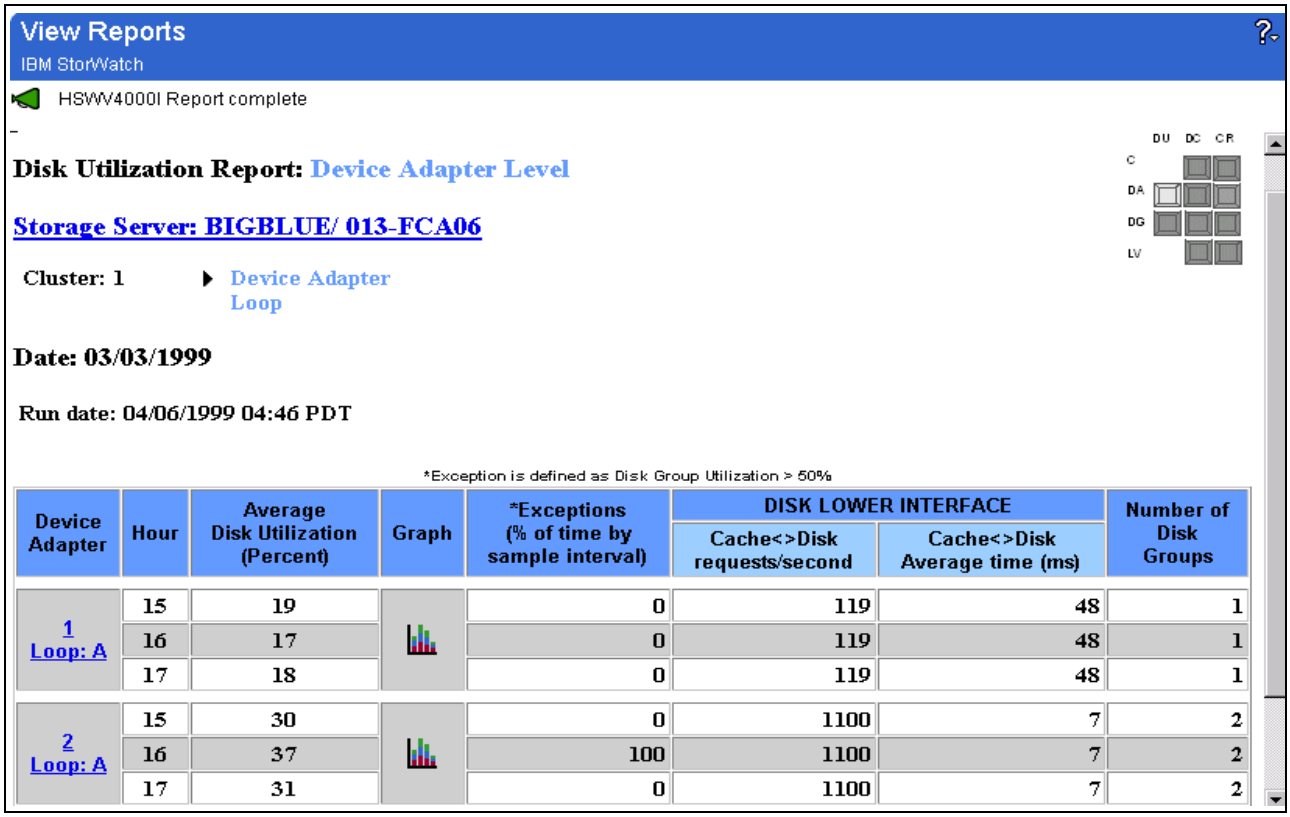

Figure 57. Example of the disk utilization report: Device adapter level

Table 13 describes each field used in the figure.

Table 13. Device adapter detail report for disk utilization

Field	Description
Device Adapter	Device adapter identification
Hour	Time of data collection interval
Average Disk Utilization (Percent)	The percentage of average disk utilization in one hour interval
Graph	Link to a graphic representation of the reports
Exceptions	The percentage of time by sample interval
Disk Lower Interface	The information in the two fields under this umbrella shows Cache<>Disk transfer in both directions
Cache<>Disk requests/second	Transfer rate in both directions in number of IO requests per second
Cache<>Disk Average time (ms)	The average transfer time of IO request for both directions
Number of Disk Groups	The number of disk groups per device adapter

2. From either of these reports, you can click the graph symbol in the Graph column. This links you to a graphical presentation of the reports (see Figure 58 on page 113).

Disk Utilization Report: DEVICE ADAPTER LEVEL

Storage Server: 013-FCA06-> Cluster: 1-> Device Adapter: 1

Date: 03/03/1999

Report Date: 05/25/1999 12:21 EST

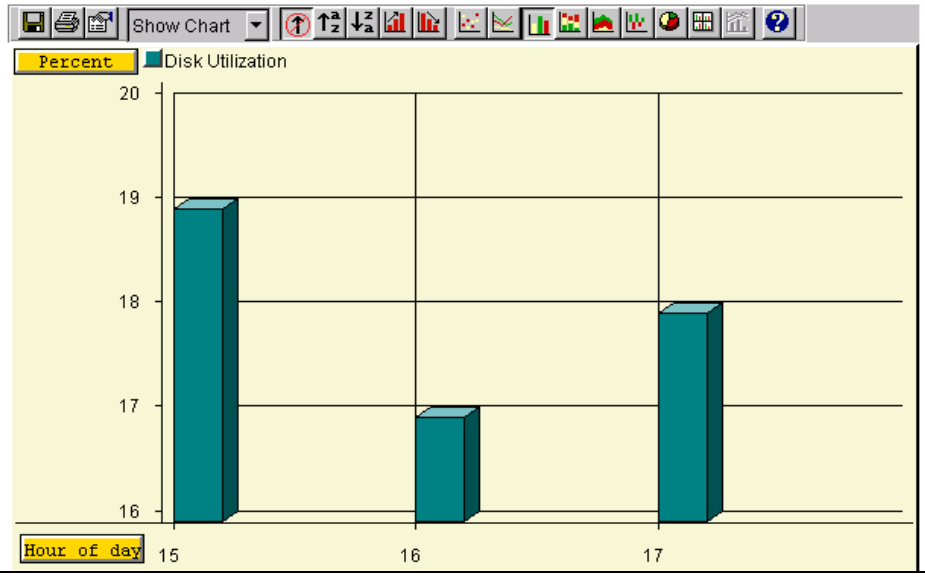

Figure 58. Example of the graph for disk utilization

3. From the Device Adapter Level report, click the highlighted Device Adapter number. This links you to the report on Disk Group Level (see Figure 59 on page 114).

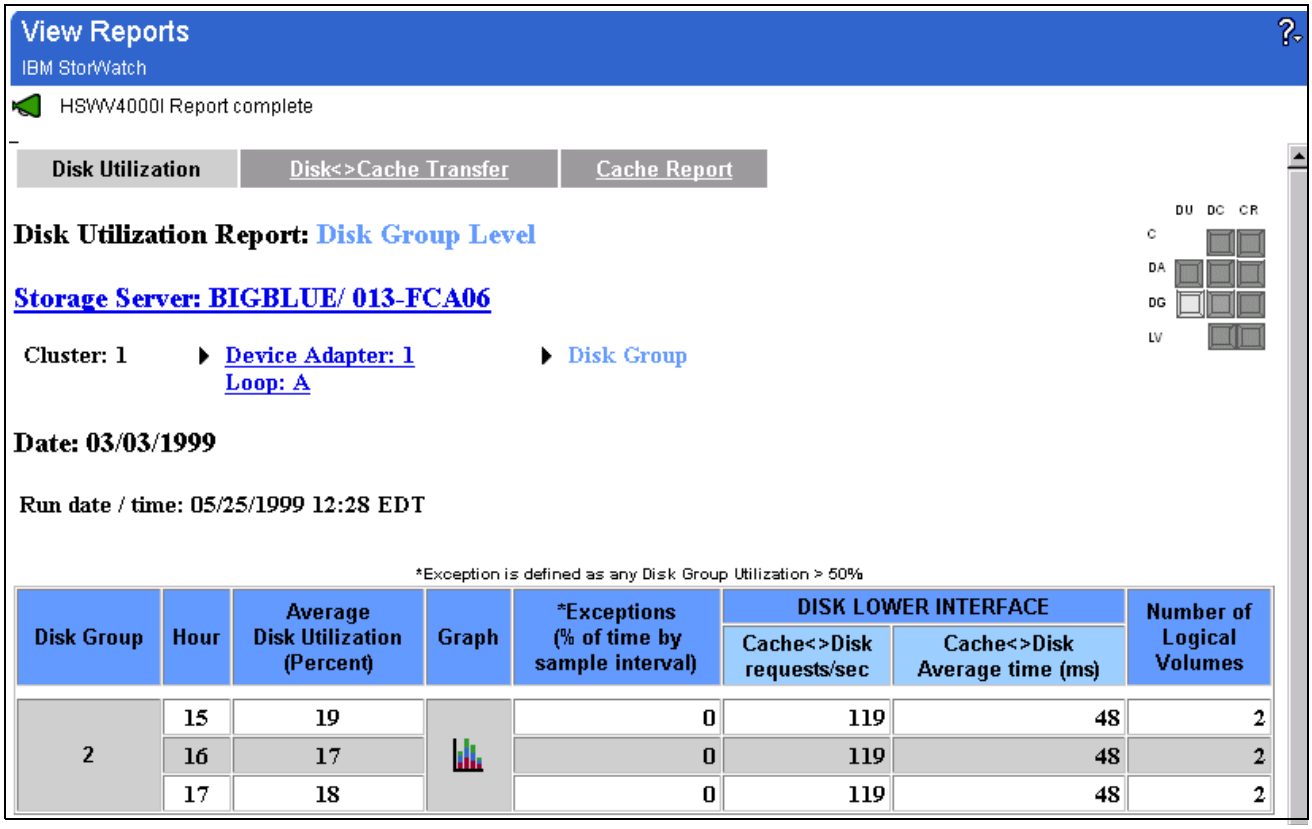

Figure 59. Example of the disk utilization report: Disk group level

Table 14 provides an explanation of the fields used.

Table 14. Disk group detail report

Field	Description
Disk Group	Disk group identification
Hour	Time of data collection interval
Average Disk Utilization (Percent)	The percentage of average disk utilization in one hour interval
Graph	Link to a graphic representation of the reports
Exceptions	The percentage of time by sample interval
Disk Lower Interface	The information in the two fields under this umbrella shows Cache<>Disk transfer in both directions
Cache<>Disk requests/second	Transfer rate in both directions in number of IO requests per second
Cache<>Disk Average time (ms)	The average transfer time of IO request for both directions
Number of Logical Volumes	The number of logical volumes per Disk Group

Analyze disk utilization
You should pay attention to:

- Bottlenecks
- Thresholds that may have been exceeded
- Very high or low percentages of usage

6.1.7.3 Disk/cache and cache/disk transfer summary

To display the disk/cache and cache/disk transfer summary report

1. Click **Cache/Disk Transfer** in the Begin with performance report type field (see Figure 55 on page 109).
2. Click **Show Report.**

The Cache Transfer Summary Report is shown, as seen in Figure 60. Table 15 provides an explanation of the fields used.

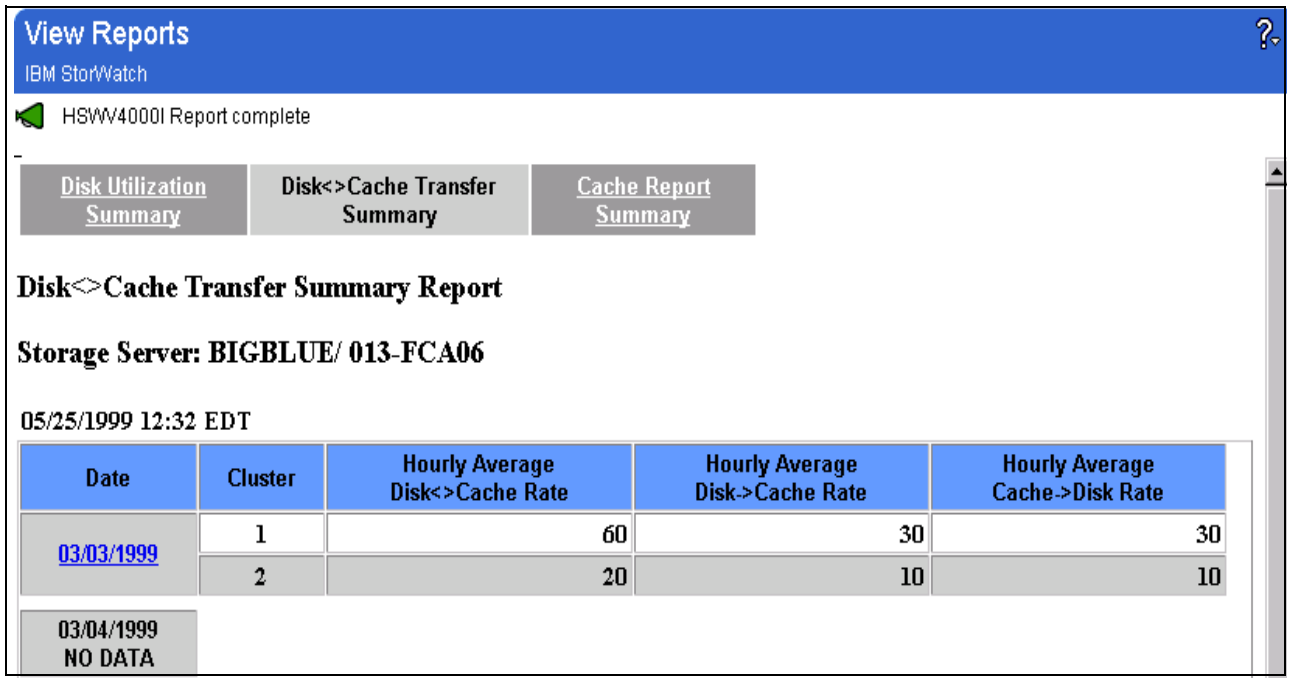

Figure 60. Example of the Disk <> Cache transfer summary report

Table 15 provides an explanation of the fields used.

Table 15. Disk<> Cache transfer summary report

Field	Description
Date	The date that data was first detected, and consequently collected by Expert
Cluster	Cluster identification
Hourly Average Disk<> Cache Rate	Total transfer rate in both directions between disk and cache
Hourly Average Disk-> Cache Rate	Transfer rate from disk to cache

Field	Description
Hourly Average Cache->Disk Rate	Transfer rate from cache to disk

To get more details on disk/cache and cache/disk transfer for your analysis:

1. On the Disk <> Cache Transfer Summary Report (see Figure 60 on page 115), click a highlighted Date field. This links you to the report on that cluster.
2. From the Cluster Level report (see Figure 61 on page 117), click a highlighted Cluster number. This links you to the report on Device Adapter Level.
3. From the Device Adapter Level report (see Figure 62 on page 118), click a highlighted Device Adapter number. This links you to the report on Disk Group Level (see Figure 63 on page 119).
4. From any of these reports, you can click the graph symbol in the Graph column. This links you to a graphic presentation of the reports.

Cluster detail report

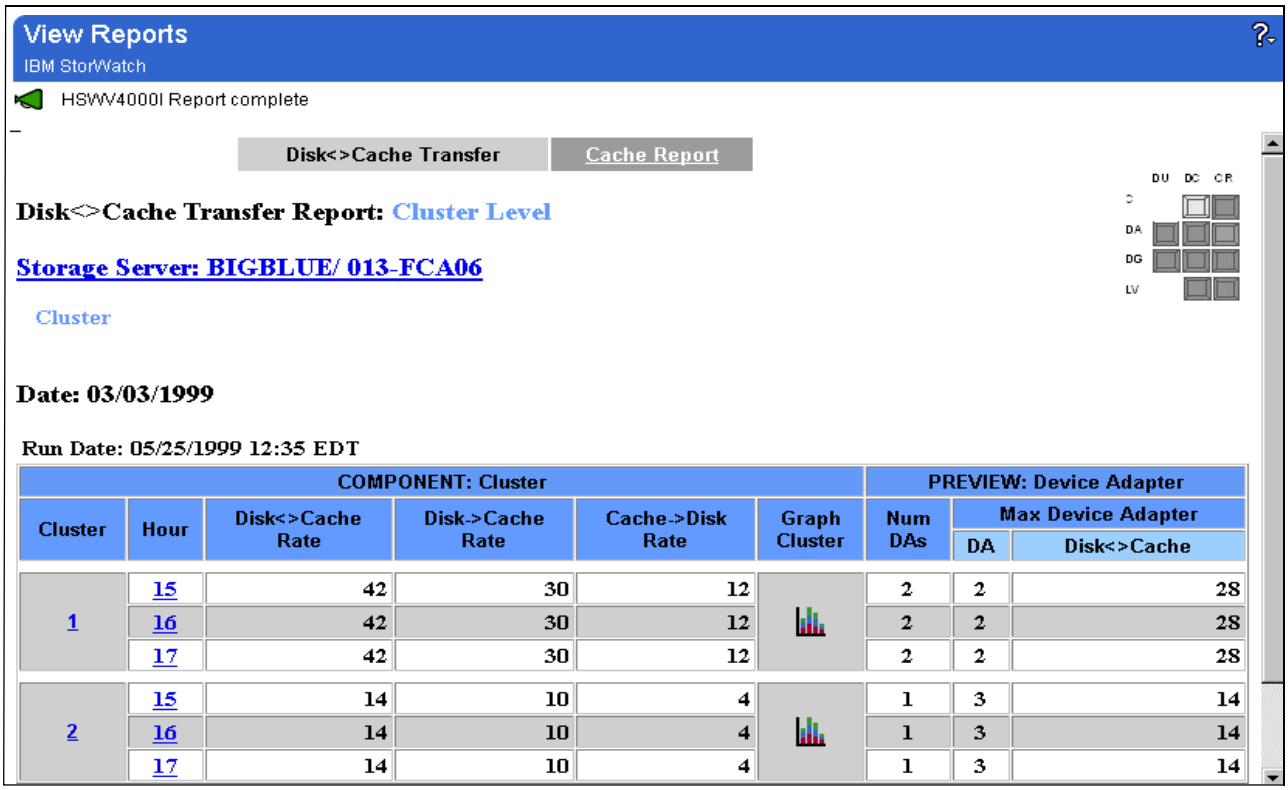

COMPONENT: Cluster						PREVIEW: Device Adapter		
Cluster	Hour	Disk<>Cache Rate	Disk->Cache Rate	Cache->Disk Rate	Graph Cluster	Num DAs	Max Device Adapter	
							DA	Disk<>Cache
1	15	42	30	12		2	2	28
	16	42	30	12		2	2	28
	17	42	30	12		2	2	28
2	15	14	10	4		1	3	14
	16	14	10	4		1	3	14
	17	14	10	4		1	3	14

Figure 61. Example of the Disk<> Cache transfer report: cluster level

Table 16 provides an explanation of the fields used.

Table 16. Disk<> Cache transfer report: cluster level

Field	Description
Cluster	Cluster identification
Hour	Time of data collection interval
Disk<> Cache Rate	Total transfer rate in both directions between disk and cache
Disk->Cache Rate	Transfer rate from disk to cache
Cache->Disk Rate	Transfer rate from cache to disk
Graph Cluster	Link to a graphic representation of the reports
Num DAs	The number of Device Adapters not involved in the highest transfer rate
Max Device Adapter	The information in the two fields under this umbrella shows the highest transfer rate in the interval
DA	The number of Device Adapters involved in the highest transfer rate
Disk<> Cache	The highest transfer rate between disk and cache in both directions per Device Adapters

Device adapter detail report

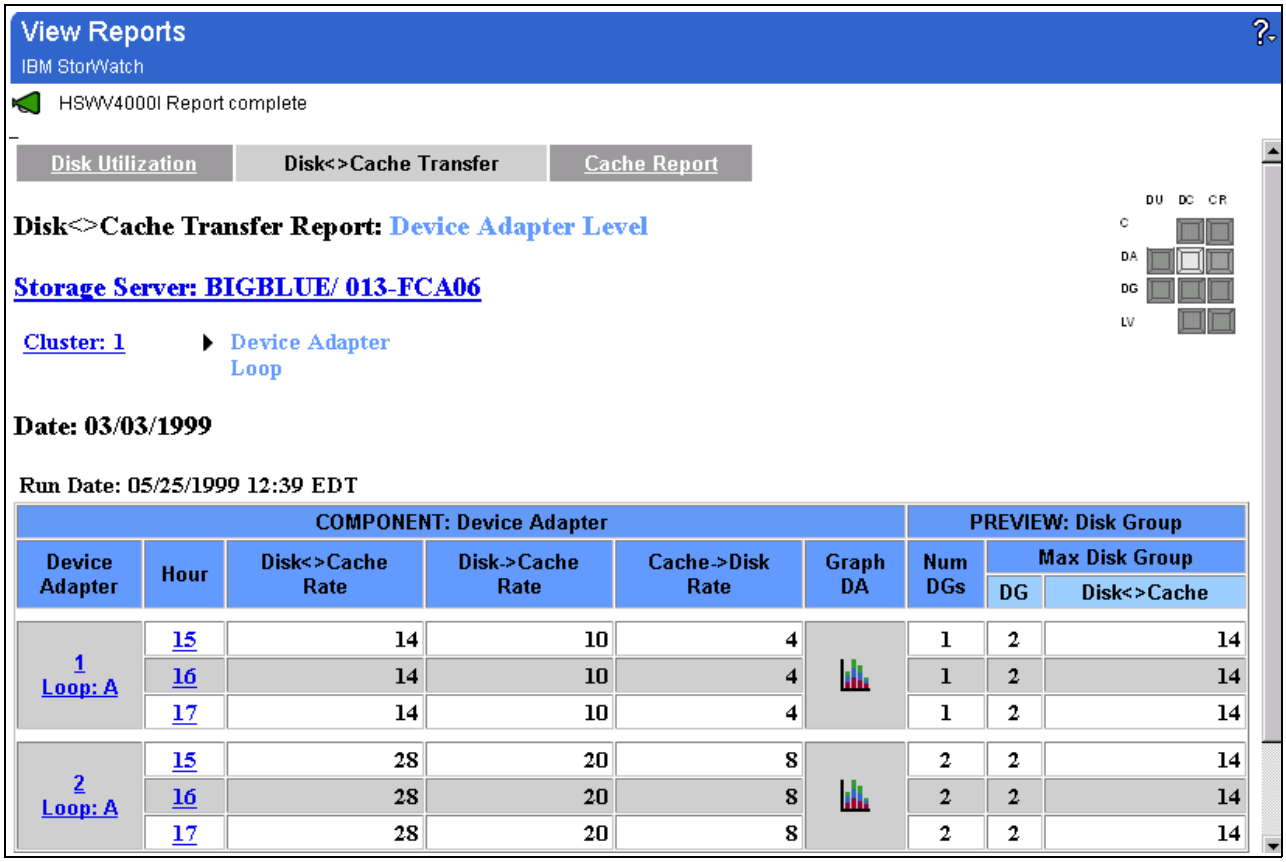

COMPONENT: Device Adapter						PREVIEW: Disk Group		
Device Adapter	Hour	Disk<>Cache Rate	Disk->Cache Rate	Cache->Disk Rate	Graph DA	Num DGs	Max Disk Group	
							DG	Disk<>Cache
1 Loop: A	15	14	10	4		1	2	14
	16	14	10	4		1	2	14
	17	14	10	4		1	2	14
2 Loop: A	15	28	20	8		2	2	14
	16	28	20	8		2	2	14
	17	28	20	8		2	2	14

Figure 62. Example of the Disk <> Cache transfer report: device adapter level

Table 17 provides an explanation of the fields used.

Table 17. Disk <> Cache transfer report: Device adapter level

Field	Description
Device Adapter	Device adapter identification
Hour	Time of data collection interval
Disk<> Cache Rate	Total transfer rate in both directions between disk and cache
Disk->Cache Rate	Transfer rate from disk to cache
Cache->Disk Rate	Transfer rate from cache to disk
Graph DA	Link to a graphic representation of the reports
Num DGs	The number of Device Groups not involved in the highest transfer rate
Max Disk Group	The information in the two fields under this umbrella shows the highest transfer rate in the interval
DG	The number of Disk Groups involved in the highest transfer rate
Disk<> Cache	The highest transfer rate between disk and cache in both directions per Disk Groups

Disk group detail report

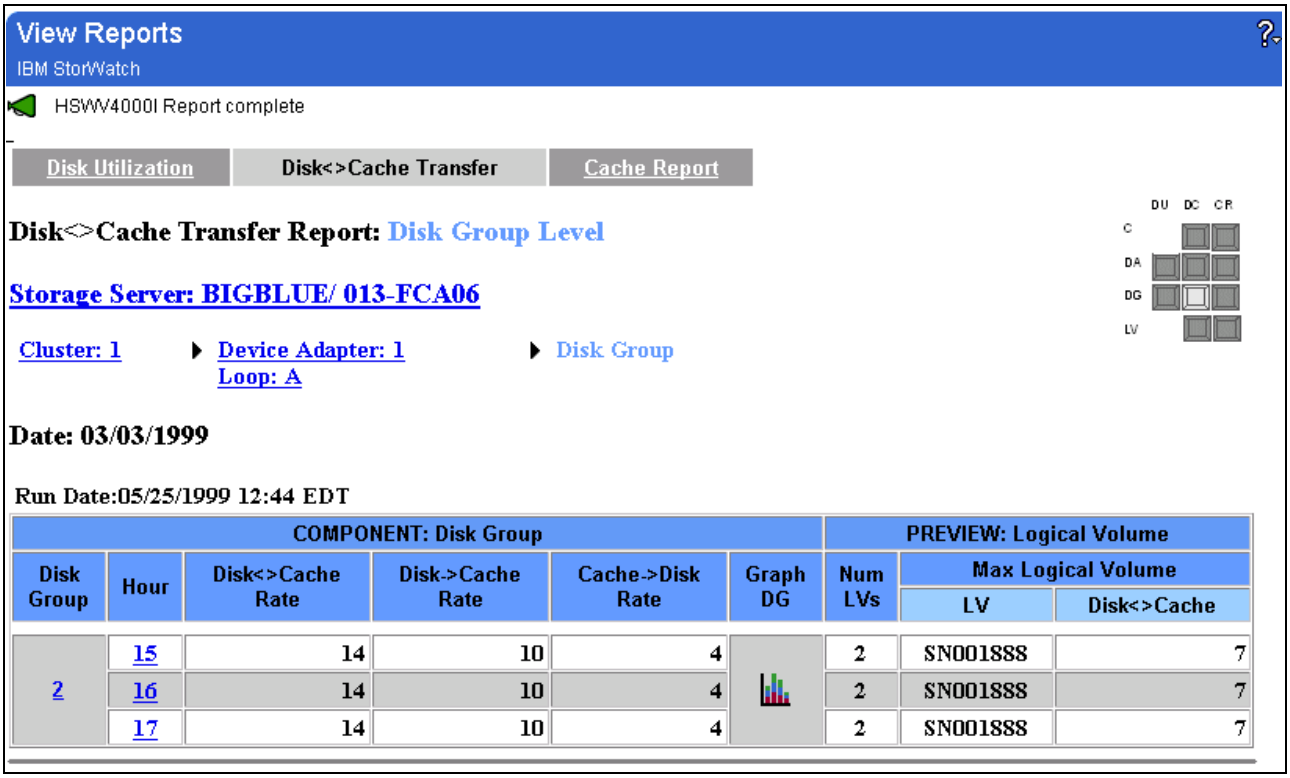
View Reports

IBM StorWatch

HSWV4000I Report complete

Disk Utilization | Disk<>Cache Transfer | Cache Report

Disk<>Cache Transfer Report: Disk Group Level

Storage Server: BIGBLUE/ 013-FCA06

Cluster: 1 ▸ Device Adapter: 1 Loop: A ▸ Disk Group

Date: 03/03/1999

Run Date:05/25/1999 12:44 EDT

COMPONENT: Disk Group						PREVIEW: Logical Volume		
Disk Group	Hour	Disk<>Cache Rate	Disk->Cache Rate	Cache->Disk Rate	Graph DG	Num LVs	Max Logical Volume	
							LV	Disk<>Cache
2	15	14	10	4		2	SN001888	7
	16	14	10	4		2	SN001888	7
	17	14	10	4		2	SN001888	7

Figure 63. Example of the Disk<>Cache transfer report: Disk group level

Table 18 provides an explanation of the fields used.

Table 18. Disk<>Cache transfer report: Disk group level

Field	Description
Disk Group	Disk group identification
Hour	Time of data collection interval
Disk<> Cache Rate	Total transfer rate in both directions between disk and cache
Disk->Cache Rate	Transfer rate from disk to cache
Cache->Disk Rate	Transfer rate from cache to disk
Graph DG	Link to a graphic representation of the reports
Num LVs	The number of logical volumes not involved in the highest transfer rate
Max Logical Volume	The information in the two fields under this umbrella shows the highest transfer rate in the interval
LV	The logical volumes involved in the highest transfer rate
Disk<> Cache	The highest transfer rate between disk and cache in both directions per logical volume

Analyze Disk to Cache Transfer
You should pay attention to:

- Bottlenecks
- Thresholds that may have been exceeded
- Very high or low percentages of usage

6.1.7.4 Cache summary report

To display cache summary report:

1. Click **Cache** in the Begin with performance report type field (see Figure 55 on page 109).
2. Click **Show Report.**

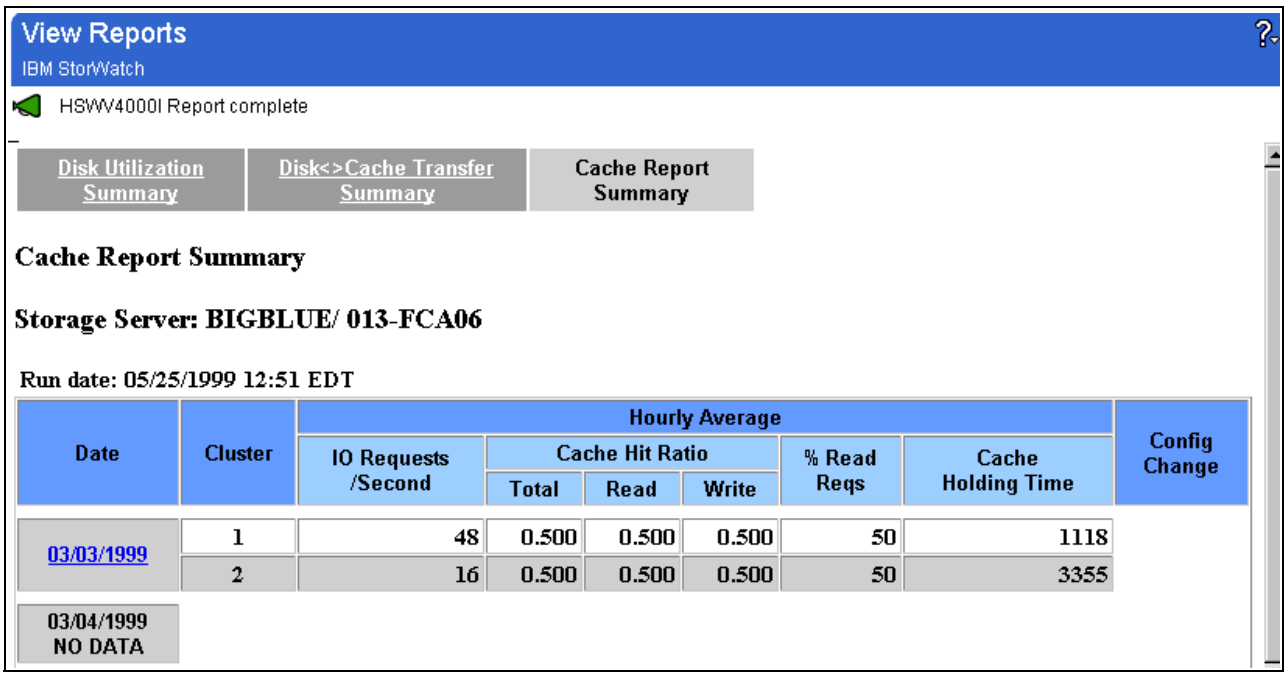

Figure 64. Example of the cache report summary

The Cache report summary provides information on the items displayed in Table 19 on page 120.

Table 19. Cache report summary

Field	Description
Date	The date that data was first detected, and consequently collected by Expert
Cluster	Cluster identification
IO Requests/Second	Hourly average number of IO requests per second
Total	The total number of cache hit ratio
Read	The total cache read hit ratio
Write	The total cache write hit ratio
% Read Reqs	The percentage of read requests

Field	Description
Cache Holding Time	The time data occupies cache
Config Change	Has the logical configuration changed during period of data collection

Below each of the following figures, there is a table that describes the fields used in the figure.

To get more details on cache for your analysis:

1. On the cache report summary (see Figure 64 on page 120), click a highlighted Date field. This links you to the report on cluster level.
2. From the cluster level report (see Figure 65 on page 122), click a highlighted cluster number. This links you to the report on device adapter level.
3. From the device adapter level report (see Figure 66 on page 123), click a highlighted device adapter number. This links you to the report on disk group level.
4. From the disk group level report (see Figure 67 on page 124), click a highlighted Disk Group number. This links you to the report on logical volume level (Figure 68 on page 125).
5. From any of these reports, you can click the graph symbol in the Graph column. This links you to a graphical representation of the reports.

6.1.7.5 Cluster detail report

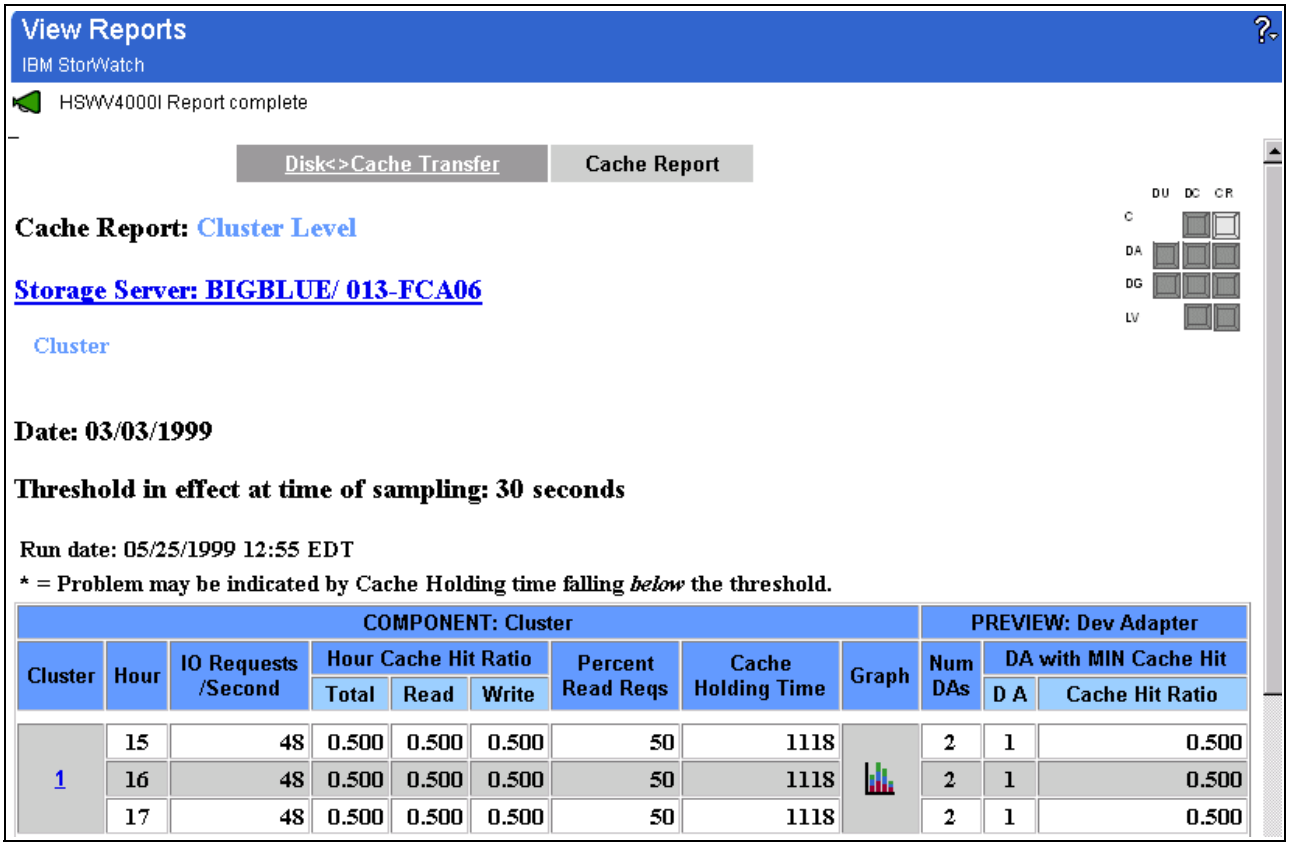

COMPONENT: Cluster									PREVIEW: Dev Adapter		
Cluster	Hour	IO Requests /Second	Hour Cache Hit Ratio			Percent Read Reqs	Cache Holding Time	Graph	Num DAs	DA with MIN Cache Hit	
			Total	Read	Write					D A	Cache Hit Ratio
1	15	48	0.500	0.500	0.500	50	1118		2	1	0.500
	16	48	0.500	0.500	0.500	50	1118		2	1	0.500
	17	48	0.500	0.500	0.500	50	1118		2	1	0.500

Figure 65. Example of the cache report: Cluster level

Table 20. Cache report: Cluster level

Field	Description
Cluster	Cluster identification
Hour	Time of data collection interval
IO requests/second	Hourly average number of IO requests per second
Total	The total number of cache hit ratio
Read	The total cache read hit ratio
Write	The total cache write hit ratio
Percent read reqs	The percentage of read requests
Cache holding time	The time data occupies cache
Graph	Link to a graphic representation of the reports
Num DAs	The number of Device Adapters not involved in the minimum cache hit rate
DA	The number of Device Adapters involved in the minimum cache hit rate
Cache Hit Ratio	The minimum cache hit rate per Device Adapter

6.1.7.6 Device Adapter Detail Report

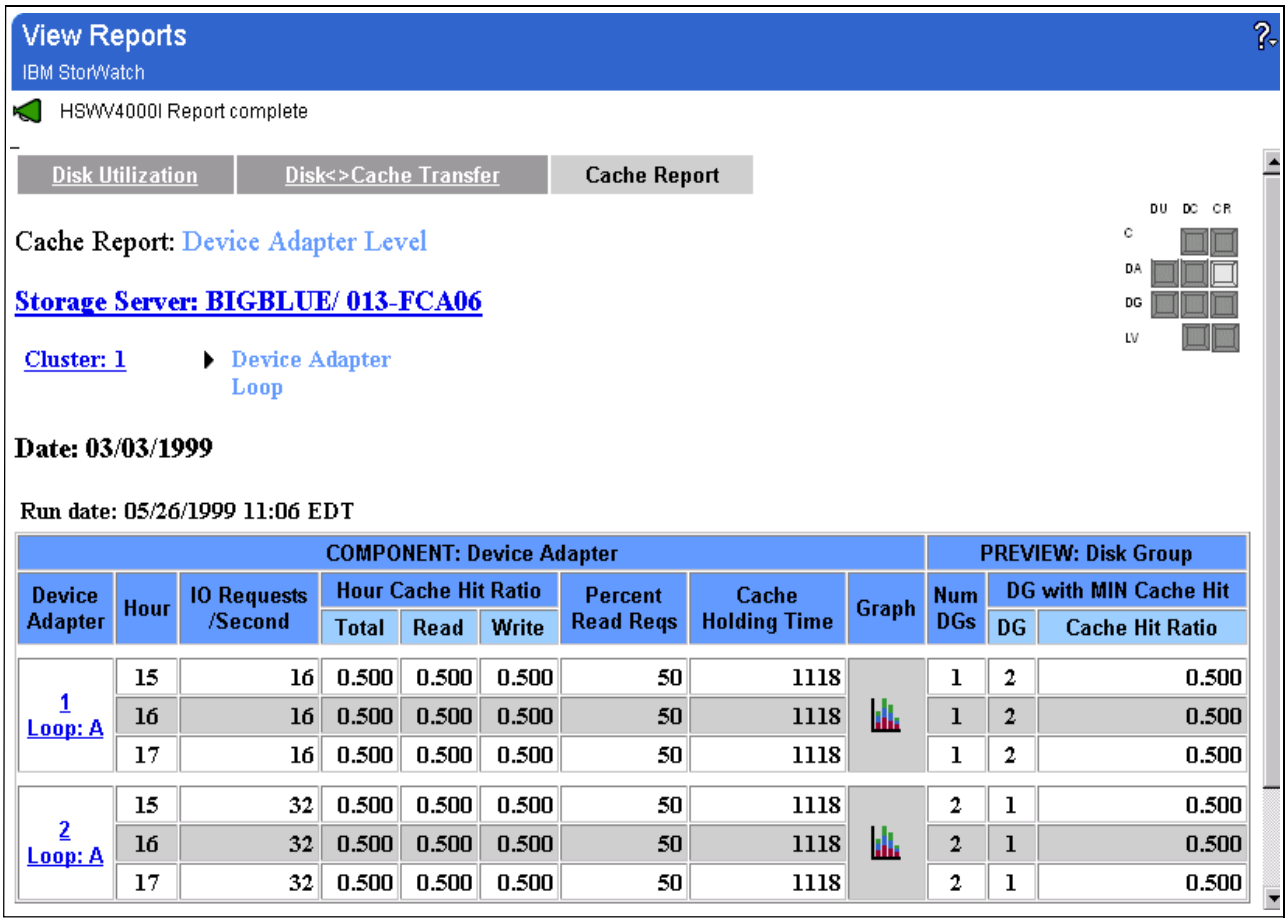

COMPONENT: Device Adapter									PREVIEW: Disk Group		
Device Adapter	Hour	IO Requests /Second	Hour Cache Hit Ratio			Percent Read Reqs	Cache Holding Time	Graph	Num DGs	DG with MIN Cache Hit	
			Total	Read	Write					DG	Cache Hit Ratio
1 Loop: A	15	16	0.500	0.500	0.500	50	1118		1	2	0.500
	16	16	0.500	0.500	0.500	50	1118		1	2	0.500
	17	16	0.500	0.500	0.500	50	1118		1	2	0.500
2 Loop: A	15	32	0.500	0.500	0.500	50	1118		2	1	0.500
	16	32	0.500	0.500	0.500	50	1118		2	1	0.500
	17	32	0.500	0.500	0.500	50	1118		2	1	0.500

Figure 66. Example of the cache report: Device adapter level

Table 21 provides an explanation of the fields used.

Table 21. Cache report: Device adapter level

Field	Description
Device Adapter	The Device Adapter identification
Hour	Time of data collection interval
IO Requests/Second	Hourly average number of IO requests per second
Total	The total number of cache hit ratio
Read	The total cache read hit ratio
Write	The total cache write hit ratio
Percent Read Reqs	The percentage of read requests
Cache Holding Time	The time data occupies cache
Graph	Link to a graphic representation of the reports
Num DGs	The number of disk groups not involved in the minimum cache hit rate
DG	The number of disk groups involved in the minimum cache hit rate
Cache Hit Ratio	The minimum cache hit rate per disk group

6.1.7.7 Disk group detail report

View Reports

IBM StorWatch

HSWV4000I Report complete

Disk Utilization | Disk<>Cache Transfer | Cache Report

Cache Report: Disk Group Level

Storage Server: BIGBLUE/ 013-FCA06

Cluster: 1 ▸ Device Adapter: 1 Loop: A ▸ Disk Group

Date: 03/03/1999

Run date: 05/26/1999 11:20 EDT

COMPONENT: Disk Group									PREVIEW: Logical Volume		
Disk Group	Hour	IO Requests /Second	Hour Cache Hit Ratio			Percent Read Reqs	Cache Holding Time	Graph	Num LVs	LV with MIN Cache Hit	
			Total	Read	Write					LV	Cache Hit Ratio
2	15	16	0.500	0.500	0.500	50	1118		2	SN001877	0.500
	16	16	0.500	0.500	0.500	50	1118		2	SN001877	0.500
	17	16	0.500	0.500	0.500	50	1118		2	SN001877	0.500

Figure 67. Example of the cache report: Disk group level

Table 22 provides an explanation of the fields used.

Table 22. Cache report: Disk group level

Field	Description
Disk Group	The disk group identification
Hour	Time of data collection interval
IO Requests/Second	Hourly average number of IO requests per second
Total	The total number of cache hit ratios
Read	The total cache read hit ratios
Write	The total cache write hit ratios
Percent Read Reqs	The percentage of read requests
Cache Holding Time	The time data occupies cache
Graph	Link to a graphic representation of the reports
Num LVs	The number of logical volumes not involved in the minimum cache hit rate
LV	The logical volume involved in the minimum cache hit rate
Cache Hit Ratio	The minimum cache hit rate per logical volume

6.1.7.8 Logical volume detail report

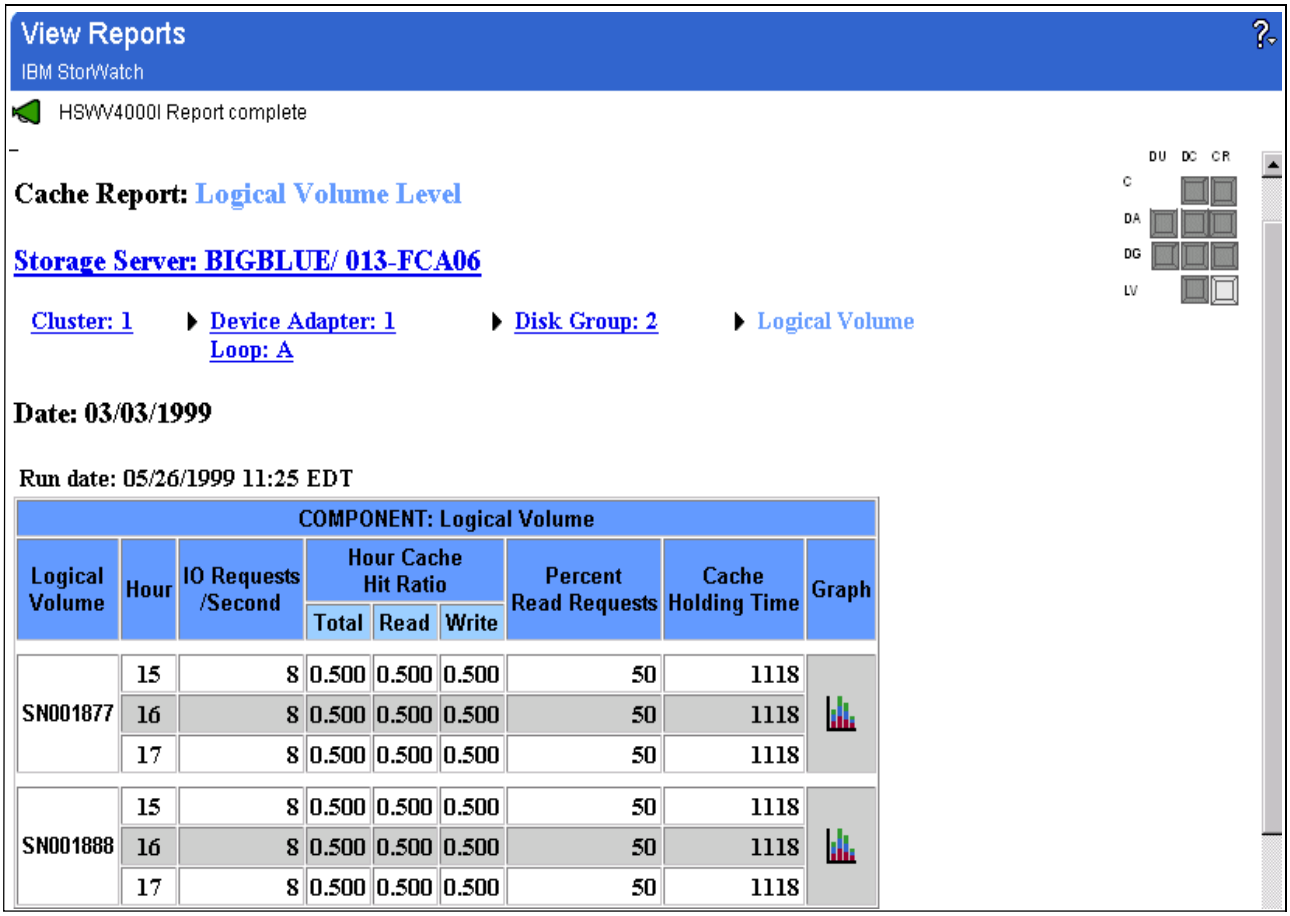

View Reports

IBM StorWatch

HSWV4000I Report complete

Cache Report: Logical Volume Level

Storage Server: BIGBLUE/ 013-FCA06

Cluster: 1 ▸ Device Adapter: 1 Loop: A ▸ Disk Group: 2 ▸ Logical Volume

Date: 03/03/1999

Run date: 05/26/1999 11:25 EDT

COMPONENT: Logical Volume								
Logical Volume	Hour	IO Requests /Second	Hour Cache Hit Ratio			Percent Read Requests	Cache Holding Time	Graph
			Total	Read	Write			
SN001877	15	8	0.500	0.500	0.500	50	1118	
	16	8	0.500	0.500	0.500	50	1118	
	17	8	0.500	0.500	0.500	50	1118	
SN001888	15	8	0.500	0.500	0.500	50	1118	
	16	8	0.500	0.500	0.500	50	1118	
	17	8	0.500	0.500	0.500	50	1118	

Figure 68. Example of the cache report: Logical volume level

Table 23 provides an explanation of the fields used.

Table 23. Cache report: Logical volume level

Field	Description
Logical Volume	The logical volume identification
Hour	Time of data collection interval
IO Requests/Second	Hourly average number of IO requests per second
Total	The total number of cache hit ratios
Read	The total cache read hit ratios
Write	The total cache write hit ratios
Percent Read Requests	The percentage of read requests
Cache Holding Time	The time data occupies cache
Graph	Link to a graphic representation of the reports

Analyze cache report
You should pay attention to:

- Bottlenecks
- Thresholds that may have been exceeded
- Very high or low percentages of usage

Part 2. Implementation in the S/390 environment

This Part discusses the processes for implementing the ESS in an S/390 environment using the ESS Specialist. As part of the logical configuration process, the ESS provides the capability to use configuration worksheets, which can be filled in by you and given to the IBM CE at installation time. This information will be translated to provide you with a logical configuration to match your requirements. The information in this book does not mention this process, but instead details the work necessary to implement the ESS manually.

For more information on the initial configuration process, please refer to Appendix A, “ESS Configuration Planner” on page 253.

Chapter 7. S/390 operating system software support

This chapter describes S/390 Operating Systems Software Support for the ESS.

7.1 General considerations

All S/390 operating systems are capable of utilizing the new disk subsystem - ESS. It is strongly recommended that you check the following general S/390 software items during the pre-installation process:

- Has the QBUCKET information been reviewed for the ESS and checked for software requirements?
- Have the minimum software pre-requisites been met?
- Has all the recommended software maintenance been installed and tested?
- Have the Independent Software Vendors (ISV) been contacted regarding their support of the ESS?
- Have any required ISV fixes been installed and tested?

Figure 69 on page 129 shows a schematic of S/390 software support.

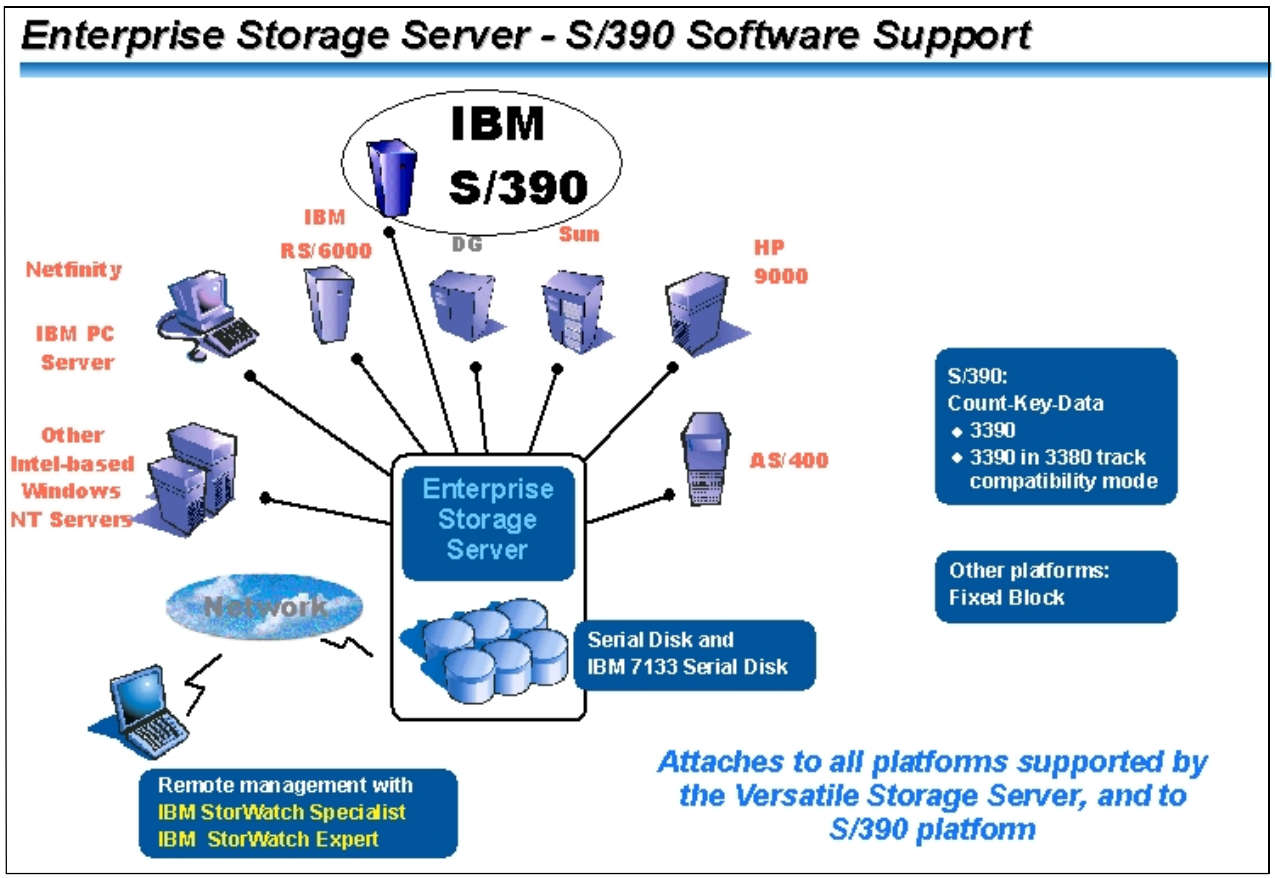

Figure 69. ESS - S/390 software support

7.2 OS/390 support

ESS and all its new functions are supported by the OS/390 operating system. Software support consists of code modifications to DFSMS/MVS functional components and related program products. DFSMS/MVS components include:

- Access Method Services (AMS)
- Device Support
- DFSMSdss
- System Data Mover

In addition to DFSMS/MVS, the following program products are also changed to support ESS:

- DFSMS Optimizer
- DFSORT
- EREP
- ICKDSF
- RMF

DFSMS/MVS software support for ESS is provided as Program Temporary Fixes (PTFs). It is available in these three types:

- Exploitation support
- Toleration support
- Transparency support

7.2.1 Exploitation support

Exploitation support allows an OS/390 system to fully utilize the ESS and all of its new functional capabilities.

7.2.1.1 Prerequisites

- OS/390 Version 2 Release 7 and higher with DFSMS/MVS 1.5.0 and higher

  Supported functions

  - HCD Configure Parallel Access Volumes (PAVs)
  - IOS support for static PAVs
  - Dynamic PAV support
  - DASD UIM support for PAVs
  - Volume level Instant Image
  - DEVSERV Query PAVs
  - Media Manager, Performance CCWs
  - XRC Suspend/Resume, unplanned outage
- OS/390 Version 1 Release 3 and up with DFSMS/MVS 1.3.0 and up

  Supported functions

  - Exploitation Support as previous except Dynamic PAV support
    - Static PAVs only

7.2.2 Toleration support

Toleration support permits you to install an ESS on a non-exploitation level of the OS/390 system. It is therefore possible for a non-exploitation OS/390 system to share the input output definition file (IODF) with other exploitation-level OS/390 systems.

7.2.2.1 Prerequisites

- DFSMS/MVS 1.1 or 1.2

  Supported functions

  - Definition of new 2105 CNTLUNIT TYPE and new PAV device types (Alias or Base)
  - Builds UCBs for base and non-PAV devices only
- RMF 6.2.0 and up

7.2.3 Transparency support

Transparency support enables you to install an ESS on a non-exploitation level of the OS/390 system. However, with transparency support a non-exploitation OS/390 system cannot share IODF with other exploitation OS/390 systems.

7.2.3.1 Prerequisites:

- DFSMS/MVS 1.1 or 1.2 or 1.3 or 1.4

  Supported function

  - Emulation of multiple 3990-6 or 3990-3 Storage Controls with up to 256 unit addresses each
- RMF 5.2.0 and up

7.2.4 OS/390 support levels

As pointed out, DFSMS/MVS and related program products provide toleration, exploitation, and transparency support for ESS. The type of support you install determines the ESS functions you will be able to use. Therefore, you must decide the ESS functions you want to use before you install any of the software support. Make sure that your system release levels are installed with the appropriate DFSMS/MVS software support.

Table 24 on page 132 and Table 25 on page 132 summarizes the support levels for OS/390 software.

Table 24. OS/390 exploitation support

Product	PTF	Integrated
OS/390 - MVS	V1R3-V2R6	V2R7
OS/390 - HCD	V1R3-V2R6	V2R7
OS/390 - RMF	V1R3-V2R6	V2R7
DFSMS	V1R3-V1R4	V1R5
ICKDSF	Rel 16	Not Available
EREP	3.5.0	Not Available
DFSORT	Rel 13	Not Available

Table 25. DFSMS/MVS support

Release	Exploitation Support	Toleration Support	Transparency Support
DFSMS/MVS 1.1	Not Available	PTFs	PTFs
DFSMS/MVS 1.2	Not Available	PTFs	PTFs
DFSMS/MVS 1.3	PTFs	Use Exploitation PTFs	PTFs
DFSMS/MVS 1.4	PTFs	Use Exploitation PTFs	PTFs
DFSMS/MVS 1.5	Integrated	Integrated	Integrated
DFSMS Optimizer	PTFs	Not Applicable	Not Applicable
DFSORT 1.3	PTFs	PTFs	Not Applicable
EREP 3.5	PTFs	PTFs	Not Applicable
ICKDSF 16	PTFs	PTFs	Not Applicable

7.2.5 Other related program products

7.2.5.1 ICKDSF

The ESS supports the following ICKDSF functions:

- AIXVOL
- ANALYZE with SCAN parameter
- BUILDIX
- CONTROL
- CPVOLUME
- INSPECT with NOPRESERVE, NOCHECK, and NOASSIGN parameters
- PPRCOPY
- REFORMAT

ESS, however, does not support ICKDSF ANALYZE (with DRIVETEST parameter), INSTALL, or REVAL.

7.2.5.2 Access method services (AMS)

The LISTDATA command provides new rank counter reports with activity information of a RAID rank. While OS/390 performance monitoring software provides a view of the logical volumes, rank information shows the physical drives activity.

7.2.5.3 EREP

EREP provides problem incident reports also for the new device type 2105.

7.2.5.4 Media manager and asynchronous operations manager

Both components takes advantage of the new performance-enhanced CCWs by limiting extent access to a minimum to increase I/O parallelism.

7.2.5.5 DFSMSdss and DASD error recovery procedures

Both components takes advantage of the new performance enhanced CCWs.

7.3 VM/ESA Support

7.3.1 CP native support

- VM/ESA supports the ESS as an emulation of multiple 3990 Model 6 Storage Controls with up to 256 unit addresses each.
- No native use of new functions

7.3.2 Guest support

- **VM/ESA 2.3.0 with an enabling APAR**
  - Guest use of read/write track CCWs only
- **VM/ESA 2.4.0 with an enabling APAR**
  - Guest use of exploitation functions

7.3.3 VM/ESA 2.4.0 support for OS/390 guests

- Parallel Access Volumes
  - Supported for guest use only
- FlashCopy
  - Supported for guest use only

7.3.4 Utilization of ESS functions

PPRC is supported in VM/ESA and managed by ICKDSF or the StorWatch ESS Specialist. FlashCopy can be managed by ESS Specialist.

Multiple allegiance and I/O Queueing are ESS hardware functions, independent of software support. In a shared environment VM/ESA can take advantage of it. The priority byte, however, is not set by VM/ESA and, therefore, I/O Queueing is not applicable.

7.4 VSE/ESA support

VSE/ESA supports the ESS as an emulation of multiple 3990 Model 6 Storage Controls with up to 256 unit addresses each. Only transparency support is available. New read/write track CCW is not supported.

7.4.1 VSE/ESA support levels

ESS is supported from VSE/ESA 2.1.0 and up. No PTFs are required, however, we recommend you contact your IBM support center.

7.4.2 Utilization of ESS functions

PPRC is supported in VSE/ESA and managed by ICKDSF or the StorWatch ESS Specialist. FlashCopy can be managed by ESS Specialist.

Multiple Allegiance and I/O Queueing are the ESS hardware functions, independent of software support. In a shared environment VSE/ESA can take advantage of it. The priority byte, however, is not set by VSE/ESA and, therefore, I/O Queueing is not applicable.

7.5 TPF support

TPF supports the ESS as an emulation of multiple 3990 Model 3 TPF Storage Controls with up to 256 unit addresses each. ESS supports the TPF Multi Path Locking Facility.

7.5.1 TPF support levels

ESS is supported with **TPF 4.1**. No PTFs are required for Transparency support.

7.5.2 Utilization of ESS functions

With applied corresponding PTFs TPF 4.1 is capable to use the new performance enhanced Read Track/Write Track CCWs.

Multiple Allegiance function is available in TPF environments as an RPQ. TPF takes advantage from Multiple Allegiance and I/O Queueing functions.

Chapter 8. S/390 host setup tasks

This chapter describes the S/390-based software configuration tasks that you will need to perform to successfully implement the ESS and its new functions.

8.1 Preparation and considerations

It is recommended that before commencing these tasks the internal ESS logical setup, or a logical configuration plan, be completed since many of the parameter values required to complete the host definition are dependent on the ESS internal settings.

The ESS Specialist is used to view the ESS internal logical construction and the appropriate address information. See Chapter 5, “Enterprise Storage Server Specialist (ESS Specialist)” on page 69 for an overview of the Specialist and Chapter 9, “ESS configuration for S/390 storage” on page 153 for information on how to define logical devices and subsystems.

It may be an advantage to have network access to the ESS and an ESS Specialist User ID with “View” authority so you can more easily cross check the ESS internal configuration settings with your I/O definitions.

When coding your I/O definitions, you will be defining multiple logical control units (LCU) and Devices internal to the ESS/2105 rather than defining the physical 2105 machine configuration, as is normally done.

Due to the hardware internal switching functions, the enhanced recoverability and internal logical structure within the ESS, you only need to define one 2105/3990 Control Unit cluster rather than the 3990 convention of defining two clusters.

Even though the ESS can support 4096 devices, 16 CU x 256 devices, be aware that an ESCON port can only support 1024 devices, the equivalent of 4 CU x 256, or 8 CU x 128. A list with the Device and Logical Path limits for the ESS can be found in section 2.2.1.2, “Logical features” on page 10.

The definition process

The order of tasks required to successfully define the ESS control units and devices to your host will be as follows:

- Define the Logical Control Unit (LCU) and appropriate type 2105 or 3990 Model 6, Model 3-TPF or Model 3.
- Define the PAV Base devices and/or standard non-PAV 3390/3380 devices, since these are mapped to physical volumes.
- Decide on the number of volumes to have PAV-enabled, and determine the maximum number of Alias per Volume, as specified in ESS settings.
- Subtract the total number of base plus non-PAV volumes from the PAV Starting Address on the LCU, as specified in the ESS settings.
- The remaining device addresses are available to be defined as PAV Alias devices on this LCU and should be equal to or more than the number of Alias required.

After activation of your new IODF you will then be able to initialize the ESS volumes.

Depending on your S/390 operating system and the levels of software, you may be restricted to a subset of available new functions. See Chapter 7, “S/390 operating system software support” on page 129 to review the exploitation levels.

8.2 OS/390 software configuration

This section describes the OS/390 configuration process.

8.2.1 IOCP/HCD

The Multiple Allegiance function requires no specific definitions within OS/390, since it is provided by ESS internal logic.

Warning: In all of the three support modes, Transparency, Toleration and Exploitation, you must make sure the IOCP/HCD control unit CUADDR, device type UNIT and device quantities match, or are within the maximum ranges as set in the logical control unit and volume/device definitions within the ESS.

Mismatches between the logical hardware configuration and device definition in HCD are reported by a DFSMS device support message IEA435I PHYSICAL DEVICE INCONSISTENT WITH LOGICAL DEFINITION. This message is generated during the Vary on-line process for the device.

8.2.1.1 Transparency definitions

The OS/390 I/O subsystem sees the ESS as up to 16 logical IBM 3990 Model 6 or Model 3 subsystems with up to 256 devices per logical subsystem.

Therefore, the IOCP/HCD definitions are as they would be if you were to define a group of “real” IBM 3990 Model 6, or a mix of Model 3 and Model 6, with up to 256 devices, including 3390 and 3380.

No Parallel Access Volume (PAV) device types, 3390B, 3390A, are permitted.

8.2.1.2 Toleration definitions

The OS/390 I/O subsystem recognizes the new ESS control unit type 2105 (3990-6 emulation), which for definition purposes, can be viewed as a logical IBM 3990 Model 6 Control Unit. Up to 16 (control units) CNTLUNIT UNIT=2105 or a mix of 2105, 3990 Model 6 or Model 3, can be defined per ESS.

The new Parallel Access Volume (PAV) device types of Base IODEVICE UNIT=3390B or 3380B and Alias UNIT=3390A or 3380A are also recognized. However, UCBs are ONLY built for the Base 3390B/3380B and non-PAV device types (3390, 3380).

The total number of devices defined (3390, 3380, Base 3390B/3380B and Alias 3390A/3380A) must not exceed 256 per CNTLUNIT statement.

The sample IOCP/HCD statements required to define the ESS in toleration level can be seen in Figure 70 on page 138. However, disregard the Alias device types (3390A) in the example.

8.2.1.3 Exploitation definitions

The OS/390 I/O subsystem recognizes the new ESS control unit type 2105 (native support), which for definition purposes can be viewed as a logical IBM 3990 Model 6 Control Unit. Up to 16 (control units) CNTLUNIT UNIT=2105 or a mix of 2105, 3990 Model 6 or Model 3, can be defined per ESS.

The new Parallel Access Volume (PAV) device types of Base IODEVICE UNIT=3390B or 3380B and Alias UNIT=3390A or 3380A are fully supported. The Alias 3390A/3380A UCBs are built above the 16 MB line. For more information on Static and Dynamic PAV considerations, see 8.2.2, “Parallel access volumes (PAV)” on page 143.

The total number of devices defined (3390, 3380, Base 3390B/3380B and Alias 3390A3380A) must not exceed 256 per CNTLUNIT statement.

Sample IOCDS

Figure 70 on page 138 is a sample IOCDS listing which defines a 2105/3990 logical control unit within an ESS. The devices defined are 16 PAV Base devices 3390B and 112 PAV Alias 3390A.

```
CHPID PATH=(8A),SWITCH=30,TYPE=CNC
CHPID PATH=(9A),SWITCH=70,TYPE=CNC
CHPID PATH=(A0),SWITCH=30,TYPE=CNC
CHPID PATH=(B0),SWITCH=70,TYPE=CNC
CHPID PATH=(C0),SWITCH=30,TYPE=CNC
CHPID PATH=(D0),SWITCH=70,TYPE=CNC
CHPID PATH=(E2),SWITCH=30,TYPE=CNC
CHPID PATH=(E6),SWITCH=70,TYPE=CNC
CNTLUNIT CUNUMBR=1A00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6                *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=0,UNIT=2105
CNTLUNIT CUNUMBR=1A80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1B00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1B80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C3,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1C00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1C80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1D00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=1D80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),              *
      LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),           *
      CUADD=B,UNIT=2105
IODEVICE ADDRESS=(1A00,016),CUNUMBR=(1A00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1A10,112),CUNUMBR=(1A00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1A80,016),UNITADD=00,CUNUMBR=(1A80),             *
      STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1A90,112),UNITADD=10,CUNUMBR=(1A80),             *
      STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1B00,016),CUNUMBR=(1B00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1B10,112),CUNUMBR=(1B00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1B80,016),UNITADD=00,CUNUMBR=(1B80),             *
      STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1B90,112),UNITADD=10,CUNUMBR=(1B80),             *
      STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1C00,016),CUNUMBR=(1C00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1C10,112),CUNUMBR=(1C00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1C80,016),UNITADD=00,CUNUMBR=(1C80),             *
      STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1C90,112),UNITADD=10,CUNUMBR=(1C80),             *
      STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1D00,016),CUNUMBR=(1D00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1D10,112),CUNUMBR=(1D00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1D80,016),UNITADD=00,CUNUMBR=(1D80),             *
      STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1D90,112),UNITADD=10,CUNUMBR=(1D80),             *
      STADET=Y,UNIT=3390A
```

Figure 70. IOCDS PAV example

HCD panels

The following are the HCD panels used to generate the IOCDS shown above. The devices defined are 16 PAV Base devices 3390B and 112 PAV Aliases 3390A, a total of 128 devices on the LCU.

Begin by defining the 2105 logical control unit as shown in Figure 71. The description fields SSID and serial number should be completed as a means of cross referencing. The ESS serial number is shown on the ESS Specialist welcome screen (see Figure 30 on page 78) and the SSID on the ESS Specialist Control Unit Storage panel. See Figure 72 on page 139.

If possible, base your SSID on the ESS serial number.

```
Specify or revise the following values.

Control unit number  . . . . 1A00  +

Control unit type  . . . . . 2105_____  +

Serial number  . . . . . FCA50_____
Description  . . . . . SSID 0501_____________

Connected to switches  . . . 30 30 30 30 70 70 70 70 +
Ports  . . . . . c0 c1 c2 c3 c0 c1 c2 c3 +

If connected to a switch, select whether to have CHPIDs/link
addresses, and unit address range proposed.

Auto-assign  . . . . . 2   1.  Yes
```

Figure 71. HCD panel control unit 2105

The control unit descriptions are shown in the ESS Specialist Control Unit panel, see Figure 72. The CU Image (CUADDR), logical control unit Emulation type (Model 6 or Model 3) and the SSID assigned are detailed. Notice that native 2105 support is shown as 3990-6 emulation.

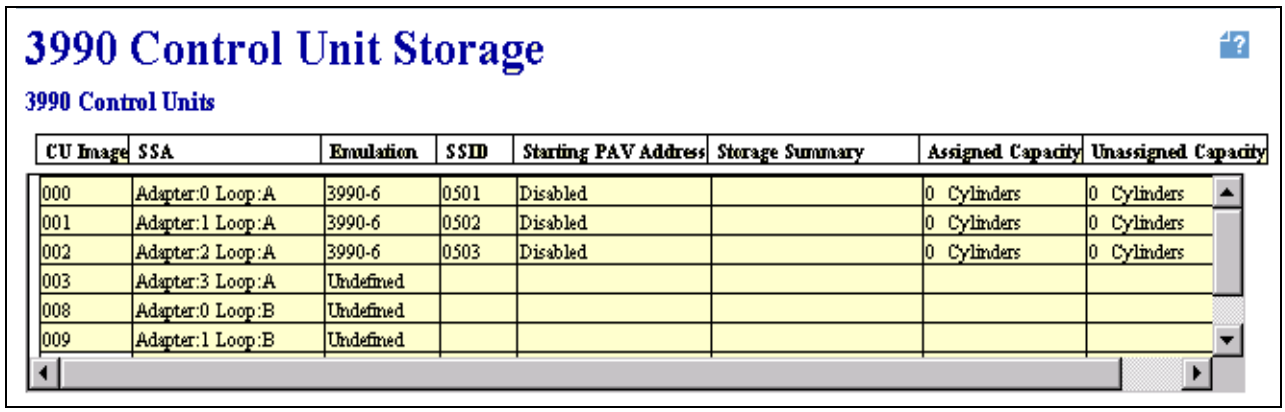

3990 Control Unit Storage

3990 Control Units

CU Image	SSA	Emulation	SSID	Starting PAV Address	Storage Summary	Assigned Capacity	Unassigned Capacity
000	Adapter:0 Loop:A	3990-6	0501	Disabled		0 Cylinders	0 Cylinders
001	Adapter:1 Loop:A	3990-6	0502	Disabled		0 Cylinders	0 Cylinders
002	Adapter:2 Loop:A	3990-6	0503	Disabled		0 Cylinders	0 Cylinders
003	Adapter:3 Loop:A	Undefined					
008	Adapter:0 Loop:B	Undefined					
009	Adapter:1 Loop:B	Undefined					

Figure 72. Specialist 3990 control unit details

The CUADD must match the “CU Image” value specified in the ESS logical definition. See Figure 72 and the HCD panel in Figure 73 on page 140 for details.

```
sssssssssssssssssssss Select Processor / Control Unit sssssssssssssssssss
                                                  Row 33 of 40 More:       >
 Command ===> _____________________________________________ Scroll ===> 0016

 Select processors to change CU/processor parameters, then press Enter.

 Control unit number . . : 1A00     Control unit type . . . : 2105

                Log. Addr. -------Channel Path ID . Link Address + -------
 / Proc. ID Att. (CUADD) +  1---- 2---- 3---- 4---- 5---- 6---- 7---- 8----

 _ SYS9X9        0_         8A.C0 E2.C1 A0.C2 C0.C3 9A.C0 E6.C1 B0.C2 D0.C3
 _ SYS911        __         _____ _____ _____ _____ _____ _____ _____ _____
 _ SYS967        __         _____ _____ _____ _____ _____ _____ _____ _____
```

Figure 73. HCD panel CUADD

The ESS logical control unit definition also has a maximum number of devices field, shown in Figure 74 as the Starting PAV Address with values of either 64, 128 or 256 selectable.

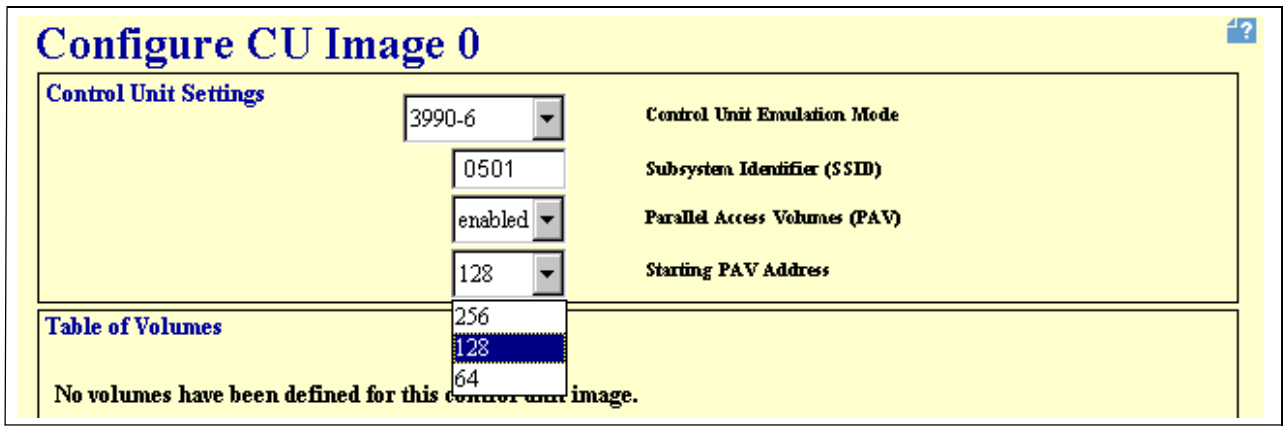

Figure 74. Control unit starting PAV address

Select the HCD device range required, as shown in the HCD panel in Figure 75. If you are defining PAV devices, the device range should equal the Starting PAV Address, in this case 128, as shown in Figure 74.

```
sssssssssssssssssssss Select Processor / Control Unit sssssssssssssssssss
                                                  Row 35 of 40 More: <     >
 Command ===> _____________________________________________ Scroll ===> 0016

 Select processors to change CU/processor parameters, then press Enter.

 Control unit number . . : 1A00     Control unit type . . . : 2105

                --------------Unit Address . Unit Range + --------------
 / Proc. ID Att. 1-------- 2------- 3----- 4----- 5----- 6----- 7----- 8-----
 _ SYS9X9        00.128 ________ _____ ______ ______ ______ ______ ______
 _ SYS911        ________ ________ ______ ______ ______ ______ ______ ______
 _ SYS967        ________ ________ ______ ______ ______ ______ ______ ______
```

Figure 75. HCD panel address range

The number of volumes allocated to the ESS logical control unit and the device types emulated internally within the logical control unit are shown on the ESS Specialist S/390 Storage panel, CU Devices window as in Figure 76. These are “real” base volumes.

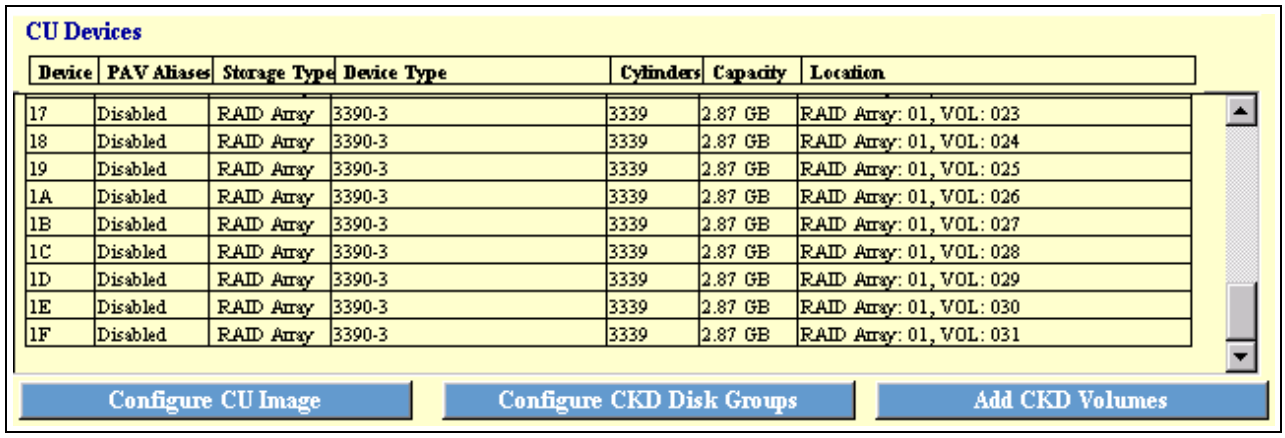

CU Devices

Device	PAV Aliases	Storage Type	Device Type	Cylinders	Capacity	Location
17	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 023
18	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 024
19	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 025
1A	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 026
1B	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 027
1C	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 028
1D	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 029
1E	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 030
1F	Disabled	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01, VOL: 031

Configure CU Image | Configure CKD Disk Groups | Add CKD Volumes

Figure 76. Specialist device selection

Start at the lowest device address available and define your PAV base and non-PAV (standard 3390) devices.

Figure 77 shows the definition of 16 PAV base device type of 3390B. Valid devices are 3390B, 3380B, 3390, 3380.

```
sssssssssssssssssssssssssssssss Add Device sssssssssssssssssssssssssssssss

 Specify or revise the following values.

 Device number  . . . . . . . . 1A00  (0000 - FFFF)
 Number of devices  . . . . . . 16__
 Device type  . . . . . . . . . 3390B_______  +

 Serial number  . . . . . . . . __________
 Description  . . . . . . . . . ________________________________

 Volume serial number . . . . . ______  (for DASD)

 Connected to CUs . . 1A00  ____  ____  ____  ____  ____  ____  ____  +


                                                              Row 1 of 1
 Command ===> ___________________________________________ Scroll ===> CSR

 Select processors to change device/processor definitions, then press
 Enter.

 Device number  . . : 1A00           Number of devices  . : 16
 Device type  . . . : 3390B

                                                Preferred Explicit Device
 / Processor ID   UA +  Time-Out  STADET  CHPID +   Candidate List
 _ SYS9X9         __    No        Yes     __        No
```

Figure 77. HCD panel add base device

The new WLMPAV parameter is shown on the device attributes panel in Figure 78. It should be set to NO for Static PAVs or non PAV devices and YES for dynamic PAVs. See 8.2.2, “Parallel access volumes (PAV)” on page 143 for more detail.

```
Row 1 of 6
Command ===> _____________________________________________ Scroll ===> CSR

Specify or revise the values below.

Configuration ID . : MVS02        DB2 Perf
Device number  . . : 1A00         Number of devices  : 16
Device type  . . . : 3390B

Parameter/
Feature     Value    P Req.  Description
OFFLINE     No               Device considered online or offline at IPL
DYNAMIC     Yes              Device supports dynamic configuration
LOCANY      No               UCB can reside in 31 bit storage
WLMPAV Yes              Device supports work load manager
SHARED      Yes              Device shared with other systems
SHAREDUP    No               Shared when system physically partitioned
```

Figure 78. HCD panel set WLMPAV parameter

The HCD panel in Figure 79 on page 143, details an Alias 3390A definition. Notice that the number of devices defined is 112. The total devices available, or required if using PAVs, is 128 (see Figure 74 on page 140) less 16 Base devices (see the panel in Figure 78). This leaves 112 available for the Alias pool. Also notice that the only device attribute you can set on an Alias is WLMPAV. Valid devices are 3390A or 3380A.

```
sssssssssssssssssssssssssssssss Add Device sssssssssssssssssssssssssssssss

 Specify or revise the following values.

 Device number  . . . . . . . . 1A10  (0000 - FFFF)
 Number of devices  . . . . . . 112_
 Device type  . . . . . . . . . 3390A_______  +

 Serial number  . . . . . . . . _________
 Description  . . . . . . . . . ________________________________

 Volume serial number . . . . . ______  (for DASD)

 Connected to CUs . . 1A00  ____  ____  ____  ____  ____  ____  ____  +


                                                                Row 1 of 1
 Command ===> _____________________________________________ Scroll ===> CSR

 Specify or revise the values below.

 Configuration ID . : MVS02        DB2 Perf
 Device number  . . : 1A10         Number of devices  : 112
 Device type  . . . : 3390A

 Parameter/
 Feature     Value    P Req.  Description
 WLMPAV        Yes Device supports work load manager
```

Figure 79. HCD panel define alias device

8.2.2 Parallel access volumes (PAV)

To invoke the Parallel Access Volume function, you first need to have ordered the appropriate feature on your ESS. This feature enables multiple I/Os to the same volume or device address from the same system image.

You can mix non-PAV device types, 3390 or 3380, with PAV devices, 3390B/3380B or 3390A/3380A, on the same LCU.

Start your PAV and non-PAV device definitions from the lowest address available and work upwards towards the Starting PAV Address. See Figure 74 on page 140.

An Alias device cannot be address 0.

Having defined all your “real” devices, you should specify the rest of the device addresses up to the Starting PAV Address to define Alias devices, 3390A or 3380A. This is required since the ESS allocates PAV Alias devices from the Starting PAV Address specified for the LCU, down towards the base devices.

The maximum number of Alias allowed per volume is set in the ESS Control Unit definitions panel. See Figure 80 on page 144. This value can be set to different values for each volume. This has no effect on the HCD definitions except that enough Alias devices need to be defined to ensure the Alias device pool matches the total PAVs Alias requested for the LCU.

Non-PAV device types do not have access to the Alias device pool.

Figure 80 shows the LCU Devices or Volume table of the Configure CU Image panel used to set the number of PAVs per logical volume.

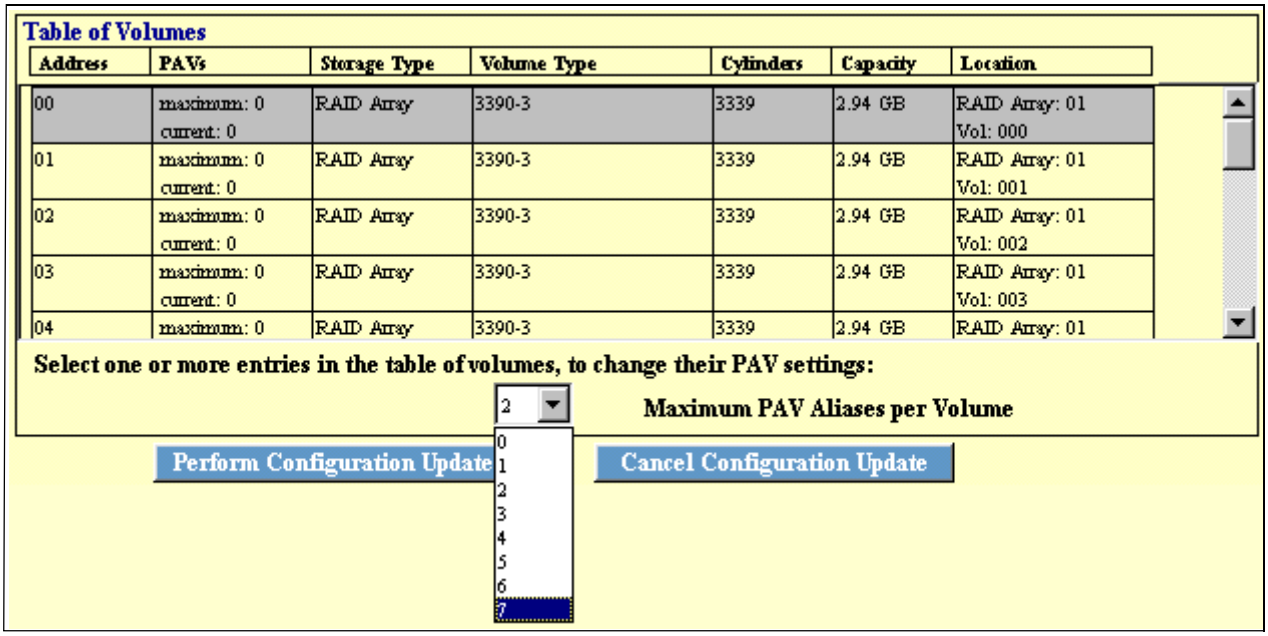

Figure 80. Setting number of PAVs per volume

8.2.2.1 Static PAVs

The static PAV devices are defined as detailed in Figure 80.

PAVs are static if you are running your sysplex Workload Manager (WLM) in compatibility mode. You must be sure to set WLMPAV=NO in your IOCP/HCD device attributes panel. The default value is YES. See examples in Figure 78 on page 142.

8.2.2.2 Dynamic PAV

To exploit the use of Dynamic PAVs, you need to be running your sysplex WLM in GOAL mode.

Dynamic PAV Alias, 3390As, are initially specified statically, as described above, then dynamically reallocated by the WLM.

You need to ensure that WLMPAV=YES is specified on the device attributes panel in HCD and that it is also set for each image in the Sysplex if not sharing IODFs.

Unrequired PAV Aliases on one volume can be reassigned to another Base volume by WLM, over and above the maximum Alias value assigned to the Base volume. Once moved to another base, the Alias will remain associated with that new base until WLM decides to move it.

You should not mix dynamic PAV and non-dynamic PAV hosts within the sysplex, since the non-dynamic PAV hosts ESS I/O activity will not be considered by WLM when it makes its Alias allocation decisions.

Aliases of an offline device will be considered unbound and WLM will select the unbound Alias as the best donor devices. If you have a device offline to one

system and online to others, you should make the device ineligible for dynamic WLM Alias management.

From WLM's Service Definition panel, you can globally enable or disable dynamic PAV Alias tuning. See Figure 81.

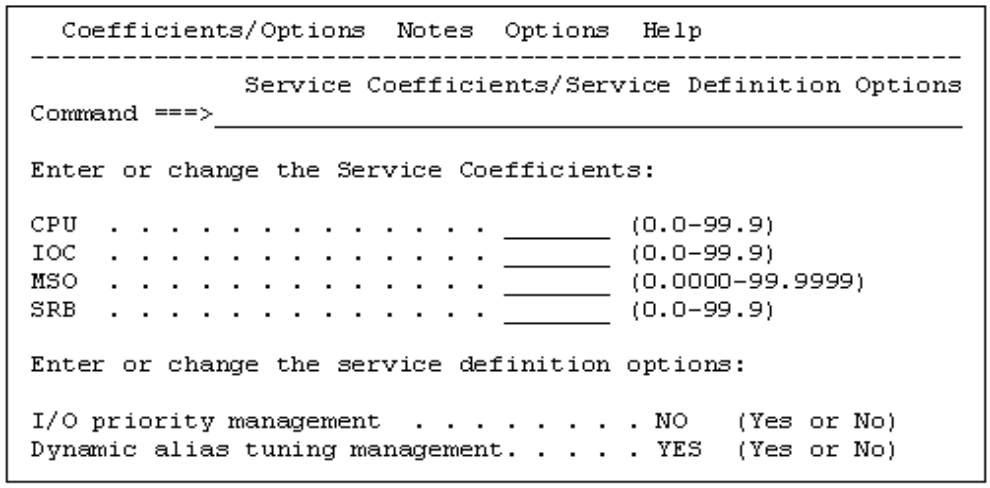

Figure 81. Activate dynamic alias tuning for WLM

8.2.2.3 Displaying the host view of PAVs

After activation of the new IODF, you will receive new base-alias information in response to the `DISPLAY M=DEV and DISPLAY M=CONFIG(xx)` commands. See the MVS System Commands manual for more information.

The `DEVSERV QPAVS` command can be used to display the status information on a device, group of devices or LCU, including PAV devices. In the sample display, Figure 82 on page 146, the UNIT NUM headings are the OS/390 device addresses and the UNIT ADDR heading details the ESS internal device address.

The UA TYPE heading shows BASE, ALIAS, or NON-PAV as the device type.

The STATUS heading contains information to highlight an invalid alias for example INV-ALIAS.

Note that the base and alias affiliations are also displayed.

The `QPAVS` command at the LCU level for the sample definitions shown previously is `DS QPAVS, SSID=0501`.

Figure 82 on page 146 shows a sample `DEVSERV QPAV` command and resulting display.

```
DS QPAV, D222,VOLUME

IEE459I 08:20:32 DEVSERV QPATHS 591
Host                         Subsystem
Configuration                Configuration
-------------                -------------
UNIT UNIT UA                      UNIT
NUM   UA TYPE       STATUS SSID ADDR TYPE
----  -- ----       ------ ---- ---- -------
D222  22 BASE              0102  22  BASE
D2FE  FE ALIAS-D222        0102  FE  ALIAS-22
D2FF  FF ALIAS-D222        0102  FF  ALIAS-22
*** 3 DEVICE(S) MET THE SELECTION CRITERIA
```

Figure 82. Devserv QPAV command

8.2.3 Custom volumes

Custom volumes are defined as 3380 or 3390 format devices types, usually with a non-standard number of cylinders. Therefore, there are no special parameters to set when defining them to the I/O Subsystem. Prior to using the volume, you must first format the volume as with all others via ICKDSF minimal init. A sample job is included in Figure 84 on page 147.

The size of the volume is set via the Add Volumes to CU panels in the ESS Specialist. Figure 83 shows a 3390 Flexvolume of three cylinders.

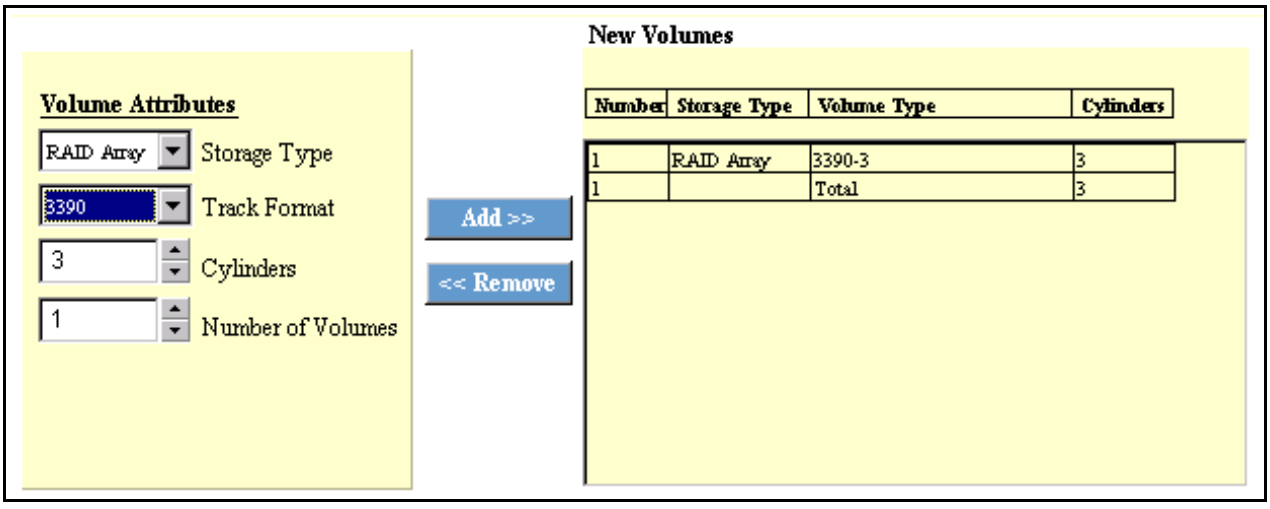

Figure 83. Flexvolume definition panel

8.2.4 Parmlib

8.2.4.1 IECIOSxx

MIH time is recommended as 30 seconds for IBM Enterprise Storage Server.

IOS will automatically set the MIH time to the value specified in the device's self description data, so there is no need to set the MIH value in this parmlib member. If any value is defined, this will override the system's determined value.

If you need to set a value, ensure you set the MIH value for all images in your Sysplex that have access to the ESS.

The MIH setting should appear as follows:

MIH=(1A00-1A5F),TIME=30

The Operator command to achieve the same is:

```
SETIOS MIH=(1A00-1A5F),TIME=30
```

8.2.5 DFSMS

There are no new or unique parameters that need to be set in DFSMS while installing the ESS, apart from including the ESS volumes into existing storage group constructs.

You may wish to consider allocating the Volumes within the ESS to a new storage group, with associated ACS routines, to take advantage of the higher performance functions available to high activity volumes (multiple allegiance or PAVs).

DFSMSdss will now invoke the FlashCopy function as described in section 8.2.7.1, “FlashCopy” on page 147.

The DFSMS Media Manager will vary the Define Extent Range processing for DB2, IMS Fast Path, VSAM, and PDSE to optimize the Multiple Allegiance and PAV performance.

8.2.6 Volume initialization

After activation of the IODF the volumes first need to be initialized with an ICKDSF minimal init. A sample job is included in Figure 84 below.

```
//INITVOL1 JOB account info...
//*-------------------------------------------------------------*
//*   THIS JOB WILL INIT A VOLUME IN AN OFFLINE MODE            *
//*   THIS IS A MINIMAL INIT                                    *
//*-------------------------------------------------------------*
//MININIT  EXEC PGM=ICKDSF,REGION=2M
//SYSPRINT DD   SYSOUT=*
//SYSIN    DD   *
 INIT UNITADDRESS(D123) NOVERIFY VOLID(MVS203) OWNER(SYSPROG) -
      VTOC(0,1,29) INDEX(2,0,15) NOVALIDATE NOCHECK
//
```

Figure 84. ICKDSF minimal init

8.2.7 Copy services

8.2.7.1 FlashCopy

To invoke the FlashCopy function, you first need to have ordered the appropriate feature (FC 183x) on your ESS. This feature enables the volume to volume T0 copy function for a source and target volume within the same Logical Control Unit.

The ESS FlashCopy function is invoked automatically by the use of the standard `DFSMSdss COPY FULL` command. The only consideration is the use of the

COPYVOLID parameter. Without COPYVOLID, the target volume label is retained, VTOC index is rebuilt and the source VVDS is copied to the target.

Naturally if the source and target volumes have the same label, the target volume will go offline, but can be set online to another system image.

DFSMSdss will do a normal copy of a source volume if the source volume has a FlashCopy relationship with another volume, and the copy is still in progress.

A sample job stream is shown in Figure 85.

```
//COPYFULL JOB . . . . .
//*
//INSTIMG  EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=V,OUTLIM=3000
//SRCVOL   DD
//TRGVOL   DD
//SYSIN    DD *
COPY FULL INDYNAM ((SRCVOL)) OUTDYNAM ((TRGVOL)) COPYVOLID
/*
```

Figure 85. DFSMSdss COPY FULL job

8.2.7.2 PPRC and XRC

The PPRC and XRC functions are defined and operated in the same way as is currently done with other IBM Storage Devices. There are some additional features to consider, for example, XRC support of the suspend/resume commands. For more details on PPRC and XRC, refer to Chapter 12, “ESS Specialist Copy Services for S/390” on page 195.

To invoke the PPRC or XRC function, you need to have ordered the appropriate features (XRC FC 181x, PPRC FC 182x) on your ESS.

You can now also define and operate PPRC and FlashCopy functions via the ESS Specialist Web browser interface.

8.2.7.3 Concurrent copy

The concurrent copy function is defined and operated in the same way as is currently done with the IBM 3990 Model 6.

There is no separate feature to order for enabling Concurrent Copy. However, if you have the FlashCopy feature installed and your source and target volumes are within the same Control Unit, FlashCopy will be automatically invoked.

8.3 VM/ESA software configuration

VM/ESA operates in Transparency mode with Guest-only support of some exploitation functions. VM/ESA therefore sees the ESS as a group of up to 16 IBM 3990 Model 6 with up to 256 devices per logical subsystem.

Therefore, the IOCP definitions are as they would be if you were to define a group of “real” IBM 3990 Model 6 with up to 256 devices including 3390 and 3380.

Guests must also define the ESS internal logical control units as IBM 3990 Model 6 and not native 2105.

The Multiple Allegiance function requires no specific definitions within VM/ESA as it is provided by ESS internal logic.

Caution: You must ensure the IOCP control unit CUADDR, device type UNIT and device quantities match, or are within the maximum ranges as set in the Logical Control Unit and Volume/Device definitions within the ESS.

A sample IOCP deck is shown in Figure 70 on page 138.

Mismatched definitions between the ESS internal logical configuration and VM/ESA can be displayed via the `QUERY PAV` command.

8.3.1 Parallel access volumes (PAV)

PAV support is available for guests only and the PAV volumes, Base 3390B and Alias 3390A, must be dedicated to a guest. For an explanation on how to code PAVs, see 8.2.2, “Parallel access volumes (PAV)” on page 143.

A PAV Base (3390B or 3380B) and non-PAV devices (3390 or 3380) can be attached to the SYSTEM, but Alias devices (3390A or 3380A) can only be attached to a guest.

If a PAV Base device is attached to the SYSTEM, then the Alias pool of devices cannot be attached to a guest.

If a PAV Base device is attached to a guest, then the Alias pool of devices must also be attached to the same guest.

8.3.1.1 Displaying the host view of PAVs

The `VM/ESA QUERY PAV` command can be used to display the base and alias affiliation details for PAV devices. This command is available for class B authorized users.

The command is in the format `QUERY PAV,ALL or QUERY PAV,rdev`.

An example of a response to the command issued against a base device is:

```
Device 01A2 is a base Parallel Access Volume device with the
following aliases: 01FE 01FF
```

An example of a response to the command issued against an Alias device is:

```
Device 01F3 is an alias Parallel Access Volume device
whose base device is: 01A4
```

8.3.2 Copy services

To invoke the PPRC or FlashCopy function, you need to have ordered the appropriate features (PPRC FC 182x, FlashCopy FC 183x) on your ESS.

The PPRC functions are defined and operated in the same way as is currently done with ICKDSF.

VM/ESA can take advantage of both FlashCopy and PPRC functions within the ESS by defining and initiating the process from the ESS Specialist Web browser interface.

Native FlashCopy is available to guests only and, therefore, the source and target volumes, dedicated full-pack minidisks, must both be ATTACHed to the VM/ESA guest. The guest invokes the function, not VM/ESA.

8.4 VSE/ESA software configuration

VSE/ESA operates in Transparency mode only and, therefore, sees the ESS as a group of up to 16 IBM 3990 Model 6 with up to 256 devices per logical subsystem.

Therefore, the I/O definitions are as they would be if you were to define a group of “real” IBM 3990 Model 6 with up to 256 devices including 3390 and 3380.

The Multiple Allegiance function requires no specific definitions within VSE/ESA as it is provided by ESS internal logic.

Caution: You must ensure the I/O definitions for your control unit CUADDR, device type UNIT and device quantities match, or are within the maximum ranges as set in the Logical Control Unit and Volume/Device definitions within the ESS.

8.4.1 Copy services

To invoke the PPRC or FlashCopy function, you need to have ordered the appropriate features (PPRC FC 182x, FlashCopy FC 183x) on your ESS.

The PPRC functions are defined and operated in the same way as is currently done with ICKDSF.

VSE/ESA can take advantage of both FlashCopy and PPRC functions within the ESS by defining and initiating the process from the ESS Specialist web browser interface.

8.5 TPF software configuration

TPF operates in Transparency mode (discounting enhanced CCW support) and therefore sees the ESS as a group of up to 16 IBM 3990 Model 3 TPF control units with up to 256 devices per logical subsystem.

Therefore, the I/O definitions are as they would be if you were to define a group of “real” IBM 3990 Model 3 TPF with up to 256 devices.

The Multipath locking facility is supported as on IBM 3990 control units for TPF environments.

The Multiple Allegiance function requires no specific definitions within TPF as it is provided by ESS internal logic. However, it is only available as an RPQ.

Caution: You must ensure the I/O definitions for your control unit CUADDR, device type UNIT and device quantities match, or are within the maximum ranges as set in the Logical Control Unit and Volume/Device definitions within the ESS.

8.5.1 Copy Services

To invoke the PPRC or FlashCopy function, you need to have the appropriate features (PPRC FC 182x, FlashCopy FC 183x) on your ESS.

TPF can take advantage of both FlashCopy and PPRC functions within the ESS by defining and initiating the process from the ESS Specialist Web browser interface.

Chapter 9. ESS configuration for S/390 storage

9.1 Introduction

Here we describe the procedures using the ESS Specialist to configure the ESS for use with S/390 hosts . The ESS Specialist may be used to do the complete storage configuration process for an ESS, defining S/390 and Fixed Block (FB) storage, or it may be used to modify an existing configuration. For example, ESS Specialist might be used to define hosts and add volumes for one of the standard array configurations that can be ordered for a new ESS.

9.1.1 ESS Specialist

The ESS Specialist is a software component that runs in the ESS clusters and provides the interface for the user to define the logical configuration of the ESS. User communication with ESS Specialist is achieved using a Web browser such as Netscape or Internet Explorer, running on a Web client. The Web client is a PC or RS6000 machine connected to one or more ESS subsystems using the TCP/IP protocol through the 10baseT Ethernet LAN adapter in each ESS cluster.

Refer to 3.8.1.1, “ESS Specialist Web client requirements” on page 32.

Also refer to Chapter 5, “Enterprise Storage Server Specialist (ESS Specialist)” on page 69 for an introduction to the ESS Specialist.

Other useful documents include:

- *The IBM Enterprise Storage Server Web Interface User’s Guide*, SC26-7346

9.1.2 S/390 LCUs

Here, we summarize the attributes of S/390 Logical Control Units (LCUs):

- The Enterprise Storage Server (ESS) can emulate up to 16 S/390 LCUs.
- Each LCU corresponds to one ESS Logical Subsystem (LSS).
- An LCU can be defined as 3990 Model 3, 3990 Model 3 TPF, or 3990 Model 6.
- Unlike a FB LSS, an LCU must be defined to the ESS before logical volumes are assigned, and to any S/390 hosts to enable host communication with those volumes.
- There can be two S/390 LCUs defined per SSA loop, one for each DA on the loop. The two DAs on a loop are in different storage clusters.
- Even numbered LCUs are managed by storage cluster 1, odd numbered LCUs are managed by storage cluster 2.
- Every installed ESCON host adapter port can address all configured S/390 LCUs.
- An LCU can have up to 256 logical volumes allocated, one per device address. The ESS can have up to 16 * 256 = 4096 S/390 logical volumes.
- If Parallel Access Volumes (PAVs) are enabled for the LCU, each PAV alias takes up one of the 256 device addresses.

9.1.2.1 Relationship between LCUs and device adapters

Figure 86 shows the correspondence between the S/390 LCUs, and the DAs and SSA loops for a single frame configuration. For a configuration which has more than 128 drives, LCU numbers 008 through to 00F would be added.

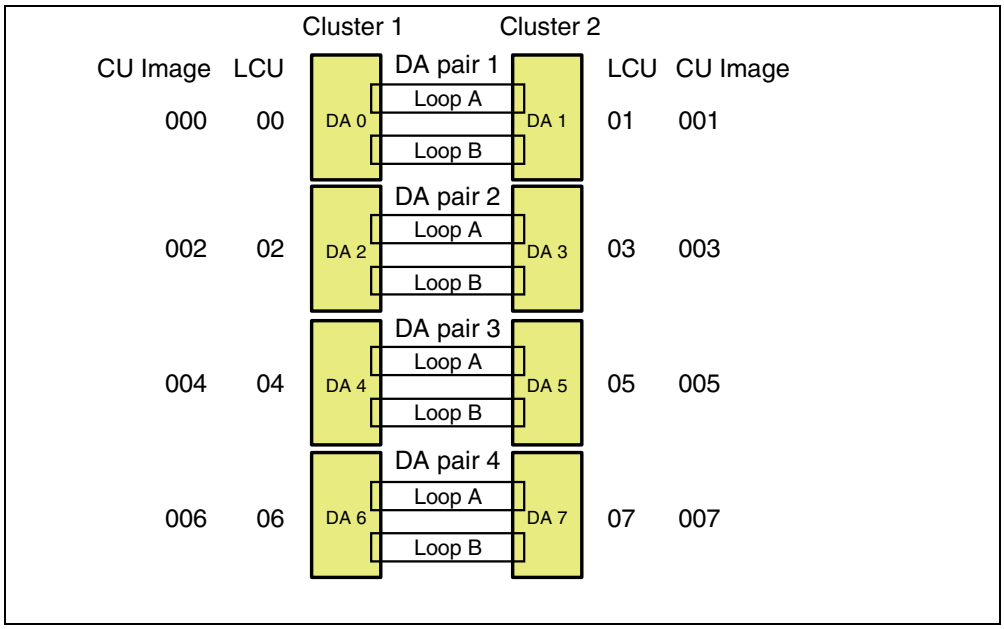

Figure 86. Relationship between S/390 LCUs and SSA loops

For example, LCU 04 is managed by storage cluster 1 and by DA 4 (in DA pair 3). The logical volumes belonging to LCU 04 are physically located on one or more (maximum of six) logical disk groups on loops A and B of DA pair 4. Because SSA loops usually contain an even number of drive bays, they can have an even number of logical disk groups defined, divided between the two storage clusters. Thus loops A and B of DA pair 3 may also have one or more logical disk groups that belong to LCU 05 and are managed by storage cluster 2 and by DA 5 (also in DA pair 3).

In the event of a cluster failover, for example, a failover of cluster 2, both LCU 04 and LCU 05 are managed by cluster 1, and all the logical disk groups on DA pair 3 are managed by DA 4. This is similar for the other SSA loops.

Figure 86 also shows the correspondence between the LCUs and the CU image numbers as displayed in the ESS Specialist S/390 Storage Panel.

9.1.3 S/390 logical volumes

Here we summarize the attributes of S/390 logical volumes:

- A S/390 logical volume is an emulated 3390 model 2, 3, or 9.
- If located in an interleaved partition of a RAID 5 array, a 3390 logical device has the same number of cylinders as the 3390 model being emulated, and the same data capacity.
  - 3390 Model 2 has 2226 cylinders, and 1.89 GB capacity.
  - 3390 Model 3 has 3339 cylinders, and 2.84 GB capacity.

    - 3390 Model 9 has 10017 cylinders, and 8.51 GB capacity.
- If located in a non-interleaved partition in a RAID array or on a non-RAID disk, logical volumes are custom volumes. That is, you can specify a logical volume with any number of cylinders from 1 to 10017.
    - A custom volume with 1 to 3339 cylinders is a 3390-3 Custom Volume.
    - A custom volume with 3340 to 10017 cylinders is a 3390-9 Custom Volume.
- If located on a non-RAID disk, a logical volume has no parity protection in the event of a disk failure.
- The track format of a logical device can be either 3390 or 3380.
    - The number of cylinders is that of the 3390 model (interleaved), or the number you specify (custom volume).
    - The data capacity of a logical device (interleaved or custom volume) is 849,960 bytes per cylinder in 3390 track format .
    - The data capacity of a logical device (interleaved or custom volume) is 712,140 bytes per cylinder in 3380 track format.
- A logical volume in 3380 track format does not emulate a 3380; it emulates a 3390 in 3380 track compatibility mode.

9.1.4 Logical disk groups

A logical disk group is the unit of allocation for physical disk drives. It contains eight disk drives.

The view of disk drive groups provided by the ESS Specialist on the Storage Allocation panel is a logical view.

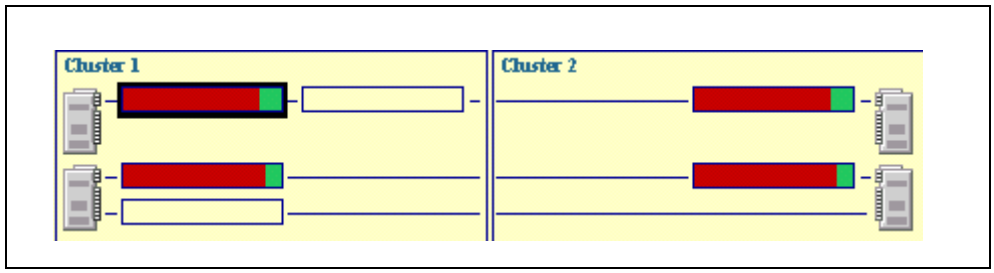

Figure 87. Logical disk groups

In Figure 87, each rectangle on the Storage Allocation panel represents a logical disk group containing eight disk drives on a loop. Only defined logical groups are shown. There may be drives on the loop that have not yet been defined into a logical group. The highlighted groups shown in the figure are assigned as S/390 storage. They are highlighted because we clicked on the ESCON Net host icon. The unhighlighted groups are configured as FB.

A group may be configured as a RAID 5 array or as individual non-RAID disks (JBODs). A RAID array or JBOD can be configured as either FB or S/390 storage. In this chapter, we consider only S/390 storage. For the definition of FB storage, refer to Chapter 15, “ESS configuration for fixed block storage” on page 211.

The physical drives on an SSA loop that belong to a logical group may reside in one or more physical drive bays. A drive bay contains eight disk drives. When

configuring a group, available drives (those not currently part of an existing group) are selected by the ESS Specialist to give the best performance characteristics. For example, the ESS Specialist may select for a RAID array the eight disks that are closest on the loop to the DA that will manage the array. The disks forming a RAID array may be distributed on both legs of the loop. See Figure 88 for details.

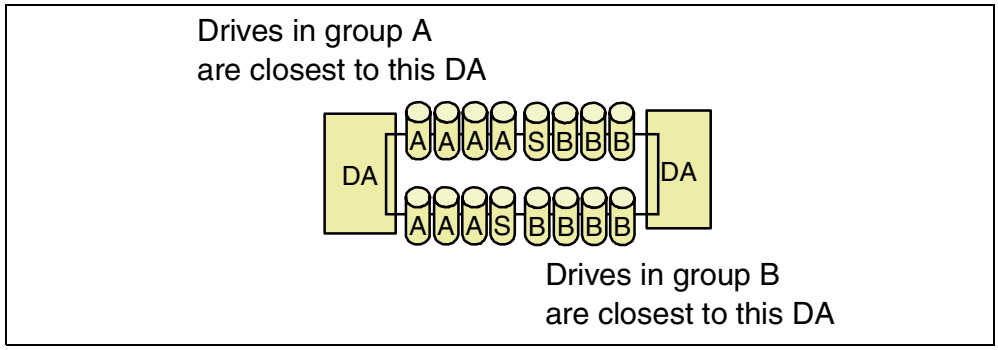

Figure 88. Example disposition of drives in RAID 5 groups.

During operation, one or more disk drives may fail over a period of time, and be replaced in the RAID array by spare drives elsewhere on the loop. The eight physical drives corresponding to a RAID logical group may thus change with time.

A drive bay is on one leg of an SSA loop. A logical disk group can contain any eight drives on the loop.

At present, it is not possible to display the disposition of physical drives versus logical disk groups using the ESS Specialist.

Summary of logical disk group attributes:

- A logical group contains eight disk drives.
- A logical disk group is contained on a single SSA loop. Groups cannot span both loops on an adapter pair.
- Disk drives are physically installed in drive bays. Each drive bay contains eight disk drives. The number of drive bays per loop can be 0,2,4, or 6.
- The maximum number of logical disk groups you can define on a loop is the same as the number of drive bays installed on the loop. That is, it can be 0,2,4, or 6.
- The odd numbered logical disk groups on the ESS Specialist Storage Allocation panel are managed by cluster 1. The even numbered groups are managed by cluster 2.
- A logical disk group can be configured in either of two ways:
  - A RAID 5 array
  - Non-RAID disks (JBODs)

Summary of RAID array attributes:

- A RAID 5 array is configured as either 6+P+S, or 7+P.
- The first two RAID 5 arrays configured on an SSA loop are 6+P+S. This ensures that two spare drives are available on the loop if two or more RAID 5 arrays are configured.

- The third and subsequent RAID 5 arrays on an SSA loop are 7+P.
- Generally, all RAID arrays in an ESS base frame are 6+P+S.
- A spare drive can be used for any RAID 5 array on that loop, either FB or S/390.
- The drive bays in the base frame are closest on the SSA loops to the DAs.

9.1.5 Configuring S/390 storage

Definition of logical volumes in the ESS for use with S/390 systems is straightforward. It involves four main tasks using the ESS Specialist:

Configure logical control units (LCUs)
In performing this task, you select the control unit emulation mode (3990-3, 3990-3 TPF, or 3990-6), specify the SSID for the LCU, and enable or disable Parallel Access Volumes (PAVs) for the LCU. This is done for one to 16 S/390 LCUs.

Configure disk groups for each LCU
Here, you select one or more logical groups of disks on the DA and loop that corresponds to the chosen LCU. For these groups further choices are made:

- Format the group as a RAID 5 array, or as individual non-RAID disks.
- For RAID 5, auto-allocate a group of interleaved 3390 Model 2, 3, or 9 logical volumes, or reserve the array for Custom Volumes.

The interleaved volumes have the standard number of cylinders for the 3390 model you specify.

Add logical volumes for each LCU
Here, you allocate logical volumes to:

- Non-interleaved RAID arrays
- Non-interleaved partitions in interleaved RAID arrays
- Non-RAID disks.

The volumes you add are custom volumes, that is, you can specify any number of cylinders per volume from 1 up to 10017.

Set number of PAVs per volume
If you want to enable PAVs for an LCU, after the logical volumes have been added you can set the number of PAV aliases for the volumes in the LCU. This value can be set for a single volume or selected groups of volumes.

9.1.5.1 Task sequence

We recommend that you do the tasks in this order:

1. Configure all the LCUs you need now, and all the LCUs you may want to have later. Addition of LSS definitions is a disruptive process which requires initialization of the ESS microcode.
2. Configure logical disk groups for the LCUs.
3. Add logical volumes for the LCUs.
4. Set the number of PAVs per volume for the allocated volumes in the LCUs.

Task (3) is not necessary if you intend to use only auto-allocated volumes, not custom volumes.

Task (4) is not necessary if you do not intend to enable parallel access volumes (PAVs) for any S/390 LCU.

9.1.6 ESS Specialist S/390 panels

Four ESS Specialist panels are used to configure S/390 LCUs and logical disks:

S/390 Storage panel: Here you can display the LCUs and logical devices that have been configured in tabular form. From this panel you can select the other S/390 panels.

Configure LCU panel: Here you configure the LCUs. This panel is also used to define the number of PAV aliases per volume.

Configure Disk Groups for CU Image panel: Here you define RAID arrays and non-RAID disks for an LCU.

Add Volumes for CU Image panel: Here you allocate logical volumes in the RAID and non-RAID space previously defined for an LCU.

Refer to Figure 89 for more information.

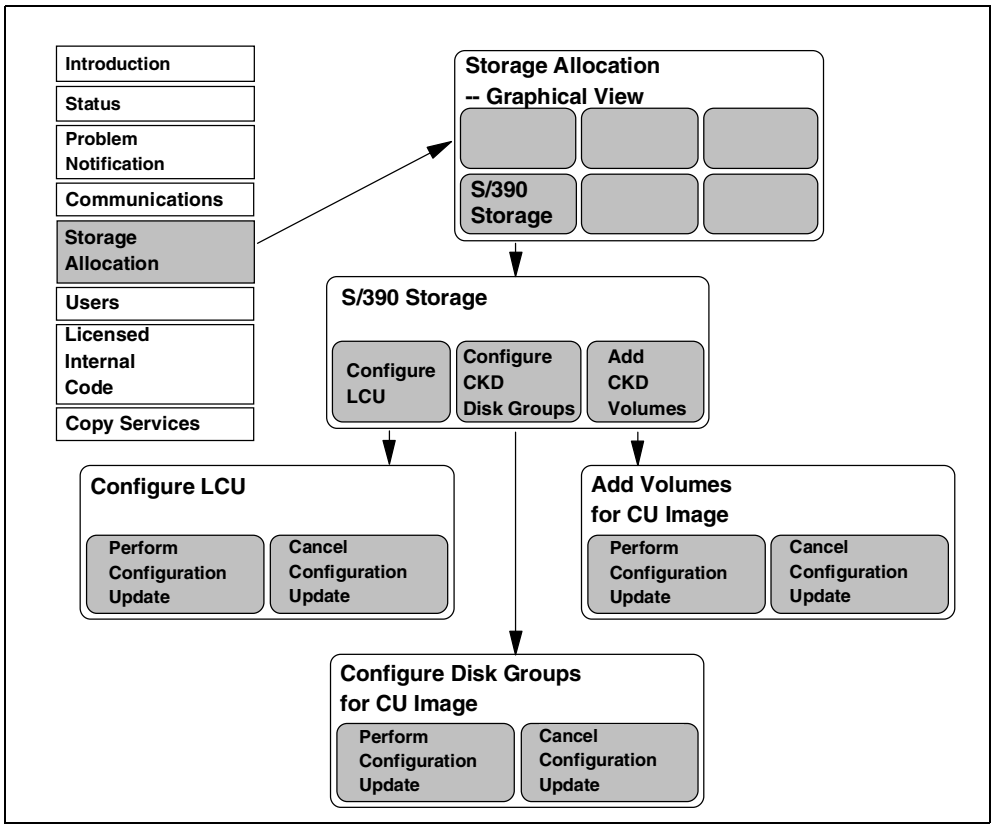

Figure 89. ESS Specialist: S/390 panels

From the Configure LCU panel, the Add Volumes for CU Image panel, and the Configure Disk Groups for CU Image panel, you can elect not to perform a previously entered configuration change by clicking on the Cancel Configuration Update button. This will return you to the S/390 Storage panel.

9.2 Start here

This section describes how to connect to the ESS Specialist and display the Storage Allocation panel.

9.2.1 Connecting to ESS specialist

Connect to the ESS Specialist by entering the hostname or TCP/IP address of either cluster as the URL in your Web browser. Refer to 5.5, “Connecting to your ESS Specialist” on page 76.

Note: If your Web browser (Netscape or MSIE) hangs or stops responding during configuration procedures, this is a problem with the browser, not ESS Specialist or the Windows operating system. You need to press **CTRL-ALT-DEL** and use the task manager to select the Web browser program. Under Status, it will probably say "not responding". End the task, then restart the browser and make the necessary selections to return to where you were.

Note: Ensure that any domain nameserver specified in the ESS is actually available. If it is not, this will appreciably slow the operation of ESS Specialist.

,

Figure 90. ESS Specialist: Introduction panel

You should see the ESS Specialist introduction panel as shown in Figure 90.

Every panel in ESS Specialist has a help button (?) near the top right corner. Selecting this button displays the ESS Specialist Help text. Select the Task Help tab for specific information on the functions you are using. Refer to 5.3, “Help” on page 74 for further information.

9.2.2 Storage allocation panel

From the Introduction Panel, select the **Storage Allocation** button.

If this is the first function selected since connecting to the ESS, you may be presented with a New Site Certificate window. Refer to 5.4, “Security” on page 75. View the series of certificates and proceed to the Username and Password Required panel (see Figure 31 on page 79). Enter a user ID with configuration authority and the password.

You should now see the Storage Allocation - Graphical View panel. Refer to Figure 91.

The icon labelled ESCON Net near the top of the frame is always present. It represents the S/390 hosts that can access S/390 logical volumes in the ESS.

ESCON connections in the ESS share these characteristics:

- All S/390 host channels use the same ESCON protocol
- Every ESCON adapter port installed in the ESS can address any logical volume (up to 256) in any of the defined S/390 LCUs (up to 8)

Only one ESCON host definition is needed, and ESS Specialist has pre-defined it as ESCON Net. The HCD definition for each S/390 host specifies which LCUs and logical volumes in the ESS can be reached by that host.

If you click on the ESCON Net icon, configured S/390 disk groups display, and the ESCON adapter ports highlight in yellow, as shown in Figure 91. This indicates that all installed ESCON ports are automatically associated with the ESCON Net host.

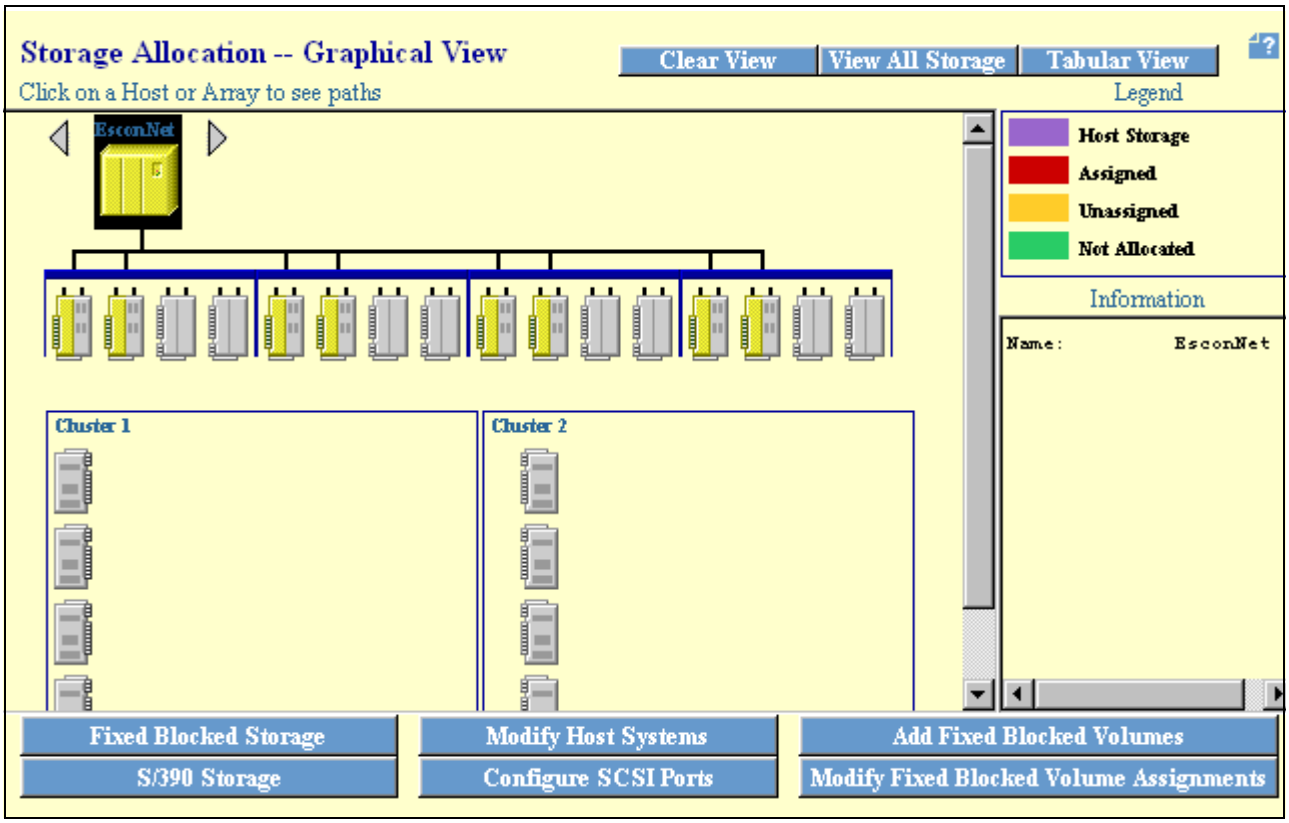

Figure 91. ESS Specialist: Storage allocation - Graphical view

In a newly installed ESS, before any disk groups have been configured, the four DAs for storage cluster 2 are shown to the left of the box in the Storage Allocation panel, as shown in Figure 91. When groups are defined, the cluster 2 DA icons move to the right side of the box, as shown in Figure 109 on page 180.

When groups are defined on only one loop of a DA pair, the single loop may appear as in Figure 87 on page 155, toward the top of the DA icons. This does not necessarily mean it is loop A. Click on one of the groups on the loop and refer to the information box. In Figure 87, the disk group shown on DA pair 1 could be on either loop A or loop B of DA pair 1. DA pair 2 has a group on both loop A and loop B.

On the Storage Allocation panel, you can select Tabular View. This displays a table showing all the configured logical volumes. For detailed information, refer to 5.6.5, “Storage Allocation panel” on page 83.

In the lower part of the Storage Allocation panel are six buttons to select functions. S/390 Storage is the only function you need to select from the Storage Allocation - Graphical View panel to configure S/390 devices in the ESS.

Note: During configuration procedures using ESS Specialist, you will be presented with these messages:

- 2108: Retrieving the current configuration data ... Please wait.
- 1533: Down level data ... Please wait a few minutes and try again.

ESS Specialist allows more than one user to be logged on simultaneously. Users may be logged onto ESS Specialist in the other cluster. To ensure it has accurate configuration data, ESS Specialist needs to read the current configuration of the machine. Message 2108 usually appears when a function is selected on one of the panels.

You see message 1533 after a configuration change ends with a successful completion message, and you click **OK** to return to the S/390 Storage panel. After message 1533 is displayed, click the **Refresh Data** button on this panel; the panel will not refresh automatically. After clicking the Refresh Data button, you receive this message:

- 2122: Retrieving the current configuration status ... Please wait.

If this is again followed by message 1533, wait a few minutes before re-trying.

Normally, message 1533 persists for about the same time as the function just successfully completed. The ESS Specialist is synchronizing configuration data in the two clusters.

Note: When using ESS Specialist, ensure that the CE terminal is not logged on via the serial port. If you see a message that one or more components are in service mode, this is probably the cause.

9.3 ESS Specialist S/390 panels

On the Storage Allocation - Graphical View panel, click on **S/390 Storage.** ESS Specialist displays the S/390 Storage panel.

9.3.1 S/390 storage panel

Figure 92 on page 163 shows the S/390 Storage panel.

Note: The array and device capacities shown in the example ESS Specialist panels in this book may differ from those you see when using ESS Specialist. You can assume that the values you see are correct. The example panels were taken while the ESS and ESS Specialist were under development.

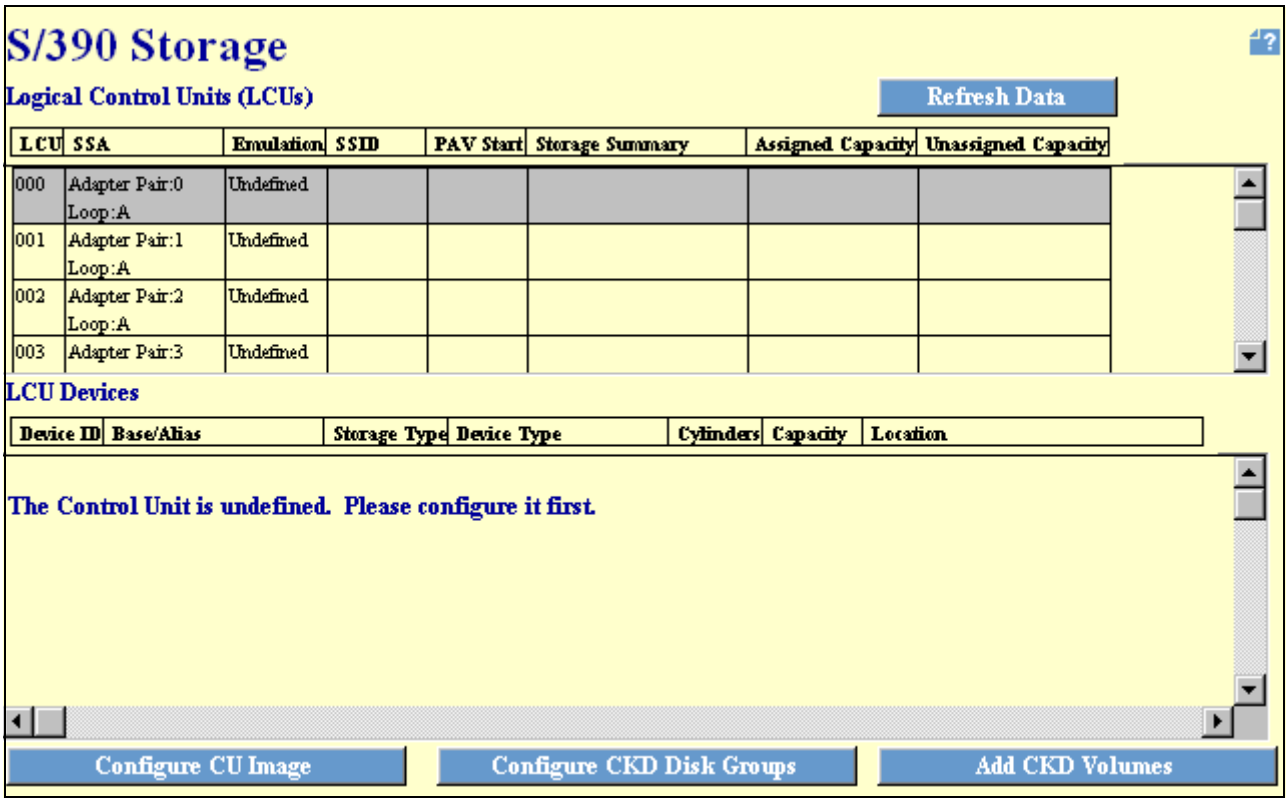

Figure 92. ESS Specialist: S/390 storage panel

On the S/390 Storage panel, you can scroll through the list of eight LCUs (000 to 007). These correspond with LCU 00 through LCU 07 in Figure 86 on page 154. Initially, they are all undefined. The correspondence between LCUs, Device Adapters, and SSA loops is shown in Figure 86 on page 154.

Note: If you want to go back to the Storage Allocation panel from here, click on the Storage Allocation button on the left of the ESS Specialist navigation frame.

Select an LCU to configure by clicking on the corresponding row in the S/390 Storage panel. The row will highlight in grey. A message appears in blue in the lower part of the panel:

The Control Unit is undefined. Please configure it first.

Then select the **Configure CU Image** button. If you see the error message in Figure 93 on page 164, you did not select a CU image before pressing the Configure CU Image button.

Figure 93. ESS Specialist: Error 1820

9.3.2 Configure logical control unit (LCU)

In this example, LCU 000 is selected, and the Configure LCU 0 panel appears as shown in Figure 94. The CU image is initially undefined.

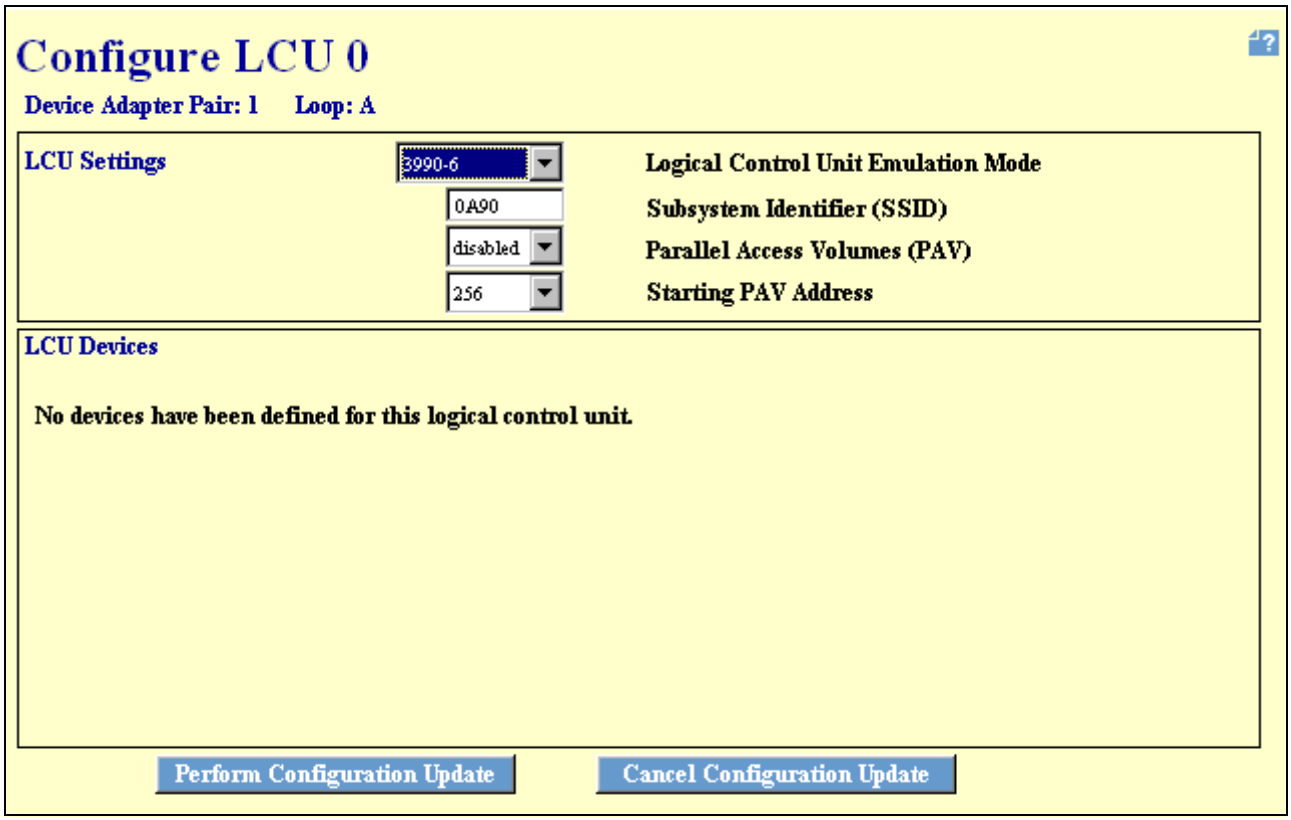

Figure 94. ESS Specialist: Configure LCU panel

9.3.2.1 LCU emulation mode

Display the pull-down menu for the Control Unit Emulation Mode, and select one of the available modes:

- 3990-6
- 3990-3
- 3990-3 TPF
- Undefined

Other options then appear on the panel.

Note: You should define all LCUs that you anticipate a need for. Later configuration of additional LCUs is disruptive because initialization of the ESS is needed to activate them.

Use Figure 96 on page 167 as a guide to selecting options and configuring the LCU.

9.3.2.2 Subsystem ID

You need to know the value of the Subsystem ID (SSID) to be assigned to each LCU. This number should be unique in the location, that is, no other storage subsystem in the location should use the same SSID.

The SSID must not have FF as the first two digits, or 00 as the last two digits. The SSID could be based on the ESS serial number, with the last digit corresponding to the LCU number.

For example, 2105-013FFA90 could have SSIDs FA90 through FA97 for the eight LCUs.

9.3.2.3 PAV starting address

If you want to enable PAVs for this CU image, you can do so now, or you can leave PAVs disabled at this time. If configuring PAVs, you will return to the Configure LCU panel after you have configured the logical volumes for the LCU, to assign PAVs for the volumes. You can enable PAVs for the LCU at that time.

If you enable PAVs for this LCU, set the PAV starting address higher than the number of base devices, and lower than the total number of base devices and PAV aliases combined, that you plan to define for the LCU.

PAVs are assigned using the highest unused address available, beginning at the PAV starting address, proceeding in decreasing address order.

With the PAV starting address set to less than the total number of bases and PAV aliases, the PAVs are assigned from the PAV starting address down to the highest base + 1, then from the PAV starting address + 1 upwards, to the limit of 256 addresses available in an LCU. This results in a contiguous set of device addresses for the LCU. You can assign a single range of UCBs in the host operating system to cover the base devices and PAV aliases.

For example, if you configure 64 base devices, and two PAVs for each base, this requires 192 addresses in the LCU. Set the PAV starting address to 128. The addresses in use for the LCU are then x'00' - x'3F' for the bases, x'7F' down to x'40', and x'80' - x'BF' for the aliases. In HCD, specify 192 UCBs starting at x'00' to cover addresses x'00' - x'BF', the first 64 of these being defined as 3390B (Base) device types, and the following 128 addresses being 3390A (Alias) device types.

Alternatively, set the PAV starting address to 256. If fewer than 256 bases and aliases are configured, this results in a gap in the address range for the LCU. For example, if you define 64 base devices in an LCU, two PAV aliases for each base, and a PAV starting address of 256, the resulting addresses are x'00' - x'3F for the base devices, and x'80' - x'FF' for the PAV aliases. In the HCD configuration, specify two address ranges for the LCU: 64 addresses starting at x'00', and 128

addresses starting x'80'. If the number of UCBs is not important, consider allocating extra PAV aliases to fill the 256 address range.

9.3.2.4 Deleting an LCU

You can delete an existing configured LCU by displaying its Configure CU Image panel, and selecting **Undefined.** When you click on **Perform Configuration Update,** the warning message shown in Figure 95 is displayed.

Figure 95. ESS Specialist: Message 1550

If you click on **Yes**, the LCU Emulation Mode changes to undefined, the LCU is deleted, with its associated disk groups, arrays, and logical volumes. Any data on those volumes is discarded.

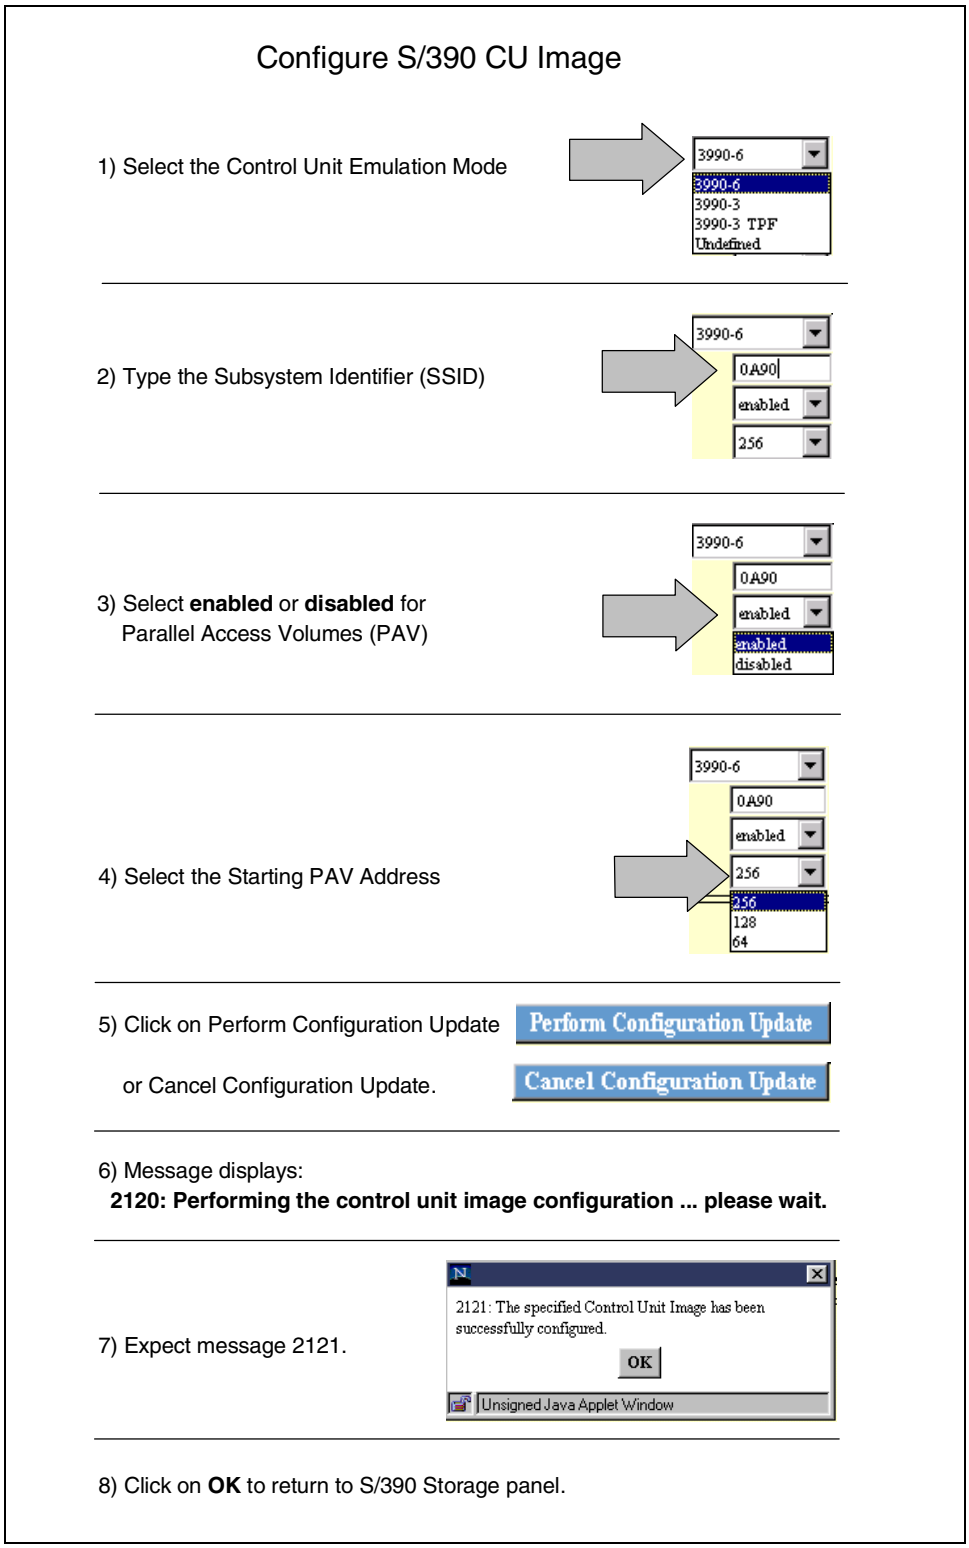

Figure 96. ESS Specialist: Procedure to configure a S/390 CU image

After successful configuration of an LCU, you are returned to the S/390 Storage panel. See Figure 96.

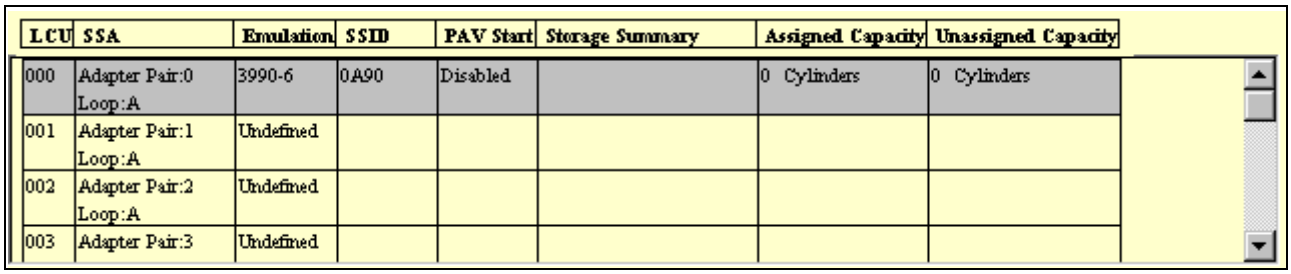

Figure 97. ESS Specialist: S/390 storage panel CU image table

Figure 97 shows the CU Image table from the S/390 Storage panel. The table now shows CU image 000 (LCU 00) defined, with the 3990-6 emulation and SSID selected previously. In this example we left PAVs disabled. Assigned capacity is zero, until you configure disk groups for the LCU.

On this panel, select the CU image you want to add disk groups to. It will highlight in grey. Click on **Configure CKD Disk Groups**.

9.3.3 Configure disk groups for an LCU

Figure 98 shows the Configure Disk Groups for CU Image panel.

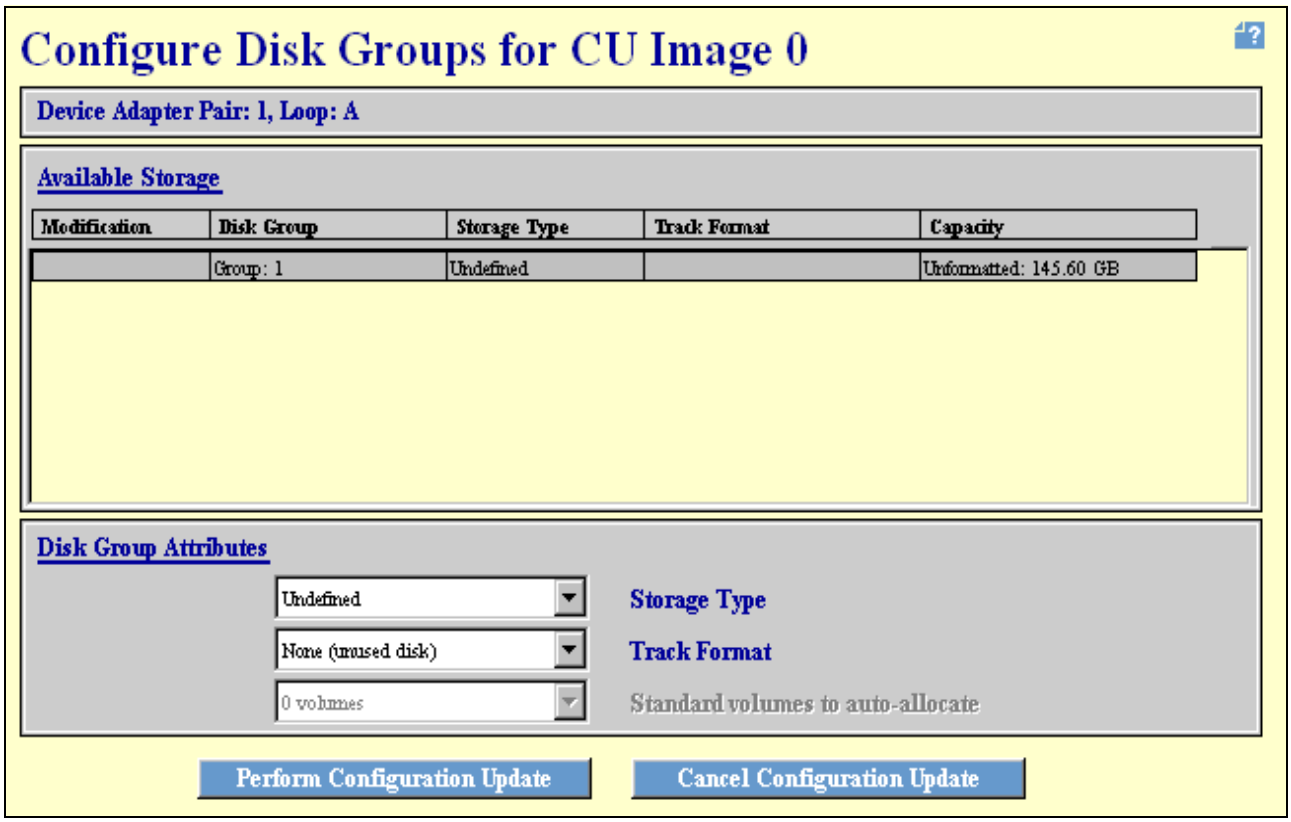

Figure 98. ESS Specialist: Configure disk groups for CU image panel

Note: The *even* numbered LCU (in cluster 1) is associated with *odd* numbered disk groups (1,3,5) on an SSA loop. The **odd** numbered LCU (in cluster 2) manages the *even* numbered disk groups (2,4,6).

Select a group to configure by clicking on the corresponding row in the Available Storage table. The row highlights in grey. Use Figure 100 on page 172 as a guide to configure the logical disk group as RAID array.

9.3.3.1 Defining Non-RAID Storage

The procedure shown in Figure 100 creates a RAID array with interleaved 3390 logical volumes. If you want to configure the group as JBODs:

1) Select non-RAID in step 1.

When you select non-RAID, the row in the table expands to several rows, one for each of the disk drives in the group. The Storage Type field in the table is non-RAID. The Capacity field has the unformatted capacity of a single disk drive.

2) Select 3380 or 3390 track format in step 2.

Non-RAID disks can then be individually selected, and the track format changed to 3390 or 3380 track mode.

3) The Standard volumes to auto-allocate is set to zero. You cannot change it. The volumes will be allocated later using the Add Volumes for CU Image panel.

5. Click on **Perform Configuration Update**.

Note: A non-RAID disk group may contain both FB and CKD disks in any combination. If a disk group is defined as non-RAID using this panel and any disks in the group are not defined as 3390 or 3380 track format, these disks are available to be defined as Fixed Block track format using the Configure Fixed Block Disk Groups panel. Similarly, if a non-RAID disk group is defined using the FB panel, any disks not defined as FB appear in the Configure Disk Groups for CU Image panel for the LCU corresponding to the disk group, and are available to be defined as S/390.

9.3.3.2 Interleaved volumes

When you create a S/390 RAID array, fixed size 3390 interleaved logical volumes are auto-allocated in the array.

Note: If you do not want to configure interleaved volumes, in step 3 select **0 volumes** to auto allocate zero standard volumes.

When volumes are auto-allocated, two partitions are created in the array. The first partition contains the interleaved volumes and takes up most of the capacity in the array. The second partition is non-interleaved. Logical volumes are added to the non-interleaved partition using the Add Volumes for CU Image panel.

How big is each partition? Initially, 5000 cylinders are reserved for the non-interleaved partition. The remaining space is available for interleaved volumes. The size of the remaining space depends on the total array capacity.

The total array capacity depends on:

- The size of the disk drives in the group (9.1 GB, 18.2 GB, or 36.4 GB)
- The array configuration (6+P+S, or 7+P)

Refer to 2.2.3.4, “S/390 CKD data capacity” on page 14 for RAID array capacities.

Logical volumes of the type specified (3390-2, 3, or 9) are allocated to the interleaved partition in groups of four, until insufficient space remains to accommodate another group. The remaining space is then added to the non-interleaved partition.

Note that the first two RAID arrays (FB or S/390) configured on any SSA loop are 6+P. Subsequent arrays, if any, are 7+P. Because a 7+P array has one extra drive, it has more capacity than a 6+P array, and the number of volumes auto-allocated for each 3390 model is larger than for a 6+P array. The unformatted, or formatted capacity of each disk group is shown in the capacity column of the Available Storage Table. In Figure 98 on page 168, this capacity is shown as 145.60 GB. This is the unformatted capacity of the eight 18.2 GB drives in this disk group.

The auto-allocated volumes have the standard number of cylinders for the 3390 model selected:

3390 Model 2 2226 cylinders

3390 Model 3 3339 cylinders

3390 Model 9 10017 cylinders

If you selected 3380 track format, the 3390-9 is not available. The 3390 volumes have the same number of cylinders as above, but the cylinders are in 3380 track format. Due to the smaller track capacity of 3380 formatted volumes, the number of volumes auto-allocated in the array for a given 3390 model may be greater for 3380 track format than for 3390 track format.

The minimum size of the non-interleaved partition is 5000 cylinders of the selected track format (3380 or 3390). The maximum size is smaller than 5000 cylinders plus 4*N, where N is the standard number of cylinders for the selected 3390 model being allocated to the interleaved partition.

Why interleave logical volumes? When logical volumes are mapped to the physical disks in a non-interleaved RAID 5 array, part of each volume occupies a band of tracks on each disk drive. Performance of individual logical devices varies according to the position of the tracks they occupy on the drives. Lower numbered volumes are located near the outer diameter of the disks, and tend to perform better than higher numbered volumes nearer the center of the disks.

In the interleaved partition, all the allocated volumes are interleaved together across the physical tracks on the disk drives in the RAID array. Logical volumes have similar performance, but lower numbered logical cylinders on the volumes may have higher performance than higher numbered cylinders because of their location nearer the outer diameter of the physical disk drives.

Best overall performance is obtained using interleaved 3390 volumes.

If you intend to allocate either 3390-2 or 3390-3 volumes to an array, you might instead consider allocating a smaller number of 3390-9 volumes, and using PAV aliases. Compared to 3390-3, configuring one third the number of 3390-9 volumes and using two PAV aliases per base uses the same number of UCBs, and the same amount of ESS storage capacity, but can give improved performance under heavy workloads due to the increased addressability of the data in the volumes. Other factors such as use of reserve enqueue may affect

performance, because each reserve serializes access to three times as much data for 3390-9 volumes compared with 3390-3 volumes.

9.3.3.3 Deleting a Disk Group

You can delete an existing configured logical disk group by selecting **Undefined** under Storage Type on the Configure Disk Groups for CU Image panel. The warning message shown in Figure 99 is displayed.

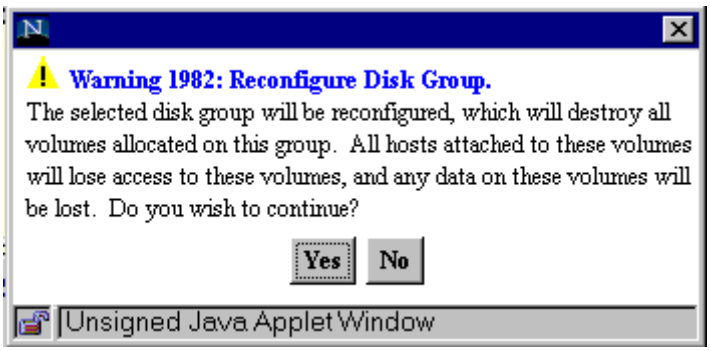

Figure 99. ESS Specialist: warning 1982

If you click on **Yes**, the Storage Type changes to undefined. When you click on **Perform Configuration Update**, the Disk Group is deleted, with its associated array and logical volumes. Any data on those volumes is discarded.

When a single non-RAID disk is undefined, the entire disk group is deleted. If any of the disks in a non-RAID group has been defined as an FB device, you cannot delete the group from here without first going to the Configure Fixed Block Disk Groups panel and changing the track format of all disks in this group to none. If all the disks in the group that are visible here have their track format changed to none, the group can be deleted from the Configure Fixed Block Disk Groups panel.

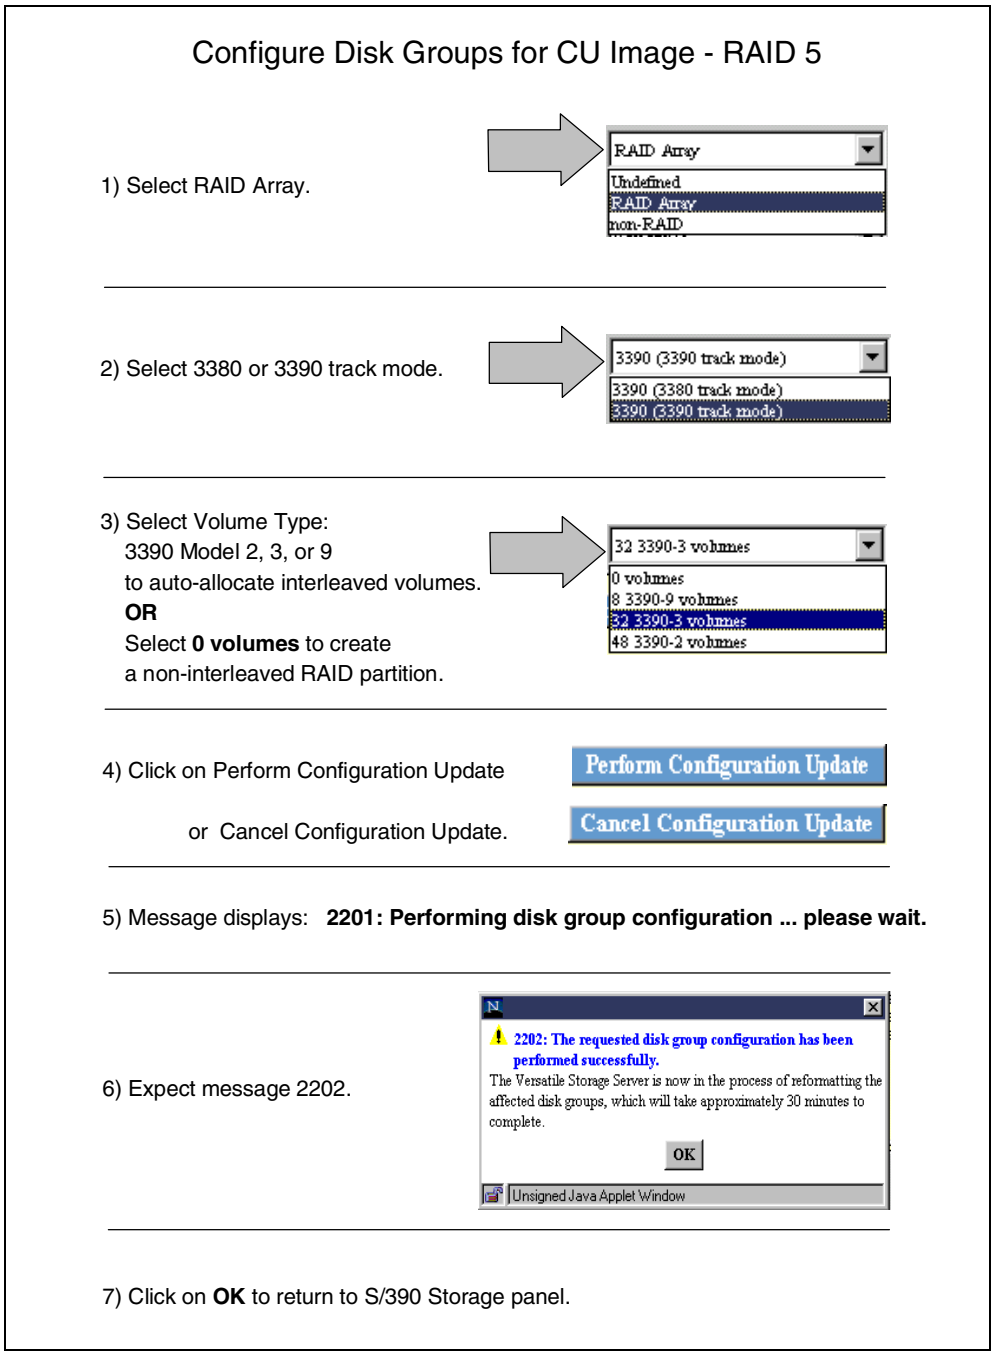

Figure 100. ESS Specialist: Procedure to configure disk groups

At step 3 in Figure 100, the number of volumes auto-allocated for each 3390 model varies, depending on the array capacity, and whether you chose 3390 track format or 3380 track format.

After clicking on **OK** at step 7 in Figure 100, you see message 2201:

2201: Performing Disk Group Configuration ... Please wait

This message displays before you are returned to the S/390 Storage panel. On the S/390 Storage panel, you can select other LCUs and configure further disk groups. The volumes auto-allocated in each array will not appear in the LCU Devices table until formatting of the disk group is complete. Several disk groups can be formatting simultaneously.

After successful configuration of a disk group containing interleaved volumes, the S/390 Storage panel appears as in Figure 101.

From the Logical Control Units section of the panel, the Assigned Capacity column contains the combined capacity of the auto-allocated volumes in the RAID array you just configured, if this is the first group you have configured for the LCU. The unassigned capacity column shows the capacity of the non-interleaved partition in that array.

The LCU devices section of the panel displays the interleaved devices that were auto-allocated during the disk group configuration. The information for each device includes the device type (3390 model), cylinders per device, and device capacity.

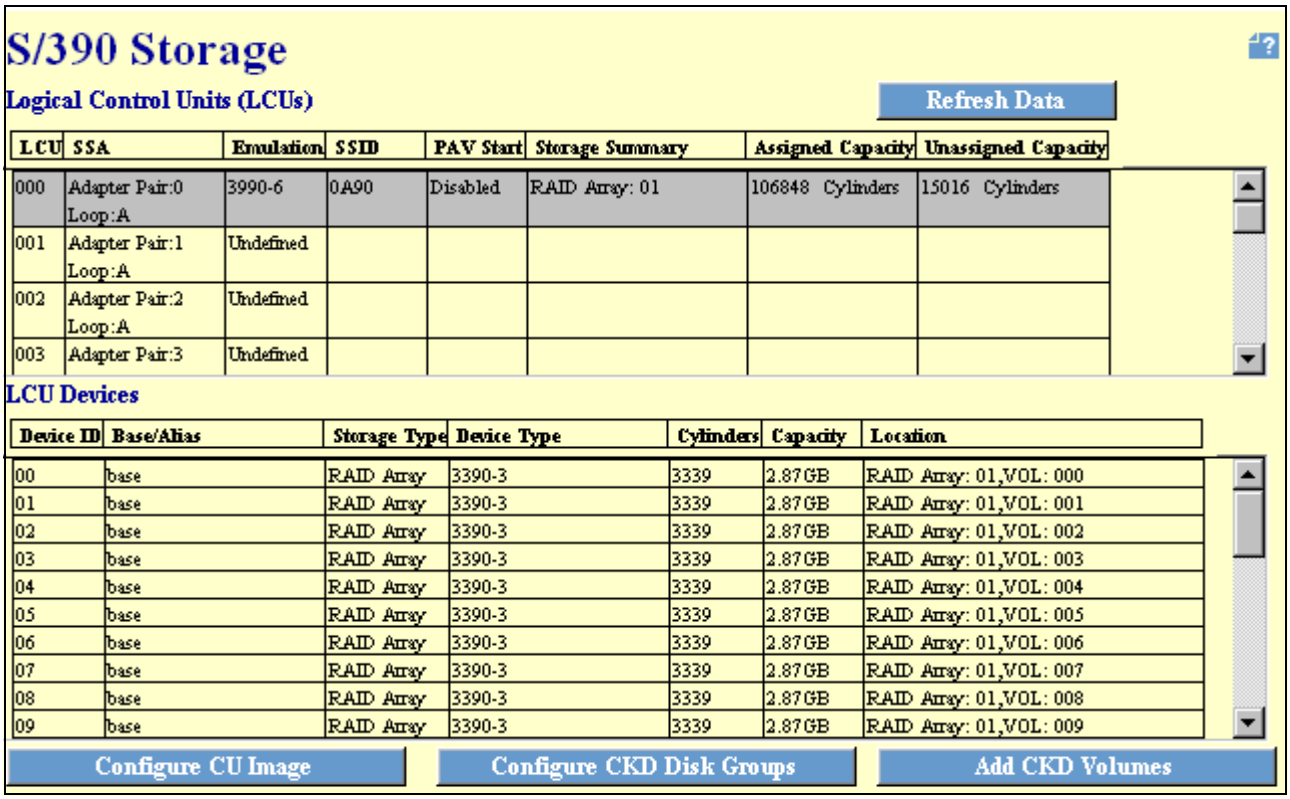

S/390 Storage

Logical Control Units (LCUs)

Refresh Data

LCU	SSA	Emulation	SSID	PAV Start	Storage Summary	Assigned Capacity	Unassigned Capacity
000	Adapter Pair:0 Loop:A	3990-6	0A90	Disabled	RAID Array: 01	106848 Cylinders	15016 Cylinders
001	Adapter Pair:1 Loop:A	Undefined					
002	Adapter Pair:2 Loop:A	Undefined					
003	Adapter Pair:3	Undefined					

LCU Devices

Device ID	Base/Alias	Storage Type	Device Type	Cylinders	Capacity	Location
00	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 000
01	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 001
02	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 002
03	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 003
04	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 004
05	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 005
06	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 006
07	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 007
08	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 008
09	base	RAID Array	3390-3	3339	2.87GB	RAID Array: 01,VOL: 009

Configure CU Image | Configure CKD Disk Groups | Add CKD Volumes

Figure 101. ESS Specialist: Interleaved volumes

Note that the cylinders per device and the device capacity are the standard values for the 3390 model specified. Interleaved devices are not Custom Volumes.

If during the disk group configuration you specified non-RAID, or RAID with zero volumes auto-allocated, the non-RAID or RAID storage will appear in the Storage Summary section of the Logical Control Unit table. No entries will appear in the

LCU Devices section of the S/390 Storage panel. Allocate the volumes later as Custom Volumes in the Add Volumes for an LCU panel.

To allocate the volumes, click on the **Add CKD Volumes** button on the S/390 Storage panel. You should then see the Add Volumes for CU Image panel.

9.3.4 Add volumes for an LCU

Figure 102 shows the Add Volumes for CU Image panel.

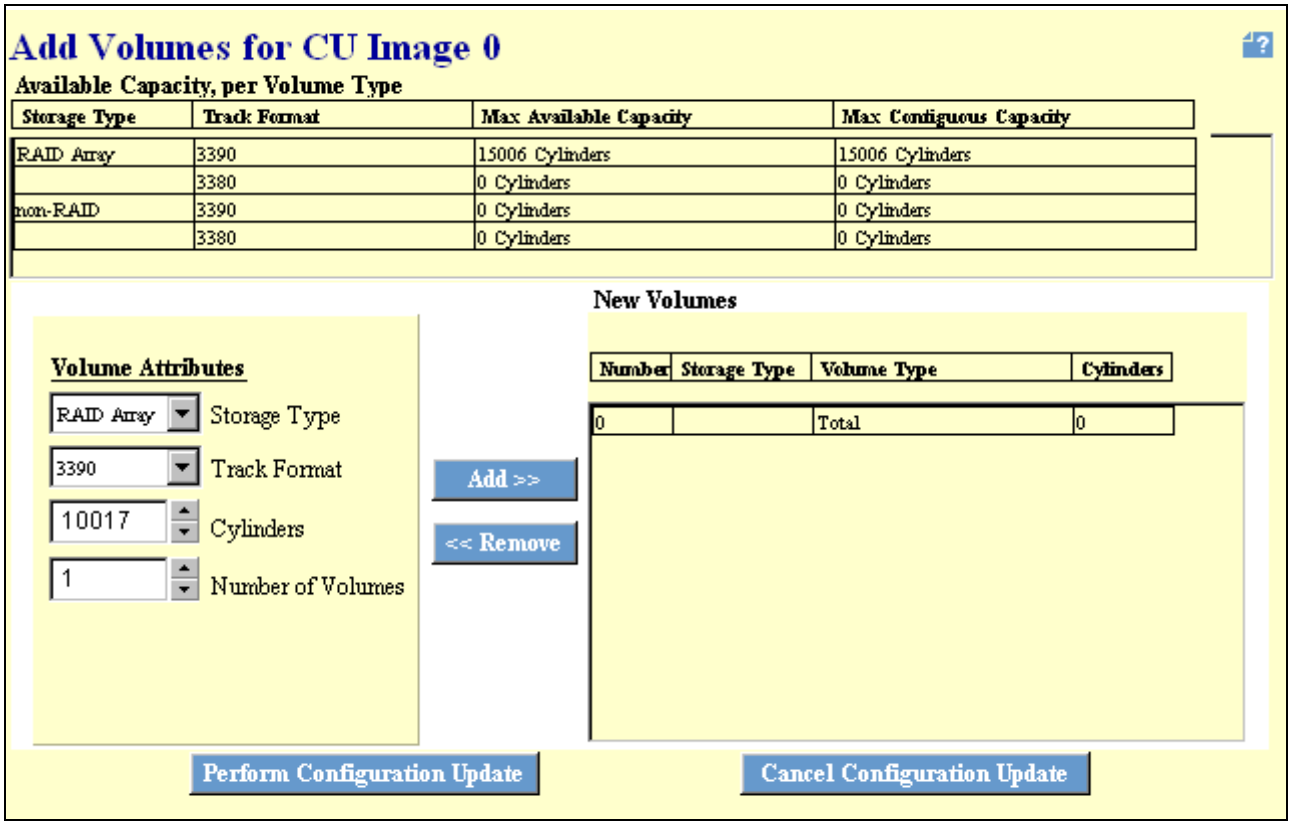

Figure 102. ESS Specialist: Add volumes for CU image panel

9.3.4.1 Storage type

The Available Capacity table shows available capacity for the CU Image, classified into four categories:

- RAID and 3390 track format
- RAID and 3380 track format
- Non-RAID and 3390 track format
- Non-RAID and 3380 track format

From the Volume Attributes section, you select the category of available space to add volumes to, by specifying RAID or non-RAID in the Storage Type field, and 3380 or 3390 in the Track Format field.

Note that if you auto-allocated 3390 track format devices in an array using the Configure Disk Groups for CU Image panel, the non-interleaved partition is available as RAID - 3390 format, and you can allocate only 3390 format custom

devices in the array. The same is true for 3380 track format. This is because a RAID array in S/390 must be formatted entirely in a single track format, either 3390, or 3380.

9.3.4.2 Number of cylinders

You set the number of cylinders per volume by typing it directly into the Cylinders field, or using the spin buttons. The maximum number you can set is 10017, or the number of cylinders in the Max Contiguous Capacity field in the table, whichever is smaller. If you have selected a category where no capacity is available, you cannot change the number from zero.

The volumes you add here are Custom Volumes, that is, you can specify any number of cylinders from 1 to 10017. If you want to emulate the standard capacities for 3390 volumes, specify:

1113 For 3390 Model 1

2226 For 3390 Model 2

3339 For 3390 Model 3

10017 For 3390 Model 9

For 3380 track format volumes, the standard cylinders per volume are:

885 For 3380-STD, 3380-D, or 3380-J

1770 For 3380-E

2655 For 3380-K

9.3.4.3 Number of volumes

The number of volumes to add to the selected category of space can be typed into the Number of Volumes field, or the spin buttons used to set it.

The available space may be distributed over several non-contiguous areas in multiple RAID arrays, or multiple non-RAID disks. A logical volume cannot span RAID arrays or non-RAID disks. The number of volumes of a given size you can add is limited to the number of segments of contiguous space of that size or larger in the category you choose. You can add volumes of different sizes in one category of available capacity, until no usable space remains.

Use Figure 103 on page 176 as a guide to adding volumes to a CU image.

9.3.4.4 Deleting a device

Prior to selecting Perform Configuration Update, you can remove a volume that has previously been added by selecting the device in the New Volumes table and clicking on the **<< Remove** button. Note that you can delete only Custom Volumes. Interleaved volumes cannot be individually deleted.

After they are created, logical S/390 volumes cannot be individually deleted. If a logical disk group is undefined on the Configure Disk Groups for CU Image panel, the volumes residing on the group are deleted, and any data on them is discarded.

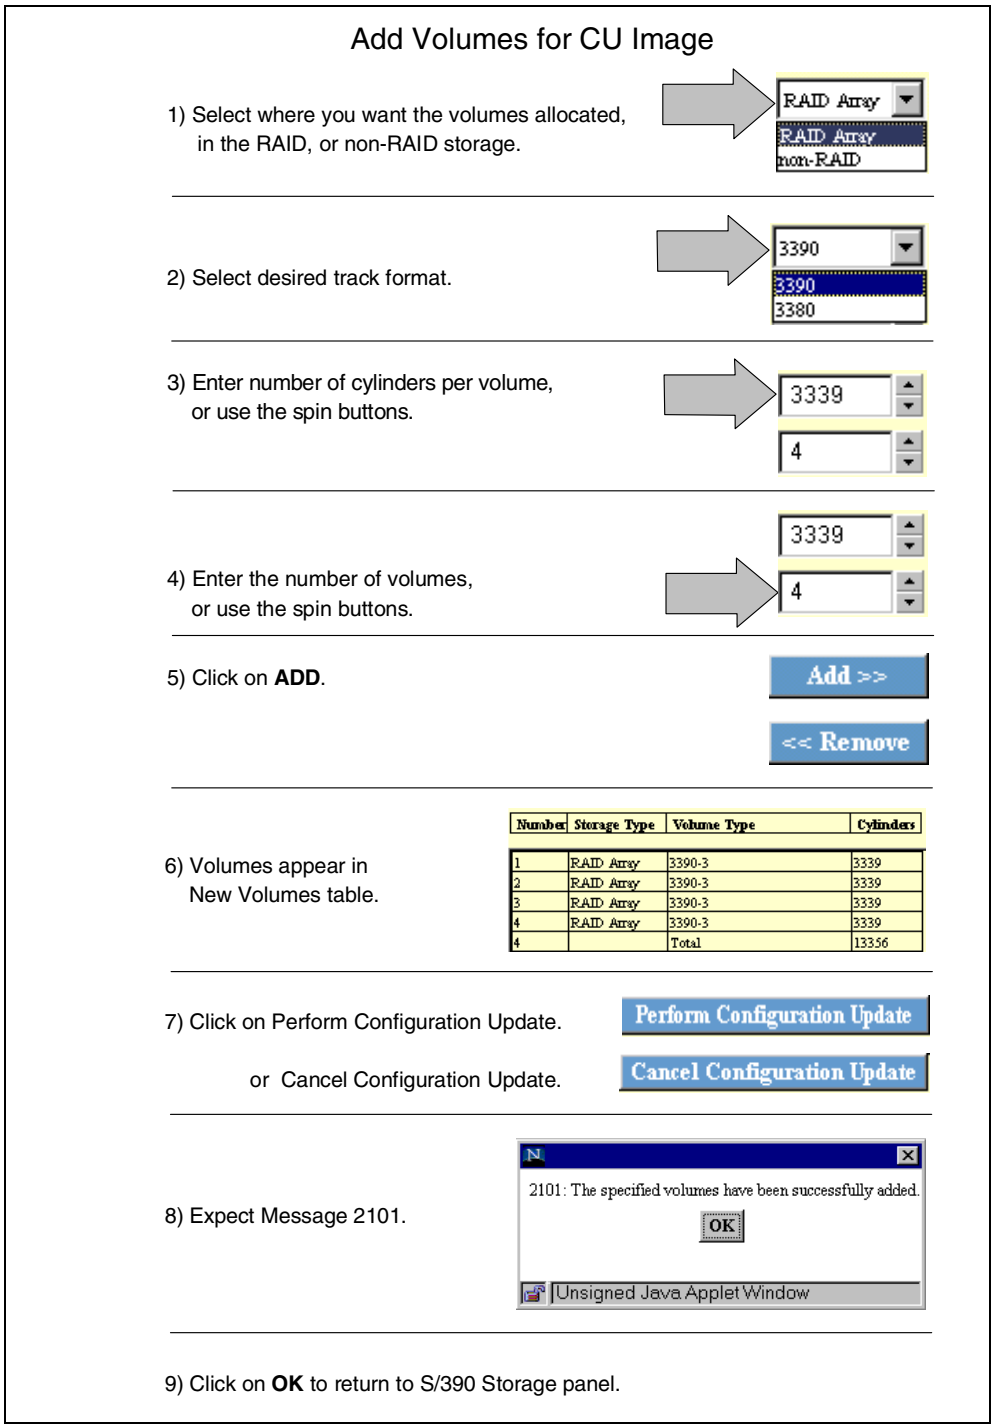

Figure 103. ESS Specialist: Procedure to add volumes

As volumes are added to a category of available space, they appear in the New Volumes table. If you specify a number of volumes that is too large for the available space, ESS Specialist displays message *Error 1901*. See Figure 104 on page 177. The volumes that can fit into the available space are allocated, and are shown in the New Volumes table.

Figure 104. ESS Specialist: error 1901

The Available Capacity table is updated, and the Cylinders field is updated to match the Max Contiguous Capacity field, if this contains less than 10017. Refer to Figure 105. The number of cylinders subtracted from the available capacity may exceed the total number of cylinders in the volumes added. This is because of overheads associated with each volume, for example, metadata.

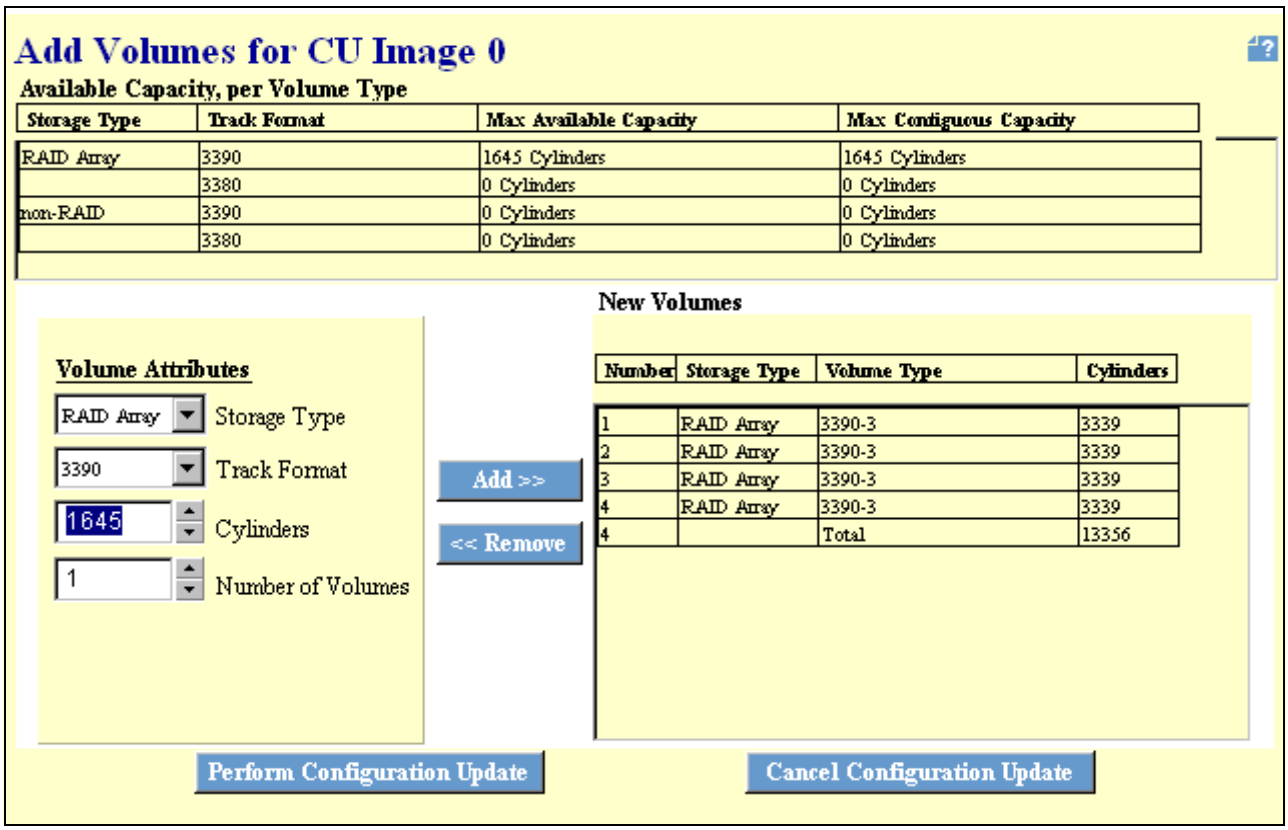

Figure 105. ESS Specialist: Adding custom volumes

When you have added the desired volumes, click on **Perform Configuration Update**. On the S/390 Storage panel, select your LCU. The LCU Device table shows the Custom Volumes you added.

Note: Logical Volumes are mapped to device IDs in the LCU in the order that volumes are created. The device IDs map to the device address as seen by the S/390 operating system. The lowest available device ID is used first. In the example, the 32 interleaved volumes map to device IDs 00 - 1F in LCU 0, and the four Custom Volumes added later map to IDs 20-23.

9.3.5 Set number of PAVs per volume

After you have defined an LCU, configured one or more disk groups, and added desired Custom Volumes, you can define the number of PAVs to assign to the volumes in the LCU. This task is not necessary for any LCU for which you intend to disable PAVs.

To assign PAVs, select the desired LCU on the S/390 Storage panel, then click on **Configure CU Image** to select the Configure LCU panel. Refer to Figure 106.

To select volumes, use the scroll bar to make the desired volumes visible in the LCU Device list, then click on a volume to select it. It will highlight in grey. Additional volumes can be selected using CTRL- left click. A range of volumes can be selected by clicking at one end of the range, then SHIFT- left clicking at the other end. A volume can be deselected using CTRL- left click. Once the desired volumes are selected, in the lower part of the panel select either:

- Add PAV aliases to the selected volumes

  Specify the number of PAV aliases to add for each selected volume by directly entering the number, or by using the spin buttons. Note that you cannot enter a non-zero value here unless PAVs are enabled for the LCU.

- Delete PAV aliases for the selected volumes.

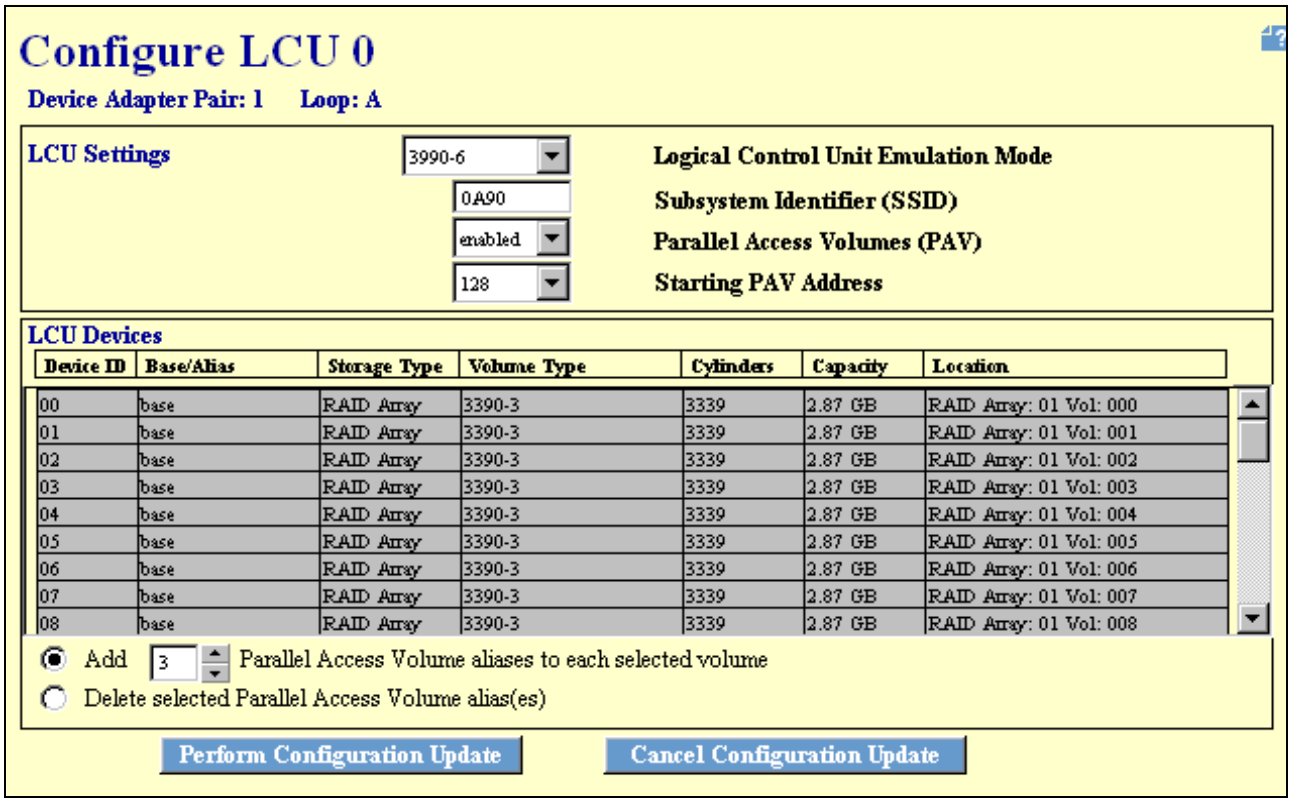

Figure 106. ESS Specialist: Configuring PAV aliases

In this example, we initially configured LCU 0 with PAVs disabled, (see Figure 94 on page 164). In Figure 106, on the Configure LCU panel, we have now enabled PAVs, and set the Starting PAV Address to 128. Refer to 9.3.2.3, “PAV starting address” on page 165 for further information.

In Figure 106 on page 178 we have selected all 32 volumes in the LCU. These are the auto-allocated volumes. There were no Custom Volumes configured in the non-interleaved partition of the array. Note that the volumes are listed in the Base/Alias column of the table as base device.

Further, we have specified that three PAV aliases be added for each of the 32 selected volumes. When **Perform Configuration Update** is clicked, this results in 128 addresses being configured for the LCU. You can select any combination of base devices and assign up to seven PAV aliases per device for each group you select, provided the total number of base devices and aliases for the LCU does not exceed 256.

Figure 107 shows the LCU Device table for LCU 0 after the aliases are configured. Note that this table can be viewed in either the Configure LCU panel, or the S/390 Storage panel.

Addresses (Device IDs) 00-1F are the 32 base devices previously configured. For each device, the Base/Alias column now has base (3 aliases exist).

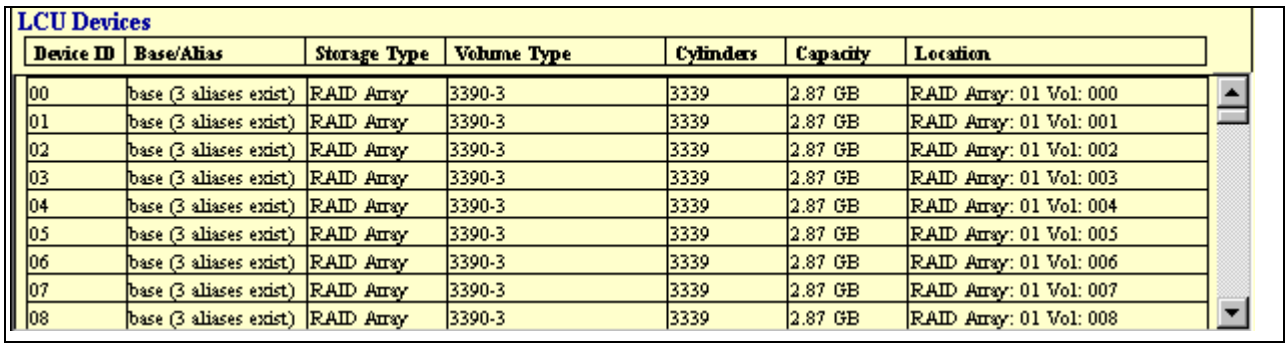

LCU Devices

Device ID	Base/Alias	Storage Type	Volume Type	Cylinders	Capacity	Location
00	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 000
01	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 001
02	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 002
03	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 003
04	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 004
05	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 005
06	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 006
07	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 007
08	base (3 aliases exist)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 008

Figure 107. ESS Specialist: LCU devices table (base devices)

Figure 108 shows the bottom of the same table. The highest device ID (7F) is the first alias for device 00. The device IDs are assigned in descending order in groups of three to the base device addresses in ascending order. The addresses x'20' - x'7F' are the 96 PAV aliases, three per base device, as specified. For the PAV aliases, the Base/Alias column has alias (base device xx).

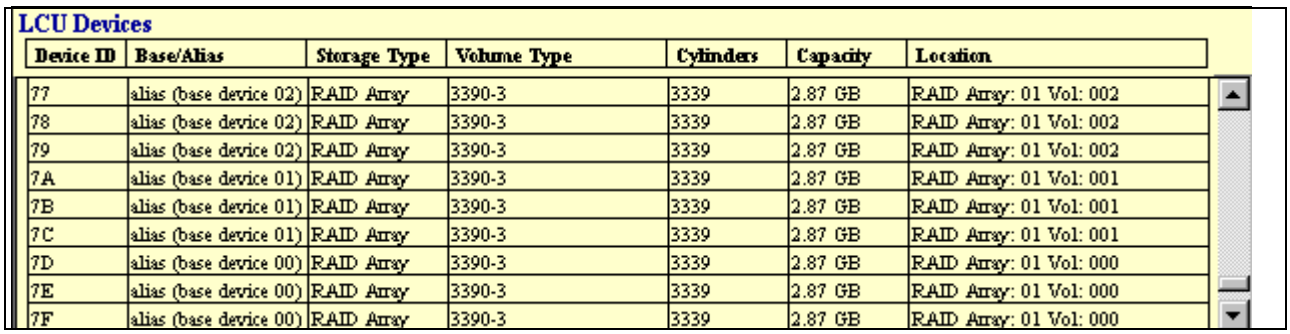

LCU Devices

Device ID	Base/Alias	Storage Type	Volume Type	Cylinders	Capacity	Location
77	alias (base device 02)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 002
78	alias (base device 02)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 002
79	alias (base device 02)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 002
7A	alias (base device 01)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 001
7B	alias (base device 01)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 001
7C	alias (base device 01)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 001
7D	alias (base device 00)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 000
7E	alias (base device 00)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 000
7F	alias (base device 00)	RAID Array	3390-3	3339	2.87 GB	RAID Array: 01 Vol: 000

Figure 108. ESS Specialist: LCU devices table (PAV aliases)

Note that if we specify a greater total number of PAV aliases than 96 (for example, seven aliases for each of the 32 base devices), the total number of addresses in the LCU exceeds the PAV starting address of 128 we chose. The alias addresses are assigned starting at x'7F', decrementing to x'20', then continuing at x'80' and

incrementing to the limit of x'FF'. Device ID 80 in this example is then the sixth alias of base device 13.

Figure 109 shows the Storage Allocation panel, after the ESCON Net icon has been clicked. The one logical disk group in LCU 0 that we configured (cluster 1, DA pair 1, loop A, group 1) is shown. Clicking on this has caused its status to be shown in the information box.

The part indicated in red corresponds to the interleaved volumes that were auto-assigned when we configured the disk group. The green part is the non-interleaved partition. The non-interleaved partition has no Custom Devices allocated.

The volumes in this S/390 group are accessible by any S/390 host through any installed ESCON port in the ESS, provided that the host HCD specifies the connection. For this LCU, the HCD must specify 128 addresses to cover the 32 base devices, and 96 PAV aliases.

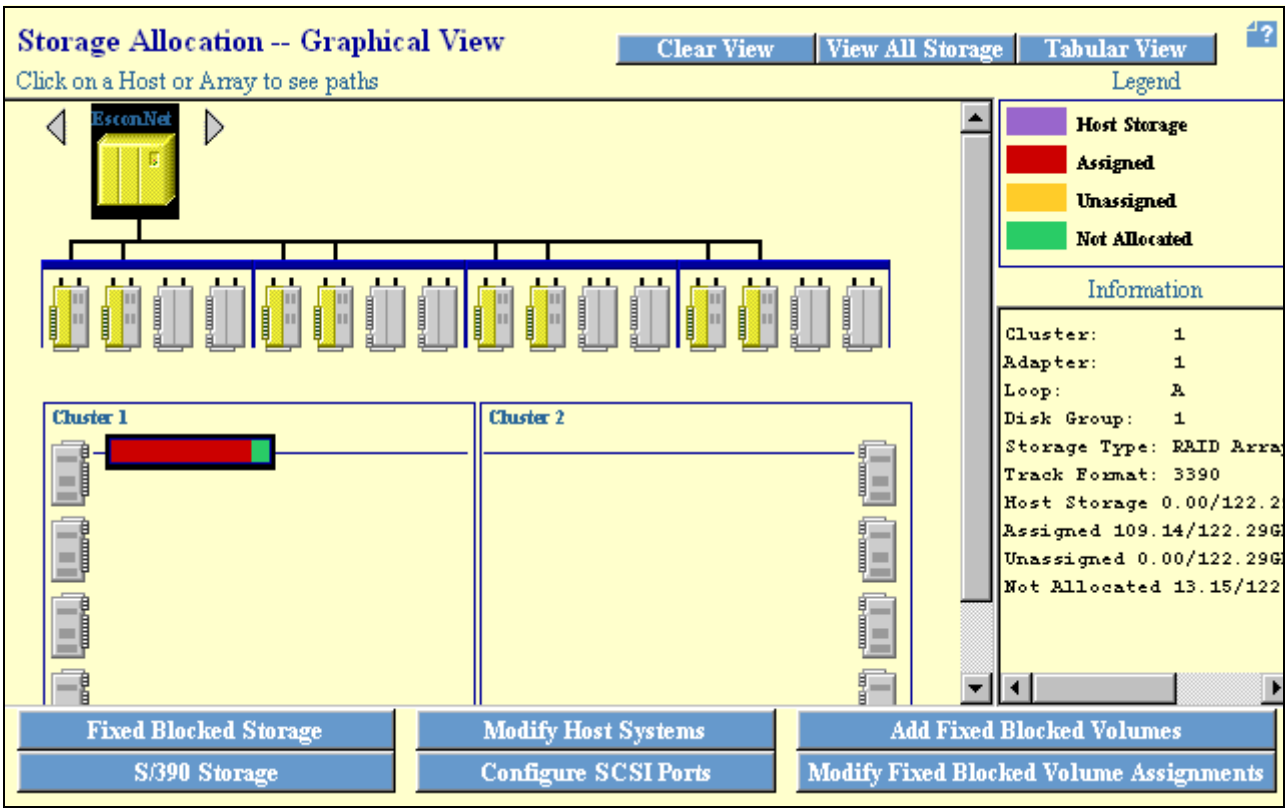

Figure 109. ESS Specialist: Panel showing configured S/390 storage

Chapter 10. Data migration

Data migration is the final installation process, when the data begins to populate the ESS. The data migration from the previous storage subsystems to ESS can be done by using standard host utilities, ESS copy services, or through the IBM migration services offering.

10.1 Migration planning

Migration planning is a key factor to the success of ESS implementation. Today, because business is very dependent on data processing, interruptions are not accepted at all. The environmental changes have to be made completely transparent to the customer. Therefore, storage administrators have to choose the non-disruptive migration methods rather than the disruptive ones. In addition, it does not matter what method was used for the migration, it should guarantee total data integrity.

Migration Steps
ESS is a very scalable subsystem that ranges from 420 GB (eight arrays of 9.1 GB DDMs) to 11.2 Terabytes (48 Arrays of 36.4 GB DDMs). ESS can simulate 16 IBM 3990 Model 6 controllers and it can address up to 4,096 logical volumes. Therefore, data can be migrated from one or several storage subsystems to a single central repository of date (ESS). The whole migration process can take hours or days to be completed and can be done in one or more steps according to the environmental complexity.

Migration Tool
DASD Magic/2 can assist you in making an effective data migration plan to ESS. DASD Magic/2 is a high-level modeling tool of effects of changes to ESS configurations, such as DDMs, SSA loops, LSSs, or different cache size in your workloads.

10.2 Planning data migration for MVS environments

There are two major considerations when planning data migration for MVS environments:

Systems-Managed Storage (SMS)
SMS implements storage policies and handles most of your storage management tasks. The data migration for ESS is greatly simplified in the SMS environment compared to non-SMS. You can create a new storage group for ESS and alter the ACS routines to reflect the changes. You can also add ESS volumes to an existing storage group.

Non-SMS
If your environment is non-SMS, you can migrate data from non-ESS volumes to ESS volumes using standard utilities like DFSMSdss. To avoid changes in JCL, you should copy the volume data by using DFSMSdss COPY with COPYVOLID (see Appendix , “DFSMSdss” on page 182).

10.2.1 Migration methods

Among several migration methods, you can choose one method - or combine more than one method - to migrate data that best fits your environment.

DFSMSdss
DFSMSdss is the faster data mover product that helps you to migrate data between *like* and *unlike* devices.

You can use the `DFSMSdss COPY` command to perform a physical volume copy. Full-volume COPY can move data only between *like* devices of equal or greater capacity (for example, from a single-capacity 3390-3 model to a triple-capacity 3390-9 model). If you specify COPYVOLID, the volume serial number of the source volume is copied to the target volume.

Logical volume copy is a data-set oriented method that allows you to move data between *unlike* devices (for example, from a 3380 to a ESS). To move a volume logically, use the `DFSMSdss COPY DATASET` command.

`DFSMSdss DUMP` and `RESTORE` commands moves data from source volumes to cartridges and then to the target volumes. The advantage of this approach is that you do not need to attach to host all storage subsystems involved in a migration at the same time. You can also use `DUMP/RESTORE` to bring ESS data from storage subsystems installed in different buildings. You can choose to dump and restore physically or logically by using or not using the DSNAME keyword in conjunction with DFSMSdss DUMP and RESTORE command.

Others host utilities
You can have on hand the following alternatives to DFSMSdss to move data sets:

- IDCAMS EXPORT/IMPORT (VSAM)
- IDCAMS REPRO(VSAM, SAM, BDAM)
- IEBCOPY (PDSs - including load module libraries - and PDSEs)
- ICEGENER (SAM) - part of DFSORT
- IEBGENER (SAM)
- Specialized data base utilities (such as for CICS, DB2 or IMS)

DFSMShsm
DFSMShsm invokes DFSMSdss to perform a data migration. So, it is a similar method to the DFSMSdss described above. DFSMShsm has control data sets and journal to help you to control the entire migration process. You can use `DFSMShsm DUMP` or `MIGRATE` commands to move data from M0 volumes.

See the *DFSMS/MVS V1R5 DFSMShsm Storage Administration Guide,* SH21-1076, and *DFSMS/MVS V1R5 DFSMShsm Storage Administration Reference,* SH21-1075.

DFSMShsm ABARS
Although the primary purpose is disaster recovery, you can also use DFSMShsm ABARS to migrate volumes to cartridges and restore then in an ESS. In the Aggregate Group you have to define the data sets you are intend to migrate to the ESS.

For reference, see *DFSMShsm ABARS and Mainstar Solutions,* SG24-5089.

XRC

XRC is a non-disruptive alternative to migrating data to ESS. Migration takes place while the application is still running. Source volumes have to reside on IBM 3990 controller or non-IBM controller with XRC support. The target volumes are those defined on an ESS. In an XRC session, System Data Mover (SDM) manages the process of copying data from source volume to the corresponding target volume. During the copy, updates on a source volume will be made on a target volume in the same order, thus maintaining update sequence consistency.

When all pairs are synchronized, you can stop the application systems, stop the XRC, recover the volumes, and restart the applications using the new volumes. The `XRC XRECOVER` command re-labels the target volumes with the source volume serials.

For more details about XRC, see *Planning for IBM Remote Copy,* SG24-2595.

10.3 Planning Data Migration for VM environments

DFSMS/VM contains services which include a data mover, an automated move process, and in interactive user interface:

- DASD Dump Restore (DDR) is a service utility shipped with VM that you can use to dump data from disk to tape, restore data from tape to disk, and copy data between like disk drive volumes. You cannot use DDR to copy data between disk devices with different track formats.
- `CMDISK` is a DIRMAINT command you can use to move minidisks from any device type supported by VM to any other type.
- `COPYFILE` is a CMS command you can use to copy files or minidisks between devices with the same or different track formats.
- `PTAPE` is a CP command you can use to dump spool files to tape and to load them from tape to disk.

10.4 Planning data migration for VSE environments

You can use several dialogs in the VSE interactive interface to set up the jobs to move data. You can reorganize your data and eliminate space fragmentation by using the backup/restore dialogs:

1. Export and import VSAM files
2. Backup and restore VSAM files
3. Backup and restore ICCF libraries
4. Backup and restore the system history file
5. Backup and restore the system residence library
6. Create a loadable tape with the system residence library and system history file ready to restore

You can also use:

- `VSE/FASTCOPY` to move volumes and files between devices with identical track formats
- `VSE/DITTO` to copy files

- VSE/POWER commands to transfer the SPOOL queue from one device to another
- VSE/VSAM to move any VSAM data set using either REPRO or EXPORT/IMPORT functions.

Other vendors also provide utilities to move data from one device to another.

10.5 IBM migration services

This is the easiest way to migrate data, because IBM will assist you along the complete migration process. In several countries IBM offers a migration service. In addition to the S/390 utilities and ESS copy services, IBM uses the Transparent Data Migration Facility (TDMF) tool to do the data migration from any previous S/390 storage subsystem to ESS. Check with your IBM Sales Representative about migration services for your specific environment and needs.

Chapter 11. Managing and monitoring the ESS

In this chapter we discuss managing and monitoring of the ESS.

11.1 General considerations

The following general items should be checked regarding systems management discipline:

- Have the system management procedures been updated to reflect the new hardware?
- Does a firewall need to be created to protect against unauthorized logins?
- Has a process to maintain the SNMP alert e-mail addresses been established?
- Have you set up a TCP/IP address on your local network for the ESS?
- Has a firewall been setup to keep unauthorized users away from the ESS?
- Have the user IDs and addresses been established for configuring, administering and viewing the ESS?

11.2 RMF

RMF reports regarding ESS performance are almost the same as for 3990. Response, connect, disconnect, PEND and IOSQ times are reported for a corresponding device address. Alias addresses for PAVs are not reported separately, but RMF will report the number of Aliases (RMF terms - exposures) that have been used by a device. For each logical control unit (LCU) RMF collects and reports cache statistics.

Additional information about usage of Alias addresses can be found in SMF record, type 74, subtype 5. RMF does not use this information.

11.3 Logical subsystem resource usage

LISTDATA and SETCACHE are AMS commands used to display logical subsystem resource usage and to control caching functions of IBM storage controllers. Both commands are changed to support ESS.

As documented in this chapter, LISTDATA now generates a new RAID RANK COUNTERS report and a modified SUBSYSTEM COUNTERS report. SETCACHE, on the other hand, is updated to produce messages indicating that ESS does not support the DFSMSdss Dual Copy function or the disabling of Cache/DFW.

11.3.1 IDCAMS LISTDATA

IDCAMS LISTDATA generates reports on activity within all IBM caching models of 3880 Storage Control, 3990 Storage Control with cache, and RAMAC Array Subsystem.

You can now use IDCAMS LISTDATA to request activity information on ESS. You will receive a modified SUBSYSTEM COUNTERS report and a new RAID RANK COUNTERS report when you issue a `LISTDATA COUNTS DEVICE`, `LISTDATA COUNTS SUBSYSTEM`, or `LISTDATA COUNTS ALL` command.

11.3.1.1 IDCAMS SUBSYSTEM COUNTERS report

The SUBSYSTEM COUNTERS report is changed to include the RAID RANK ID of a logical volume. This identifier is required for a cross reference to the new RAID RANK COUNTERS report. Figure 110, Figure 111, and Figure 112 present a SUBSYSTEM COUNTERS report including this identifier.

```
                  2105 STORAGE CONTROL MODEL 01
                         SUBSYSTEM COUNTERS REPORT
                  VOLUME XXXF50 DEVICE ID X'48'
                           SUBSYSTEM ID X'0057'
                                CHANNEL OPERATIONS
                 ......SEARCH/READ...... .............WRITE.........
                      TOTAL  CACHE READ      TOTAL     DASDFW CACHE
REQUESTS
  NORMAL             29439       28924       8268            0
  SEQUENTIAL             0           0       2920            0
  CACHE FAST WRITE       0           0          0          N/A
TOTALS               29439       28924      11188            0
REQUESTS                    CHANNEL OPERATIONS
  INHIBIT CACHE LOADING                        0
  BYPASS CACHE                                81
TRANSFER OPERATIONS                  DASD/CACHE   CACHE/DASD
  NORMAL                                    448            0
  SEQUENTIAL                                  0          N/A
DASD FAST WRITE RETRIES                       0
DEVICE STATUS     CACHING:          ACTIVE
                  DASD FAST WRITE: DEACTIVATED
                  DUPLEX PAIR:      NOT ESTABLISHED
```

Figure 110. IDCAMS subsystem counters report (Part 1)

```
RAID RANK ID X'yy'
                              LEGEND
              SUBSYSTEM COUNTERS LEGEND
VOLUME              - VOLUME SERIAL NUMBER FOR WHICH THE DATA IS GATHERED
DEVICE ID           - CHANNEL CONNECTION ADDRESS OF THE DEVICE ON WHICH
                      THE I/O WAS DONE
SUBSYSTEM ID        - SUBSYSTEM TO WHICH THE DEVICE IS ATTACHED
CHANNEL OPERATIONS - A CHANNEL PROGRAM OR PART OF A CHANNEL PROGRAM
                      BEGIN WITH A LOCATE RECORD OR SEEK COMMAND AND
                      ENDING AT NEXT LOCATE RECORD OR SEEK COMMAND OR AT
                      THE END OF CHANNEL PROGRAM
    SEARCH/READ     - OPERATIONS CONTAINING AT LEAST ONE SEARCH OR READ
                      COMMAND BUT NO WRITE COMMANDS
      TOTAL         - ALL SEARCH/READ OPERATIONS
       CACHE READ   - ALL SEARCH/READ OPERATIONS WHICH DID NOT REQUIRE
                      DATA MOVEMENT TO OR FROM DASD
    WRITE           - OPERATIONS CONTAINING AT LEAST ONE WRITE COMMAND
       TOTAL - ALL WRITE OPERATIONS
       DASDFW       - WRITE OPERATIONS WHICH ARE ADDRESSED TO A DEVICE IN
                      FAST WRITE MODE AND WHICH DO NOT CONTAIN A DEFINE
                      EXTENT COMMAND OR WHICH DO NOT INHIBIT DASD FAST
                      WRITE IN DEFINE EXTENT COMMAND
       CACHE WRITE - ALL CACHE FAST WRITE OR DASD FAST WRITE OPERATIONS
                      WHICH DID NOT REQUIRE CONCURRENT DATA MOVEMENT TO
                      OR FROM DASD
  NORMAL            - OPERATIONS WHICH DO NOT CONTAIN A DEFINE EXTENT
                      COMMAND OR SPECIFY NORMAL CACHE REPLACEMENT IN THE
                      DEFINE EXTENT COMMAND BUT NOT THE CACHE FAST WRITE
                      ATTRIBUTE
  SEQUENTIAL        - OPERATIONS WHICH SPECIFY SEQUENTIAL ACCESS IN THE
                      DEFINE EXTENT COMMAND BUT NOT THE CACHE FAST WRITE
                      ATTRIBUTE
  CACHE FAST WRITE - OPERATIONS WHICH SPECIFY THE CACHE FAST WRITE
                      ATTRIBUTE IN THE DEFINE EXTENT COMMAND BUT NOT THE
                      BYPASS CACHE INHIBIT CACHE LOADING ATTRIBUTES
TOTALS              - TOTAL FOR EACH VERTICAL LINE ITEM
INHIBIT CACHE LOADING - OPERATIONS WHICH SPECIFY THE INHIBIT CACHE
                      LOADING ATTRIBUTE IN THE DEFINE EXTENT COMMAND AND
                      CONTAIN AT LEAST ONE SEARCH, READ OR WRITE COMMAND
BYPASS CACHE        - OPERATIONS WHICH SPECIFY THE BYPASS CACHE ATTRIBUTE
                      IN THE DEFINE EXTENT COMMAND AND CONTAIN AT LEAST
                      ONE SEARCH, READ OR WRITE COMMAND
TRANSFER OPERATIONS - THE NUMBER OF TRACKS TRANSFERRED TO OR FROM DASD
  DASD/CACHE        - THE NUMBER OF TRACKS TRANSFERRED FROM THE DASD TO
                      THE CACHE
```

Figure 111. IDCAMS subsystem counters report (Part 2)

```
IDCAMS SYSTEM SERVICES                                          TIME:
                            LEGEND
              SUBSYSTEM COUNTERS LEGEND
 CACHE/DASD        - THE NUMBER OF TRACKS TRANSFERRED FROM THE CACHE TO
                     DASD
  NORMAL           - DATA MOVEMENT OPEARTIONS EXCLUDING SEQUENTIAL
  SEQUENTIAL       - 'NEXT TRACK' DATA MOVEMENT OPERATIONS IN SEQUENTIAL
                     ACCESS MODE
DASD FAST WRITE RETRIES - NUMBER OF RETRIES REQUIRED BEFORE SPACE WAS
                          AVAILABLE IN THE NVS FOR DASD FAST WRITE DATA
DEVICE STATUS      - STATUS OF THE ADDRESSED DEVICE FOR THE FOLLOWING
  CACHING          - THE DEVICE CACHING STATE, ONE OF THE FOLLOWING
    ACTIVE         - CACHING ACTIVATED
    DEACTIVATION PENDING - TRANSFER OF MODIFIED DATA TO DASD FAILED
    DEACTIVATED    - CACHING DEACTIVATED
  DASD FAST WRITE  - THE STATE OF THE FAST WRITE DATA WHICH WILL BE
                     STORED ON DASD, IS ONE OF THE FOLLOWING
    ACTIVE         - DASD FAST WRITE IS ALLOWED
    DEACTIVATION PENDING - TRANSFER OF MODIFIED DATA TO DASD FAILED
    DEACTIVATED    - DASD FAST WRITE IS DISABLED
  DUPLEX PAIR      - THE SIMPLEX/DUPLEX STATUS OF THE DEVICE WITH THE
                     STATE OF THE PAIR WHEN DUPLEX
    NOT ESTABLISHED - DEVICE IS SIMPLEX
    PRIMARY        - DEVICE IS THE PRIMARY OF A DUPLEX PAIR
    SECONDARY      - DEVICE IS THE SECONDARY OF A DUPLEX PAIR
    ACTIVE         - DUPLEX PAIR AVAILABLE
    PENDING        - COPY TO ESTABLISH A DUPLEX PAIR IN PROGRESS
    SUSPENDED      - SUSPENDED DUPLEX BY HOST COMMAND OR BY SUBSYSTEM
  ADDRESS OF PRIMARY|SECONDARY - CHANNEL CONNECTION ADDRESS OF THE
                     OTHER DEVICE IN THE DUPLEX PAIR
 PINNED DATA EXISTS, FAST WRITE SUSPENDED
                   - ONE OR MORE TRACKS FOR THIS VOLUME ARE PINNED AND
                     BOTH CACHE FAST WRITE AND DASD FAST WRITE ARE
                     SUSPENDED FOR THIS VOLUME
 PINNED DATA, FAST WRITE NOT SUSPENDED
                   - ONE OR MORE TRACKS FOR THIS VOLUME ARE PINNED AND
                     BOTH CACHE FAST WRITE AND DASD FAST WRITE ARE
                     NOT SUSPENDED FOR THIS VOLUME
 DATA IN FAILED NVS - DATA FOR THIS DEVICE IS IN THE FAILED NONVOLATILE
                     STORAGE
****               - INDICATION OF AN INVALID OR UNDEFINED BIT
                     COMBINATION IN THE DEVICE STATUS BYTES
RAID RANK ID
                   - THE RAID RANK ID ASSOCIATED WITH THIS VOLUME
IDCAMS SYSTEM SERVICES                                          TIME
IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
IDC0002I IDCAMS PROCESSING COMPLETE, MAXIMUM CONDITION CODE WAS 0
```

Figure 112. IDCAMS subsystem counters report (Part 3)

11.3.1.2 IDCAMS RAID RANK COUNTERS report

The RAID RANK COUNTERS report is produced to support ESS. Generated after the SUBSYSTEM COUNTERS report, the RAID RANK COUNTERS report contains data on logical, not physical, volumes in the RAID rank. Figure 113 shows a RAID RANK COUNTERS report.

```
IDCAMS SYSTEM SERVICES                                          TIME:

                2105 STORAGE CONTROL  MODEL 01
                      SUBSYSTEM COUNTERS REPORT
                              RAID RANK COUNTERS
                                 SUBSYSTEM ID X'ssid'
RAID RANK ID X'yy'
DEVICE ADAPTER ID X'y'
NUMBER OF HDDS IN RAID RANK:        999
HDD SECTOR SIZE                     9999999999

                                    RAID RANK OPERATIONS
REQUESTS                     I/O REQUESTS   RESP/TIME   SECTOR REQUESTS
  READ                          9999999999  999999.999        999999999
  WRITE                         9999999999  999999.999        999999999
IDCAMS SYSTEM SERVICES                                          TIME:
IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
                             LEGEND
             RAID RANK COUNTERS LEGEND
SUBSYSTEM ID       - SUBSYSTEM TO WHICH THE RAID RANK IS ATTACHED
NUMBER OF HDDS     - NUMBER OF HARD DISK DRIVES IN RAID RANK
HDD SECTOR SIZE    - SIZE IN BYTES OF A PHYSICAL HDD IO REQUEST
RANK OPERATIONS    - OPERATIONS ASSOCIATED WITH A RAID RANK
  I/O REQUESTS     - NUMBER OF I/O REQUESTS
  RESP/TIME        - AVERAGE RESPONSE TIME (MS)
  SECTOR REQUESTS  - NUMBER OF PHYSICAL HDD IO REQUESTS
  READ             - OPERATIONS CONTAINING AT LEAST ONE SEARCH OR READ
                     COMMAND BUT NO WRITE COMMANDS
  WRITE            - OPERATIONS CONTAINING AT LEAST ONE WRITE COMMAND
```

Figure 113. IDCAMS RAID rank counters report

11.3.1.3 IDCAMS SETCACHE

Traditionally, you use the IDCAMS SETCACHE command to allow or prohibit access to IBM caching models of 3880 and 3990 storage controllers. For example, you can use SETCACHE to set Dual Copy and Cache/DFW.

However, ESS Cache/DFW is ON by default, and you are not allowed to modify it. In addition, ESS does not support the Dual Copy function. You will receive a message indicating these changes if you issue any of the Cache/DFW or Dual Copy commands as shown below.

Figure 114 displays the message given in response to the SETCACHE SETSECONDARY command.

```
   SETCACHE SETSECONDARY(D00E) FILE(FILEX)
IDC31562I THE SETSECONDARY PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 114. SETCACHE SETSECONDARY response message

Figure 115 contains the message given in response to the SETCACHE SUSPENDPRIMARY command

```
   SETCACHE SUSPENDPRIMARY   FILE(FILEX)
IDC31562I THE SUSPENDPRIMARY PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 115. SETCACHE SUSPENDPRIMARY response message

Figure 116 presents the message given in response to the SETCACHE SUSPENDSECONDARY command

```
SETCACHE SUSPENDSECONDARY  FILE(FILEX)
IDC31562I THE SUSPENDSECONDARY PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 116. SETCACHE SUSPENDSECONDARY response message

Figure 117 shows the message given in response to the SETCACHE RESETTODUPLEX command

```
   SETCACHE RESETTODUPLEX  FILE(FILEX)
IDC31562I THE RESETTODUPLEX PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 117. SETCACHE RESETTODUPLEX response message

Figure 118 covers the message given in response to the SETCACHE REESTABLISHDUPLEX command.

```
   SETCACHE REESTABLISHDUPLEX(D00E)  FILE(FILEX)
IDC31562I THE REESTABLISHDUPLEX PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 118. SETCACHE REESTABLISHDUPLEX response message

Figure 119 shows the message given in response to the SETCACHE RESETTOSIMPLEX command.

```
   SETCACHE RESETTOSIMPLEX  FILE(FILEX)
IDC31562I THE RESETTOSIMPLEX PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 119. SETCACHE RESETTOSIMPLEX response message

Figure 120 contains the message given in response to the SETCACHE DFW OFF command

```
   SETCACHE DFW OFF  FILE(FILEX)
IDC31562I THE DASDFASTWRITE PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 120. SETCACHE DFW OFF response message

Figure 121 includes the message given in response to the SETCACHE NVS OFF command

```
   SETCACHE NVS OFF  FILE(FILEX)
IDC31562I THE NVS PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 121. SETCACHE NVS OFF response message

Figure 122 presents the message given in response to the SETCACHE DEVICE OFF command

```
 SETCACHE DEVICE OFF  FILE(FILEX)
IDC31562I THE DEVICE PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 122. SETCACHE DEVICE OFF response message

Figure 123 displays the message given in response to the SETCACHE SUBSYSTEM OFF command

```
 SETCACHE SUBSYSTEM OFF  FILE(FILEX)
IDC31562I THE SUBSYSTEM PARAMETER IS NOT AVAILABLE FOR THE SPECIFIED
IDC31562I SUBSYSTEM OR DEVICE
IDC3003I FUNCTION TERMINATED. CONDITION CODE IS 12
```

Figure 123. SETCACHE SUBSYSTEM OFF response message

11.4 DEVSERV command

The DEVSERV command gives you the status of DASD and tape devices. The response is a display of basic status information about a device, a group of devices, or storage control units, and can optionally include a broad range of additional information. With the new parameter, QPAVS, you can display base volumes and their aliases.

11.5 Messages and codes

This section documents new and changed messages and codes in support of ESS.

11.5.1 IEA435I

The existing message IEA435I generated by DFSMS Device Support is updated to identify the following error conditions:

IEA435I	**PHYSICAL DEVICE INCONSISTENT WITH LOGICAL DEFINITION,**
	[PHY=devtype LOG=devtype]
Explanation:	Mismatches between the hardware configuration and device definition in HCD exist.
Problem Determination:	One of the following occurred:
PHY=*devtype* **LOG=***devtype*	During the processing of VARY,ONLINE command, system services determined that the physical device type was not compatible with the logical device type defined in the I/O configuration. The physical (PHY) and logical (LOG) device type are provided.
Module:	DFMSdfp
System Programmer Response:	Correct the HCD device definitions to match the physical device type attached to the system.

11.5.2 IEE459I

The existing message IEE459I is expanded to include the new `DEVSERV QPAVS` command.

11.6 ESS Expert

The IBM Enterprise Storage Server Expert (Expert) is a software tool which provides storage resource management functions for IBM storage servers. For more information see 6.1, “Using ESS Expert” on page 93.

11.7 Simple Network Management Protocol (SNMP)

StorWatch adopts industry standard interface SNMP as the management protocol for storage systems. The ESS generates SNMP traps and supports a read-only Management Information Base (MIB). The ESS generates both generic and product-specific SNMP traps. Product-specific traps provide information on problems that are detected by the ESS that require corrective action. For details, see Chapter 4, “Availability” on page 39.

Chapter 12. ESS Specialist Copy Services for S/390

Copy Services provides a Web-based interface for establishing and managing PPRC and FlashCopy. PPRC operates at the volume level from one logical control unit (LCU), to the volume level within another LCU. FlashCopy also operates at the volume level, however, the source and target volumes are within the same LCU.

PPRC and Flashcopy are used as a disaster recovery tool, data backup tool, and for creation of test data or for data migration.

12.1 Copy Services considerations

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

There are a few restrictions to consider:

- Only one FlashCopy at a time can be active on a volume, however, you can perform a PPRC concurrently with FlashCopy on the same volume.
- The primary and secondary volumes must reside within an ESS, for both PPRC and FlashCopy.
- If you intend to manage PPRC using the Copy Services component of the ESS Specialist, then Ethernet, TCP/IP connectivity is needed between the two participating ESS subsystems (primary and secondary) and the Web browser initiating and managing the PPRC activities.
- When using Copy Services, you are required to nominate via the CONFIGURATION panel, one ESS as the Copy Services Primary Server which is the central place for collection of all information. A second ESS can be nominated as the backup server. You will need to carefully plan which ESS you would nominate in which role to ensure maximum availability.
- The source and target volumes must be the same size.

You will also need to check that the CRIT=YES - Light or Heavy options have been set correctly in the ESS VPD by your hardware specialist. See 2.3.1.4, “VPD settings” on page 19.

You will need to have the hardware engineer check that all the ESS cluster hostnames were added to the hostname list during installation. Otherwise the cluster hostname will not be found when defining the Copy Services primary.

12.2 PPRC

This section describes the configuration rules and invocation processes.

12.2.1 Rules for configuring PPRC links

The following are the PPRC configuration rules for establishing Escon connections between two ESS subsystems.

- Up to eight Escon links are supported between two ESS subsystems.

- A primary ESS can be connected via Escon links to up to four secondary ESS subsystems.
- A secondary ESS subsystem can be connected to any number of primary ESS subsystems, limited by the number of Escon links available.
- PPRC links are unidirectional, because the ESCON port at the primary ESS is reconfigured to act like an ESCON channel in a host S/390 processor. The primary ESCON port is dedicated to PPRC, and cannot connect to an S/390 host while PPRC paths are established.
- If an ESCON port on the secondary ESS is connected to an ESCON director, it can connect to either a PPRC primary ESS subsystem or an S/390 host. It does not need to be dedicated to PPRC.
- The ESCON protocol has been streamlined with less handshaking and larger frames transmitted between ESS. This has enabled the PPRC Escon links to be extended to up to 103 Klm.

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

12.2.2 How to invoke PPRC

To set up the paths and pairs, you need to know the last five digits of the serial number and logical subsystem number of the primary and secondary resources. These details can be found using the ESS Specialist or the `OS/390 DEVSERV QDASD` command.

To invoke PPRC you need to perform the following tasks from the Copy Services main menu:

1. Using the PATHS panel, establish a logical path between the primary ESS LCU and host adapter, and the target LCU and its host adapter.
2. Use the VOLUMES panel to find and select the source and then target volume PPRC pairs and Establish, Suspend or Terminate the data transfer. You can optionally establish both paths and volume pairs from this panel.
3. Alternatively, you can select the CONTROL UNIT panel to initiate or remove PPRC relationships between all volumes on an LCU.
4. You can save previously defined PPRC path and pair definitions as tasks. Using the TASKS panel, you can select and run the pre-saved set of tasks.

During steps 1, 2 and 3, once you have selected the resources that you are working with, a right mouse click on the target resource starts a wizard to guide you through the selection of the appropriate PPRC functions.

To assist you to remove unrequired resource information from the panels, a FILTER button is provided to enable display of selected resources. For example, show only:

1. S/390 or Open volumes
2. Source or target volumes
3. Physical or logical ESS

Each panel also has an INFORMATION button which will display PPRC status and other general information about the selected LCUs, volumes or paths.

Detailed information regarding each panel described above can be found in the *Web Interface Users Guide*, SC26-7346.

12.3 FlashCopy

This section describes the configuration rules and invocation processes.

12.3.1 How to Invoke FlashCopy

FlashCopy can be invoked from OS/390 by DFSMSdss COPY FULL (see 8.2.7, “Copy services” on page 147) or by using the Copy Services panels.

To invoke FlashCopy with Copy Services, use the VOLUMES panel to find and select the source and target volumes. Once presented with the wizard window, select Establish or Withdraw the FlashCopy pairs.

As with PPRC you can save your FlashCopy definitions as tasks and run them at any time by using the TASKS panel. You can also FILTER your displays and use the INFORMATION button.

Once a relationship has been established between a FlashCopy source and target volume, a background task commences that copies the entire source volume to the target. If initiating the FlashCopy using the ESS Specialist, you can suppress this copy task by specifying a NOCOPY option.

If NOCOPY has been specified, any data about to be updated on the source volume is first copied to the target volume. Hence, the target volume only contains pre-updated data, not a complete volume copy. The T0 copy of the source is still available for use as long as the source target relationship exists.

This relationship must be terminated (WITHDRAW) using the ESS Specialist. If NOCOPY was not specified, the relationship ends automatically once the background source copy has been completed.

Part 3. Implementation in the Open Systems environment

This Part of the book discusses the processes for implementing the ESS in an Open systems environment using the ESS Specialist. As part of the logical configuration process, the ESS provides the capability to use configuration worksheets, which can be filled in by you and given to the IBM CE at installation time. This information will be translated to provide you with a logical configuration to match your requirements. The information in this book does not mention this process, but instead details the work necessary to implement the ESS manually.

For more information on the initial configuration process, please refer to Appendix A, “ESS Configuration Planner” on page 253.

Chapter 13. Open Systems Software support

This chapter covers the software that you require when attaching an SCSI host system to the ESS. The list of supported systems, shown in Figure 124, is expected to be expanded and changed over time. IBM will consider supporting other UNIX and Intel based servers. Please consult with your local IBM Storage Specialist or Business Partner to make arrangements for this support. For the latest news visit the Web site:

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

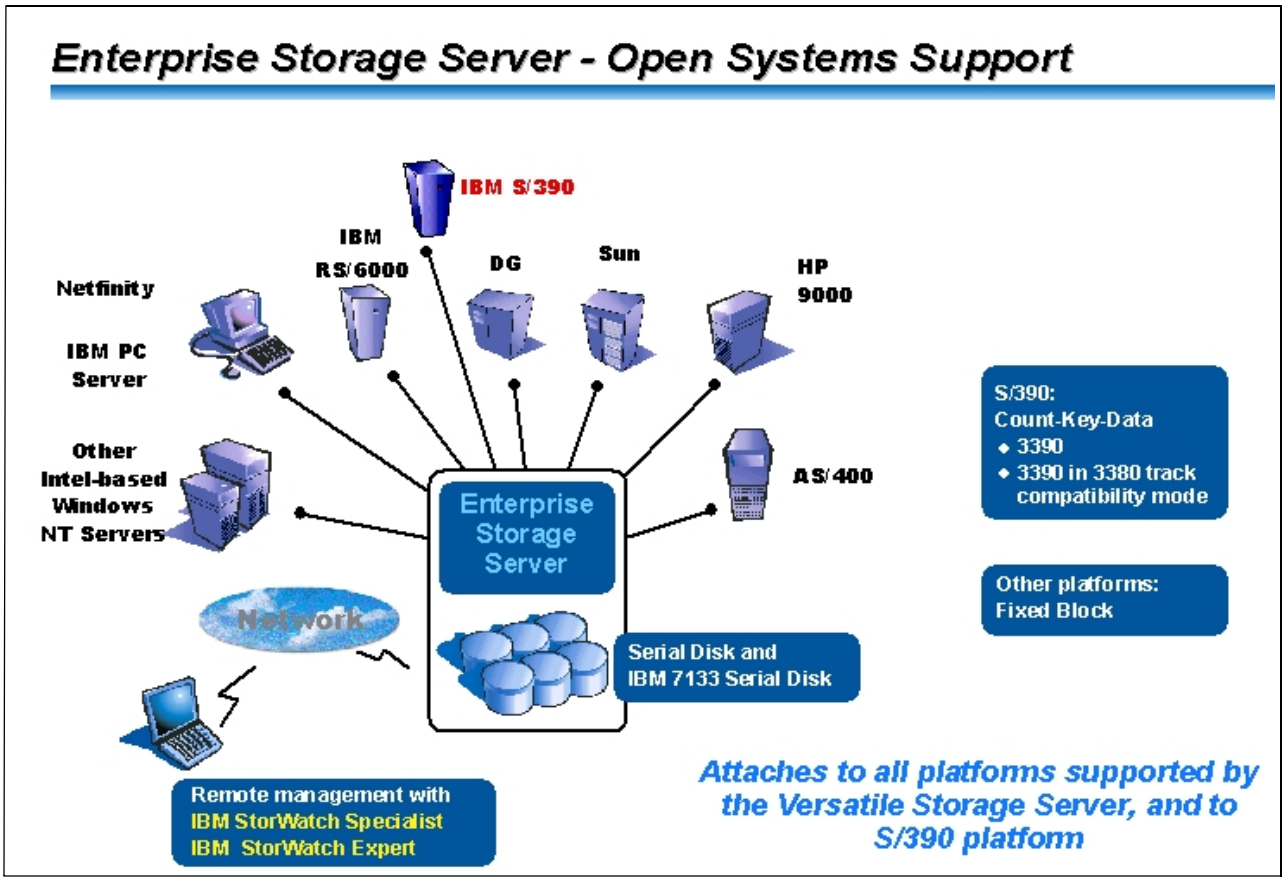

Figure 124. ESS - Open Systems Support

13.1 General considerations

During the preinstallation process it is strongly recommended that you check with assistance your IBM Systems Support Center for assistance with the following general and specific Open Systems software items

13.1.1 General items

- Has the QBUCKET information been reviewed for the ESS and checked for software requirements?
- Have the minimum software pre-requisites been met?
- Has all the recommended software maintenance been installed and tested?

- Have the Independent Software Vendors (ISV) been contacted regarding their support of the ESS?
- Have any required ISV fixes been installed and tested?

13.1.2 Specific items

13.1.2.1 ESS

- Are there any considerations for volume manager software (Veritas VxVM or HP SAM or Sun Solaris DiskSuite)?

13.1.2.2 Data General AViiON

- Are you at the minimum level of DG/UX to support the ESS?
- Are any fixes required for the DG/UX software?
- Are the installed software products at levels which support ESS? You should check with your software vendors.

13.1.2.3 Hewlett Packard 9000

- Are you at the minimum level of HP/UX to support the ESS?
- Are any fixes required for the HP/UX software?
- Are the installed software products at levels which support ESS? You should check with your software vendors and also with Hewlett Packard.

13.1.2.4 IBM AS/400

- Are you at the minimum level of OS/400 to support the ESS?
- Are any PTFs needed for OS/400?
- Are the installed software products at levels which support ESS?

13.1.2.5 IBM Netfinity server

- Do you have Windows NT installed?
- Has Service Pack 3 been installed too?
- Are the installed software products at levels which support ESS? You should check with your software vendors.
- If the IBM Data Path Optimizer will be installed, has Windows NT Service Pack 4 been installed too

13.1.2.6 IBM PC server

- Do you have Windows NT installed?
- Has Service Pack 3 been installed too?
- Are the installed software products at levels which support ESS? You should check with your software vendors.
- If the IBM Data Path Optimizer will be installed, has Windows NT Service Pack 4 been installed too?

13.1.2.7 IBM RISC System/6000

- Are you at the minimum level of AIX/6000 to support the ESS?
- Are any PTFs needed for AIX/6000?

- Are the installed software products at levels which support ESS? You should check with your software vendors.
- If the IBM Data Path Optimizer will be installed, has Windows NT Service Pack 4 been installed too?

13.1.2.8 IBM RISC System/6000 SP

- Are you at the minimum level of AIX/6000 to support the ESS?
- Are any PTFs needed for AIX/6000?
- Are the installed software products at levels which support ESS? You should check with your software vendors.

13.1.2.9 Intel-based PC server

- Do you have Windows NT installed?
- Has Service Pack 3 been installed too?
- Are the installed software products at levels which support ESS? You should check with your software vendors.
- If the IBM Data Path Optimizer will be installed, has Windows NT Service Pack 4 been installed too?

13.1.2.10 Sun SPARCserver/SPARCcenter

- Are you at the minimum level of Solaris to support the ESS?
- Are any fixes needed for Solaris?
- Are the installed software products at levels which support ESS? You should check with your software vendors.

13.1.2.11 Sun Ultra Enterprise Series

- Are you at the minimum level of Solaris to support the ESS?
- Are any fixes needed for Solaris?
- Are the installed software products at levels which support ESS? You should check with your software vendors.

13.2 AS/400

The ESS support is provided in OS/400 Version 3 Release 1 and above. The OS/400 support requires additional PTFs. Table 26 and Table 27 summarizes the support level for OS/400.

Table 26. AS/400 support

Types	Models	SCSI Adapters	Operating Systems
Advanced Series 9406	300, 310, 320	FC 6501	OS/400 - V3R1, V3R2
Advanced Series 9406	500, 510, 530, 720, 730, 740	FC 6501	OS/400 - V3R6, V3R7, V4R1, V4R2, V4R3, V4R4
e-Series 9406	620, 640, 650, S20, S30, S40	FC 6501	OS/400 - V4R1, V4R2, V4R3, V4R4

Table 27. OS/400 - Required PTFs

OS/400 Version and Release	PTF
V3R1	SF44131
V3R2	SF44132
V3R6	SF44126
V3R7	SF44127
V4R1	SF44113
V4R1.4	SF44475
V4R2	SF44114

13.3 Data general

The Data General DG/UX 4.2 supports ESS. Prior to installation, you should contact IBM System Support Center for related PTFs. Table 28 summarize the current support.

Table 28. Data general support

Type	Models	SCSI Adapters	Operating System
AViiON	4900, 5000	Adaptec AHA-2944UW AHA-4944W	DG/UX - 4.2

13.4 DEC

DEC will support the ESS. Table 29 shows the support structure.

Table 29. DEC support

Type	Models	SCSI Adapter	Operating Systems
DEC	2100, 4400, 8400	KZPBA-CB	DEC VMS - 7.2 DEC UNIX - 4.0D, 4.0E

13.5 HP 9000-800

The HP Unix 10.20 and above supports ESS. IBM recommends that corresponding General Release Patch Bundle is installed on your HP/UX operating system prior to attachment of the ESS. For details, contact IBM Systems Support Center. Table 30 summarizes the current support.

Table 30. Hewlett Packard support

Models	SCSI Adapters	Operating Systems
D-class	HP-HSC A4107A	HP/UX - 10.20, 10.30, 11.00
K-class	HP-HSC A2869A	HP/UX - 10.20, 10.30, 11.00
E, G, H, I, K, T	HP-PB A28696A	HP/UX - 10.20, 10.30, 11.00
V Series	PCI A4800A	HP/UX - 10.20, 10.30, 11.00

13.6 Intel Based-PC servers

Table 31 summarizes Intel Based-PC Server support for ESS.

Table 31. Intel Based-PC servers support

Type	SCSI	Operating Systems
Pentium Pro or later processors	Adapters	Microsoft Windows NT
200 MHz processor or faster 128 MB memory or greater	Adaptec: AHA-2944UW Buslogic: BT-958D Mylex BT-958D QLogic: QLA1240D Symbios SYM8751D	Windows NT Server 4.0 (with Service Pack 3 or greater), Windows NT Server 4.0, Enterprise Edition (with Service Pack 3 or greater

Adaptec adapter AHA-2944UW is used on Netfinity Servers (IBM Part Number 59H3900).

IBM Data Path Optimizer is supported on Windows NT Server 4.0, Enterprise Edition (Service Pack 3 or 4).

13.7 Novell Netware

Table 32 shows the prerequisites for Novell Netware.

Table 32. Novell Netware support

Types and Models	SCSI Adapters	Operating Systems
Systems which support the specified adapter and operating systems	Adaptec AHA-2944UW Symbios SYM8751D	Novell Netware - 4.2 and 5.0

13.8 RS/6000 and RS/6000 SP

The AIX 4.2.1, 4.3 and above supports ESS. Table 33 summarize RS/6000 and RS/6000 SP support.

Table 33. RS/6000 and RS/6000 SP support

Types	Models	SCSI Adapters	Operating Systems
7012	G40, 397	FC 2412	AIX - 4.2.1, 4.3 and above
7013	590, 591, 595, 59H, J30, J40, J50	FC 2412	AIX - 4.2.1, 4.3 and above
7015	R20, R24, R30, R40, R50, 99J, 99K	FC 2412	AIX - 4.2.1, 4.3 and above
7024	E20, E30	FC 6207	AIX - 4.2.1, 4.3 and above
7025	F30, F40, F50	FC 6207	AIX - 4.2.1, 4.3 and above
7026	H10, H50, H70	FC 6207	AIX - 4.2.1, 4.3 and above
7017	S70 (+equivalents)	FC 6207	AIX - 4.3 and above
9076	202, 203, 204, 205, 206, 207, 208, 2A2, 2A3, 2A4, 2A5, 2A7, 2A8, 302, 303, 304, 305, 306, 307, 308, 3A2, 3A3, 3A4, 3A5, 3A7, 3A8, 3B2, 3B3, 3B4, 3B5, 3B7, 3B8, 402, 403, 404, 405, 406, 407, 408	FC 2412	AIX - 4.2.1, 4.3 and above PSSP - 2.2, 2.3 and 2.4
9076	500, 550, 50H, 55H	FC 6207, 6209	AIX - 4.2.1, 4.3 and above PSSP - 2.4

AIX support requires additional PTFs as follows:

- AIX 4.2.1 - PTF IX62304
- AIX 4.3.2 - No PTFs Required

13.8.1 HACMP Support

- HACMP 4.2.2 on:
  - AIX 4.2.1
  - AIX 4.3.2
  - **Note:** If you are using CRM feature with HACMP 4.2.2, AIX PTF U458552 is required on AIX 4.2.1 and 4.3.1
- HACMP 4.3 on:
  - AIX 4.3.2

13.8.2 IBM Data Path Optimizer

IBM Data Path Optimizer is supported on:

- AIX 4.2.1
- AIX 4.3.2

13.9 Sun

The Sun Solaris 2.5.1 and above supports ESS. Table 34 summarizes Sun support.

Table 34. Sun support

Types	Models	SCSI Adapters	Operating Systems
SPARCserver	1000, 1000E	X1062A, X1065A	Solaris - 2.5.1, 2.6, 2.7
SPARCcenter	2000, 2000E	X1062A, X1065A	Solaris - 2.5.1, 2.6, 2.7
Ultra Enterprise	3000, 3500, 4000, 4500, 5000, 6000, 10000	X1062A, X1065A	Solaris - 2.5.1, 2.6, 2.7
Ultra 2		X1062A, X1065A	Solaris - 2.5.1, 2.6, 2.7
Ultra	30, 60, 450, 5500, 6500	X6541A	Solaris - 2.5.1, 2.6, 2.7

Note: Solaris 2.7 is also called Solaris 7.

Sun adapters types are as follows:

- X1062A: Sbus, SCSI-2 Fast/Wide Differential
- X1065A: Sbus - Ultra SCSI
- X6541A: PCI - Ultra SCSI-2, Fast/Wide Differential Dual Channel

Chapter 14. Open Systems host setup tasks

The ESS is designed to handle open systems hosts such as the RS/6000 with AIX, AS/400 with OS/400, PC Servers with Windows NT, Netware, and many other UNIX based systems. All of these systems will need SCSI generic devices, or specific disk emulations so that they can address the ESS disk storage.

14.1 General

It is recommended that before commencing these tasks the internal ESS logical setup, or a logical configuration plan, be completed since many of the parameter values required to complete the host definition are dependent on the ESS internal settings.

You should ensure that your definitions match the ESS internal addressing for the target ID and LUNs. The values are assigned by the ESS Specialist during internal logical configuration and can be viewed in the Storage Allocation - Tabular view, see Figure 40 on page 87.

It may be an advantage to have network access to the ESS and an ESS Specialist User ID with View authority so you can more easily cross check the ESS internal configuration settings with your I/O definitions.

14.2 AIX

The RS/6000 with AIX supports up to 32 LUNs per target ID for the Ultra wide SCSI adapters, and a maximum of eight LUNs with SCSI-2 fast wide adapters.

The AIX operating system contains entries in its object distribution manager database to identify the ESS. However, the AIX operating system accesses the ESS through its generic SCSI DDM and, therefore, sees the ESS LVs as a hdisk.

Once the ESS has been connected to your processor, run the 2105inst script, provided on a CD-ROM with the ESS. On successful completion of the script, you can run the configuration manager to detect the devices and add the hdisks with type IBM 2105B01.

If you connect the ESS to your processor and reboot your AIX system before you have run the 2105inst script, the ESS devices will be detected and listed as hdisk type of Other devices. You will need to delete these disks and then run 2105inst.

Once the hdisks are available, you can put a file system on the ESS LUNs.

14.2.1 Copy Services

14.2.1.1 FlashCopy

AIX marks disks with a Physical Volume ID (PVID). FlashCopy creates an identical copy of a volume, including the PVID. You will need to change the PVID of the target or copy volume if you want to use this volume from the same host.

This can be achieved by first clearing the PVID and then assigning a new unique one. For example:

chdev -l disk -a pv=clear

chdev -l disk -a pv=yes

14.3 UNIX

Data General and HP 9000 systems support up to eight LUNs, 0 to 7.

SUN systems with SOLARIS 2.5.1 and 2.6 will support up to eight LUNs per target, 0 to 7. SUN Ultra B supports up to 32 LUNs, 0 to 31. You should check each individual system for their specific support.

Once the ESS has been connected to your processor, run the 2105inst script, provided on a CD-ROM with the ESS. On successful completion of the script, you can run the configuration manager to detect the devices and add the hdisks with type IBM 2105B01.

14.4 OS/400

The AS/400 supports only six target IDs and eight LUNs ranging from 0 to 7.

The AS/400 target address is always six.

OS/400 sees the ESS Logical Volumes as 9337-580/590/5AC/5BC emulated devices of 4.2, 8.6, 16 or 32 GB.

See Chapter 3 in *Host Systems Attachment Guide,* SC26-7296*,* for more details*.*

14.5 Windows NT

PC Servers with Windows NT 4.0 will support up to eight LUNs per target.

Before attaching a host with Windows NT, set the basic input/output system (BIOS) for the SCSI adapters to “disabled”. This will ensure optimum performance is achieved. See Chapter 2 in the *Host Systems Attachment Guide,* SC26-7296, for more details*.*

Chapter 15. ESS configuration for fixed block storage

In this chapter we describe the procedures using ESS Specialist to configure fixed block storage in the ESS for use in the Open Systems environment. The ESS Specialist may be used to do the complete storage configuration process for an ESS, defining S/390 and fixed block (FB) storage, or it may be used to modify an existing configuration. For example, Specialist might be used to define hosts and add volumes for one of the standard array configurations that can be ordered for a new ESS.

15.1 Introduction

ESS Specialist is similar to the VS Specialist used to configure FB storage in the earlier VSS, but it has additional capability to facilitate configuration of features not available in the VSS:

- S/390 storage (CKD volumes accessible from the ESCON host adapters)
- Non-RAID storage (JBODs)

On the ESS Specialist Storage Allocation panel, two additional buttons appear:

Fixed block storage To configure RAID and non-RAID FB storage

S/390 storage This displays configured S/390 storage in tabular form, and leads to three further panels that are used to configure LCUs, RAID and non-RAID S/390 storage, and CKD volumes.

Refer to Chapter 5, “Enterprise Storage Server Specialist (ESS Specialist)” on page 69 for an introduction to the ESS Specialist.

Refer also to 3.8.1.1, “ESS Specialist Web client requirements” on page 32.

Here are some useful documents:

- *The IBM Enterprise Storage Server Web Interface User’s Guide*, SC26-7346

15.1.1 Logical subsystem (LSS)

The ESS uses the concept of the Logical Subsystem to internally manage logical volumes (LVs). A fixed block (FB) logical volume in the ESS is known to host systems as a LUN (logical unit number) associated with a Target ID in an SCSI host adapter port. This is just like any other SCSI disk storage. Open host systems do not need to know about logical subsystems in the ESS. However, the internal organization of the ESS does impose some rules on how logical devices are configured.

In the ESS, logical volumes reside in RAID arrays, or non-RAID disks on SSA loops. Each SSA loop has from two to six logical disk groups of eight drives (not to be confused with the physical drive bays, or 8-packs in which the disks are installed). There are two initiators (Device Adapters) on each SSA loop, one for each storage cluster. Each DA manages half of the logical disk groups on a loop. Two loops are connected to each DA. The DA is associated with an LSS managed by the cluster. Refer to Figure 125 on page 212. The eight logical subsystems for FB storage in the ESS are numbered from 10 to 17.

For example, in Figure 125, DA 2 manages up to three logical disk groups on DA pair 2, loop A, and up to three disk groups on loop B. Those disk groups, and the logical volumes on them are associated with LSS 12, in cluster 1. The disk groups on DA pair 2, loops A and B that are managed by DA 3 are in LSS 13, managed by cluster 2.

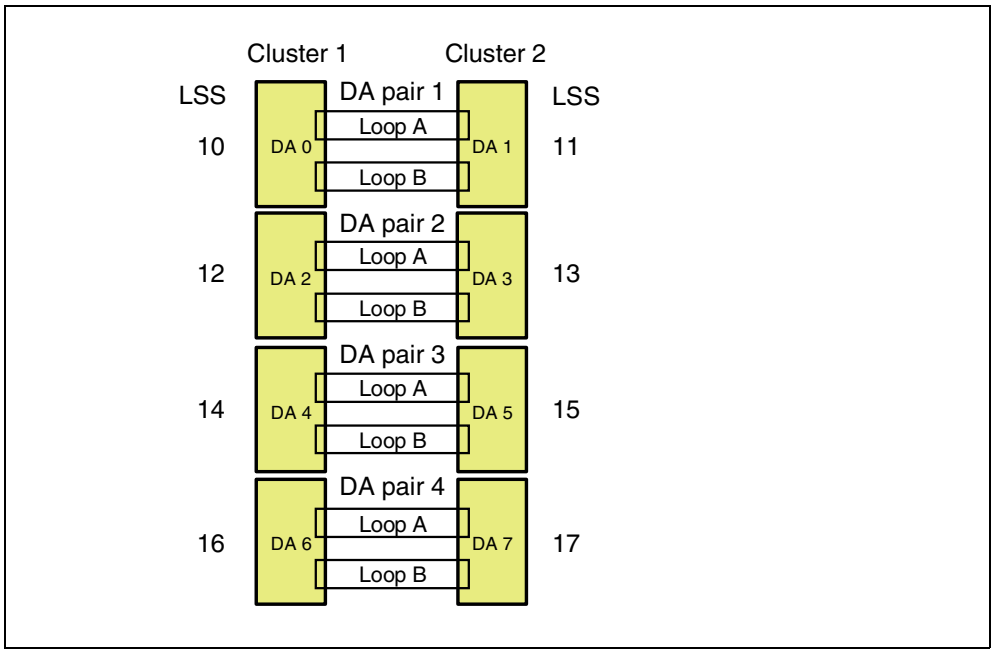

Figure 125. Relationship between SCSI LSSs and SSA loops

If the ESS has S/390 storage configured as well as FB storage, there is a second, independent set of LSSs that manages the S/390 logical volumes. These are numbered 00 to 07. See Figure 86 on page 154. Each of them corresponds to one S/390 Logical Control Unit (LCU). A rank (RAID array or JBOD) may be formatted for FB and managed by an LSS, or for S/390 and managed by an LCU. Unlike an FB LSS, an S/390 LCU must be defined to the ESS, and also to any S/390 hosts that use logical volumes in the LCU.

15.1.1.1 Summary of LSS attributes:

- An LSS is uniquely associated with one storage cluster, one DA, and half of the logical disk groups on each loop. The FB ranks in those disk groups, and the logical volumes on the FB ranks, are managed by the LSS.
- An LSS may have from zero up to six logical disk groups.
- There can be up to 16 FB logical subsystems in an ESS (and from zero to eight S/390 LCUs as well).
- An LSS can have up to 256 logical volumes (LVs) assigned.
- An LSS may have as few as one LV assigned. That is, a single LV may contain virtually all the FB capacity in an entire array, and this array may occupy the only logical disk group in the LSS.
- The even numbered LSS, the even numbered DAs, and the odd numbered disk groups on each SSA loop are managed by cluster 1. The alternate resources are managed by cluster 2.

Open host systems do not need to be aware of the internal logical structure of the ESS, but logical subsystems are visible to the ESS user in several ways:

- When managing PPRC using the Copy Services function of the ESS Specialist, you can specify a group of volumes either by ESS subsystem, or by LSS.
- Using Flashcopy, volumes can be copied only to other volumes in the same LSS.
- Each target ID in an ESS SCSI port maps to one LSS. That is, all addressable logical unit numbers (LUNs) under a target ID map to logical volumes in the same LSS. This means that when volumes from more than one LSS are added to an SCSI port, a separate target ID in the port is assigned for each LSS.

15.1.2 FB logical volumes and SCSI LUNs

A fixed block logical volume is often referred to as a LUN, because it maps to a single LUN in any SCSI port to which it has been added.

- Unlike S/390 volumes, fixed block volumes must be configured to an ESS SCSI port before a host can address them.
- A fixed block volume can be configured to two or more SCSI ports. This allows multiple paths from a host to the volume, and it allows hosts on different SCSI buses to address the volume.
- An LSS can appear as only one target ID at any SCSI port. A volume can appear as only one target/LUN per port.
- The volume may appear as a different Target/LUN in each SCSI port to which it is added. Targets and LUNS are assigned to a port in the order that volumes are added. The target/LUN assigned for a volume depends on what has been previously assigned to that port for the same, and other logical subsystems.
- One target ID on the SCSI bus must be reserved for a host initiator. Additional targets may be used by other hosts or devices on the bus. An ESS SCSI port can have up to 15 addressable target IDs. Each target ID maps to a single LSS. Any one SCSI port cannot address volumes from more logical subsystems than the number of target IDs available.
- Because ESS supports the SCSI 3 protocol, each addressable target ID in a SCSI port can support up to 64 LUNs. This allows any one ESS SCSI port to support up to 960 LUNs, of the maximum of 4096 logical volumes that can be configured in an ESS.
- UNIX and NT hosts may not support 64 LUNs per target. The number of LUNs addressable per target is limited to the number supported by the host operating system (typically 1, 8, or 32).

15.2 Fixed block storage configuration

In this section, we describe the procedures needed and the ESS Specialist panels used to configure FB storage.

15.2.1 Task sequence

We recommend that you do the necessary steps for configuration of FB volumes in the order indicated in Figure 126. On the ESS Specialist Storage Allocation

panel there are six buttons in the lower part of the panel. The S/390 Storage button is not needed here. It is used only to configure S/390 storage.

The functions you will use are:

Configure fixed block disk groups	Configure the fixed block RAID and non-RAID storage you require
Modify host systems	Define the UNIX, Windows NT, and AS/400 hosts that will address logical volumes in the ESS
Configure SCSI ports	Associate each SCSI port with the hosts it will be cabled to, and define for each port the target IDs on the SCSI bus that are already used
Add fixed block volumes	Create logical volumes (LUNs) and associate them with SCSI ports
Modify fixed block volume assignments	Associate logical volumes with additional SCSI ports

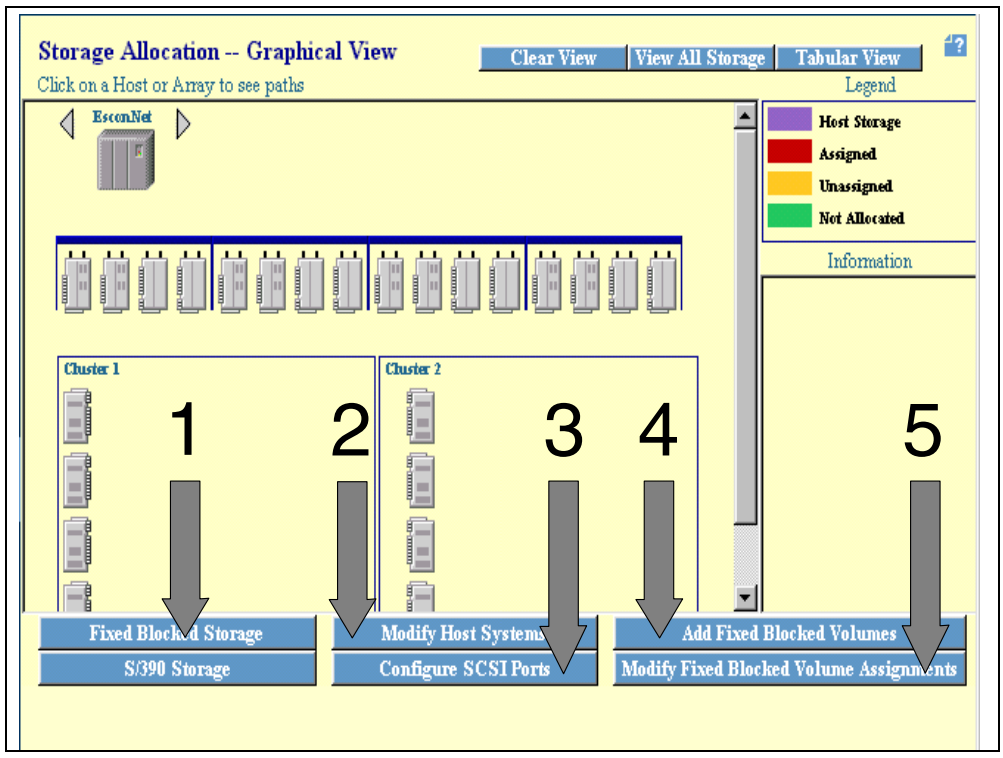

Figure 126. Recommended configuration sequence for FB storage

15.2.2 Connecting to ESS specialist

Refer to 9.2, “Start here” on page 159 for information on how to connect to the ESS Specialist, and select the Storage Allocation panel.

On the Storage Allocation panel, if you click on **View All Storage**, groups may appear even though the Configure fixed block Disk Groups panel shows no defined disk groups. The disk groups you see are configured as S/390 storage.

On the Storage Allocation panel, click on **Fixed Block Storage**.

Note: When using ESS Specialist, ensure that the CE terminal is not logged on via the serial port. If you see a message that one or more components are in service mode, this is probably the cause.

15.2.3 Fixed block storage

The Configure fixed block Disk Groups panel appears as shown in Figure 127 on page 216. In this panel is a table showing all the available disk groups in the ESS. Any disk groups that are already configured as S/390 are not shown. Use the scroll bar to view the entire table. You can select a disk group on SSA loop A, or loop B on one of the four device adapter (DA) pairs in the ESS. Each SSA loop can have up to six disk groups.

For the selected disk group, choose one of:

RAID Array Creates a RAID 5 array (6+P+S, or 7+P) using eight disk drives on the specified SSA loop

Non-RAID Defines a group of disk drives as individual, non-RAID ranks

Undefined Undefines an existing group

When you select RAID, a RAID 5 array is defined. The Storage Type and Track Format fields in the table are RAID Array and Fixed Block. The Capacity field has the formatted array capacity. In Figure 127, adapter pair 3, loop A, group 1 has a capacity of 52.62 GB. This is the formatted capacity of a 6+P RAID array of 9.1 GB disk drives.

If you select non-RAID, the row in the table expands to several rows, one for each of the disk drives in the group. The Storage Type and Track Format fields in the table are non-RAID and None. The Capacity field has the unformatted capacity of a single disk drive. In Figure 127, adapter pair 3, loop a, group 2 appears as eight non-RAID disks, each of 9.1 GB capacity.

Non-RAID disks can then be individually selected, and the track format changed to fixed block. A non-RAID disk group may contain both FB and CKD disks in any combination. If a disk group is defined as non-RAID using this panel and any disks in the group are not defined as FB track format, these disks are available to be defined as CKD track format using the S/390 Storage panels. Similarly, if a non-RAID disk group is defined using the S/390 panels, any disks not defined as 3390 appear in the Configure fixed block Disk Groups panel and are available to be defined as FB.

In Figure 127 on page 216, none of the non-RAID disks has yet been defined as fixed block.

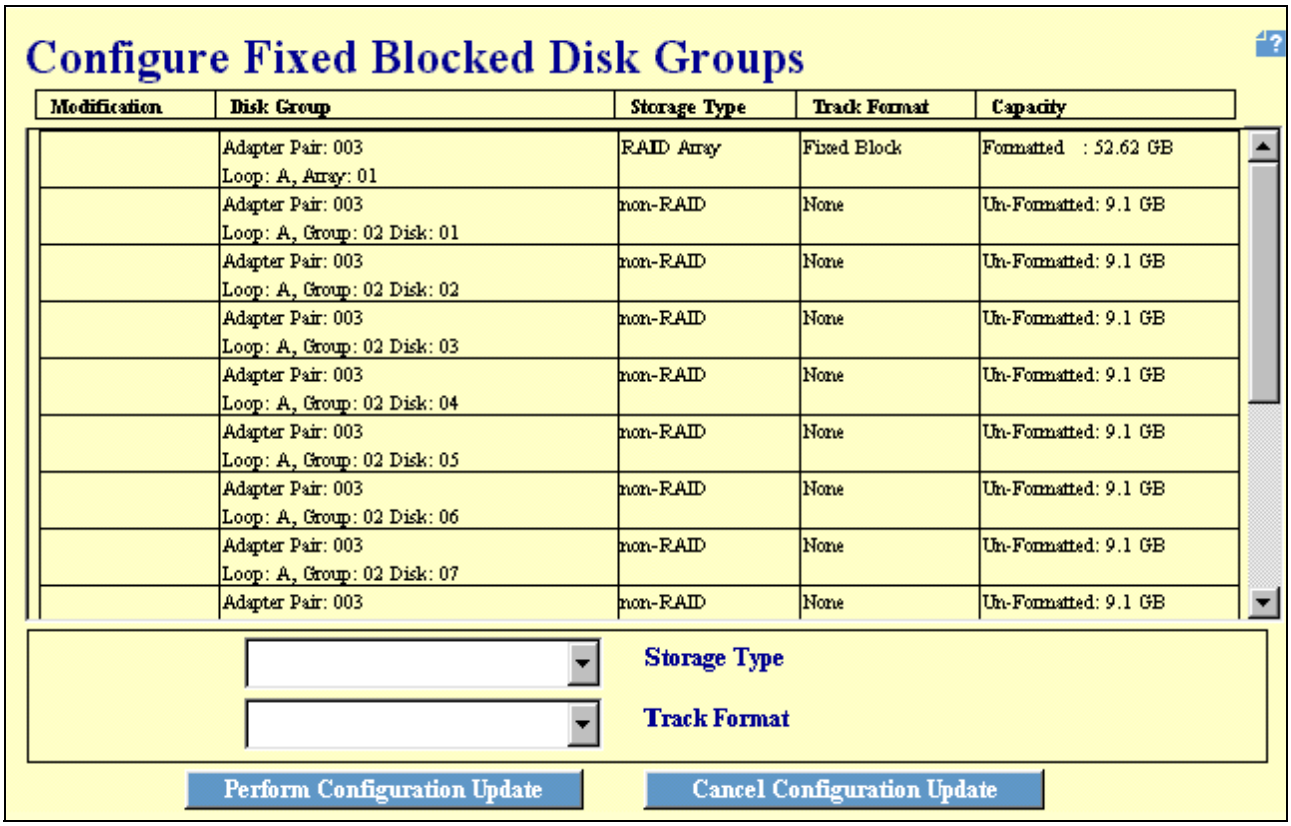

Figure 127. ESS Specialist: Configure fixed block disk groups

When the required groups are defined in the panel, each array or non-RAID disk you have created has Defined in the Modification column of the table. Click on **Perform Configuration Update**. Message 2202 is presented to indicate successful disk group configuration. See Figure 128.

Figure 128. ESS Specialist: Message 2202

Click on **OK** to return to the Storage Allocation panel.

15.2.3.1 Deleting a disk group

You can undefine an existing array, or a group of non-RAID disks by selecting **Undefined**. Warning 1802 is presented. See Figure 129 on page 217.

Figure 129. ESS Specialist: Warning 1802

Note that all logical volumes and data on the disk group will be discarded if you proceed. To cancel the reconfiguration, click on **Cancel Configuration Update** to return to the Storage Allocation panel now.

If you want to proceed with the deletion, click on **Yes**. The disk group has **Undefined** in the Modification column. Click on **Perform Configuration Update**. Message 2202 is presented as shown in Figure 128 on page 216.

Click on **OK** to return to the Storage Allocation panel.

When a single non-RAID disk is undefined, the entire disk group is deleted. If any of the disks in a non-RAID group has been defined as a S/390 device, you cannot delete the group from here without first going to the Configure CKD Disk Groups panel and changing the track format of all disks in this group to none. If all the disks in the group that are visible here have their track format changed to none, the group can be deleted from the Configure CKD Disk Groups panel.

15.2.4 Modify host systems

Open hosts must be defined to the ESS in order to access ESS logical volumes.

To define a SCSI host to the ESS, on the Storage Allocation -- Graphical View panel, select **Modify Host Systems**. The panel appears as shown in Figure 130 on page 218. Icons for previously defined hosts appear at the top of the panel. In this example an icon for a system called SUN is shown, along with the Escon Net host (which always appears). In the list of defined hosts, Escon Net and the SUN Solaris host are shown.

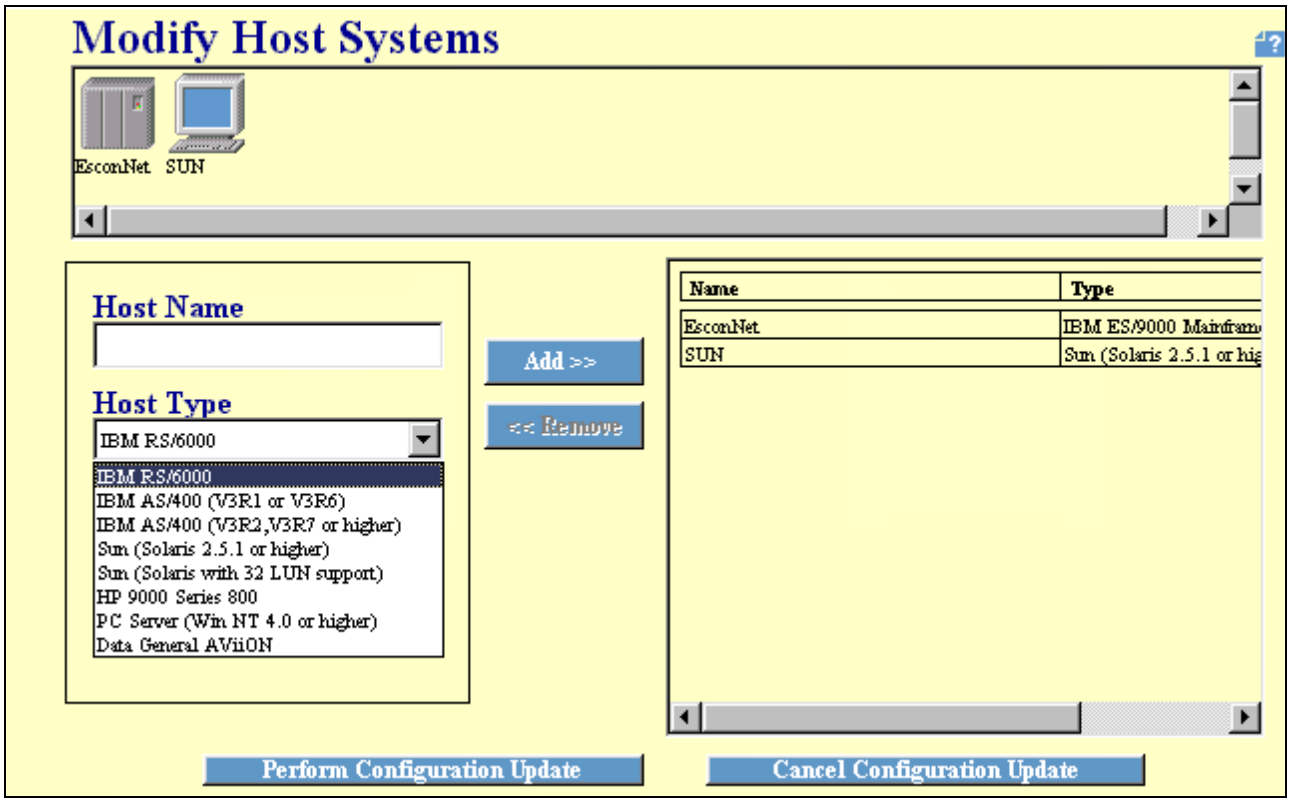

Figure 130. ESS Specialist: Modify host systems panel

When you create FB logical volumes and associate them with specific SCSI ports, the ESS Specialist uses the host type information to determine how many LUNs can be created for each target ID in the port. The port uses the defined SCSI parameters for the hosts when communicating with a host on the bus. See the list of host types in the pull-down in Figure 130.

The ESS conforms to the SCSI-3 standard, and supports up to 16 target IDs on a SCSI bus. At least one target ID must be reserved for a host system. The ESS SCSI ports support up to 15 internal target IDs. Each target ID can have up to 64 associated LUNs. Hosts using the SCSI-2 standard support 16 target IDs per bus, but only eight LUNs per target ID.

Although the ESS supports LUNs 0-63 for each target ID, the number of LUNs actually used is controlled by a Max_Lun parameter in the host definition. Max_Lun has a range of 0-63, but currently there are no open hosts that support access to more than 32 LUNs. Windows NT with Fix Pack 4 installed is an exception, in that it supports 256 LUNs per ID.

It is important to define the correct host type for each SCSI bus connection. This affects how targets and LUNs will be created by the ESS when you add volumes. Currently, AIX versions supported by the ESS support 32 LUNs per target, whereas current versions of HP-UX, Sun Solaris 2.6.0 base version and lower, and Windows NT (without Fix Pack 4) support eight LUNs. If you attach 32 volumes to an RS/6000 host at one port, one target ID with 32 LUNs are created.

For other host types that support eight LUNs per target, you can attach 32 or more volumes to a host through a single port. Multiple target IDs, each with eight

LUNs can be assigned. There is an important limitation: Each target ID assigned to a port corresponds to a different LSS. This means, for example, that you cannot add more than eight volumes from a single disk group to such a host at one port. For 32 volumes, you need to do one of the following:

1. You can distribute the volumes over four disk groups in different LSSs. The volumes can be added to a single port, and appear as groups of eight LUNs under four target IDs.
2. Or, the volumes can be in one disk group, or one LSS and be added to four different ports, all connected to the same host. One target ID/ eight LUNs are assigned in each port. The host must be cabled to all four ports. At least two SCSI adapters are needed in the host to do this.
3. Another alternative is to distribute the 32 volumes over two groups in different LSSs, 16 volumes each. For each group, add eight volumes to each of two ports attached to the host. The ports can be on a single SCSI bus.

You need to plan the storage allocation process carefully to ensure that you can actually define your required configuration, or modify it as needed.

ESS supports up to four hosts per SCSI bus. Windows NT and AIX support up to four servers on the same SCSI bus. HP-UX allows two hosts to share a bus. Sun servers and AS/400 servers allow one host on a bus.

15.2.4.1 Host name

To add a SCSI host on the Modify Host Systems panel enter the name (128 characters max) by which you will identify the host in the Host Name field. If the host is LAN connected, we suggest that you use the server's TCP/IP hostname (including domain).

15.2.4.2 Host type

Select the host type from the pull-down list. See Figure 131.

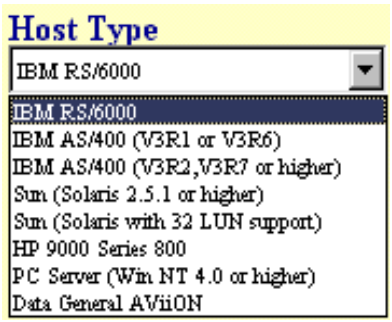

Figure 131. Host types list

The host type should be one of the following:

- IBM System/390 with MVS, VM, VSE, or TPF operating systems
- IBM RISC System/6000, including RS/6000 Series Parallel Complex
- IBM AS/400, AS/400 Advanced Series, and AS/400e Series

  For AS/400 hosts, an additional distinction is made based on the installed operating system:

- AS/400 hosts running the OS/400 operating system Version 3 Release 1, or Version 3 Release 6, only support one type of volume - the 9337-580 (4.19 GB).
- AS/400 hosts running the OS/400 operating system Version 3 Release 2, Version 3 Release 7, or a higher level, support two types of volumes - the 9337-580 (4.19 GB) and the 9337-590 (8.59 GB).

- Sun Sparc or Ultra:

  For Sun hosts, an additional distinction is made based on the installed operating system:

  - Sun hosts running the Solaris operating system at level 2.5.1 or the Solaris 2.6.0 base level support a maximum of eight logical unit numbers (LUNs) per SCSI target ID. This limits the number of volumes that the ESS can assign to those hosts.
  - Later levels of Solaris support up to a maximum of 32 LUNs per device.

    This 32 LUN support is also 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 Series 800
- All Intel-based PC servers running the Windows NT operating system level 4.0 or higher with Service Pack 3 or greater, or with Server Enterprise Edition 4.0 Service Pack 3 or greater
- Data General AViiON

The ESS may support additional SCSI host systems to those in this list. For current information on supported hosts, refer to:

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

If your host does not appear in the list, select a host type similar to yours. You can edit the SCSI parameters for the host later on the Configure SCSI Ports panel.

Click the **Add** button. The hostname appears in the List of Defined Hosts. You can add up to 64 hosts if required.

15.2.4.3 Removing a host

You can remove any of the previously defined hosts that are no longer required by selecting the host name from the list of defined hosts or clicking on the host icon at the top of the panel and clicking on the **Remove** button.

Notes:

1. The action of removing a system from this list will not affect machine operation. Therefore, you can use this process (remove and add) to modify a host name or type entry.
2. You should remove a host system from the list only after you disconnect it physically from the storage server.

To activate the changes, click on the **Perform Configuration Update** button. Message 1502 indicates a successful update of the Host System List. See Figure

132 on page 221. If you want to return to the Storage Allocation -- Graphical View panel without activating any of the changes, click on the **Cancel Configuration Update** button.

Figure 132. ESS Specialist: Message 1502

Click on **OK** to return to the Storage Allocation - Graphical View panel. Icons for the newly added hosts appear at the top of this panel.

15.2.5 Configure SCSI ports

To define the hosts that will attach to a port, on the Storage Allocation -- Graphical View panel, click on **Configure SCSI Ports**. The panel appears as shown in Figure 133 on page 222.

The Configure SCSI Ports panel identifies the attachment of one or more host systems to an ESS host port. A host system may attach to multiple ports via a separate SCSI adapter and bus cable for each port. A host system may also attach to two ESS ports using a single SCSI adapter and SCSI cable. In this case, because the ESS ports provide electrical termination of the bus, they must be at the two ends of the cable. The two ESS ports may be in the same, or different ESS subsystems.

The panel is used also to reserve target IDs for initiators and devices on the bus external to the ESS port. The remaining IDs are available for allocation to the ESS port when you add logical volumes.

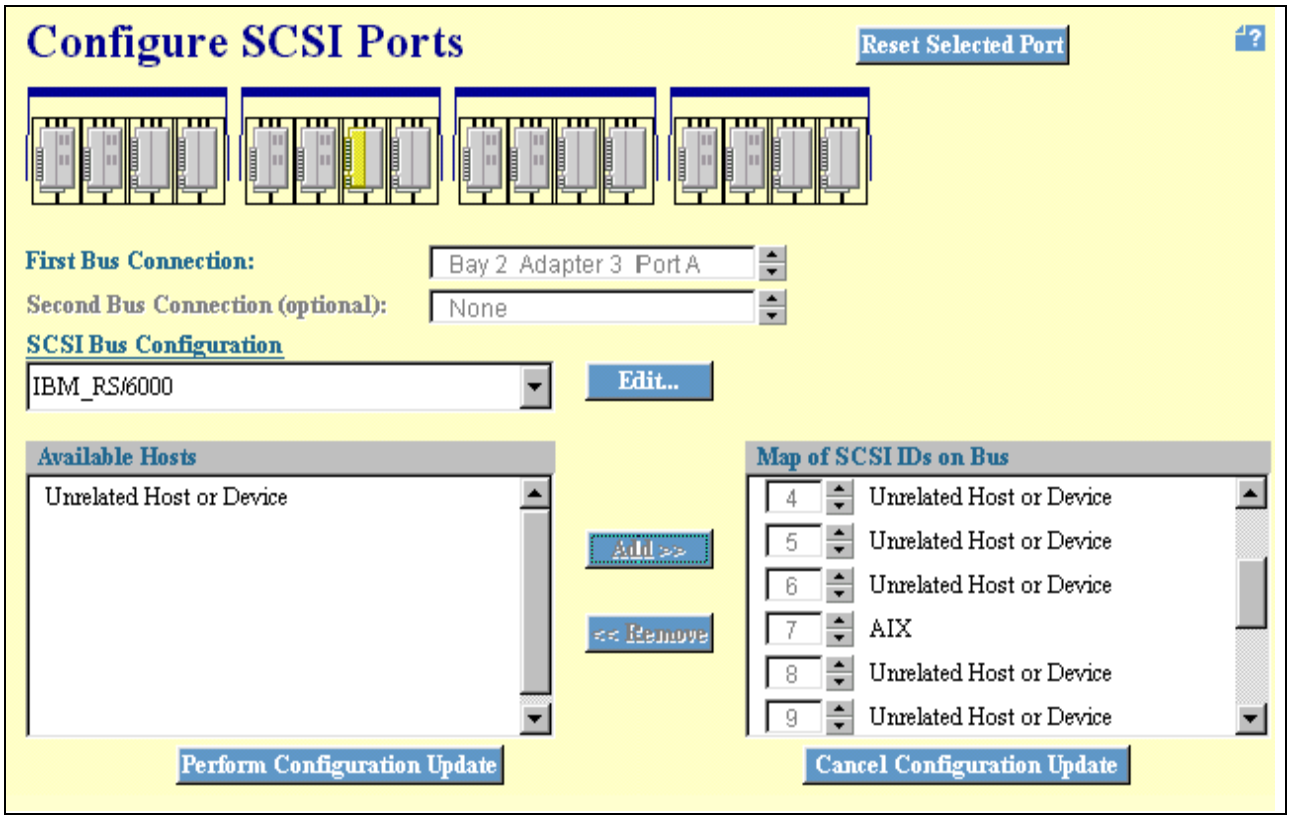

Figure 133. ESS Specialist: Configure SCSI ports panel

15.2.5.1 Selecting a SCSI port

The first row of the Configure SCSI Ports panel contains icons for the installed host adapters. The ESCON adapter icons have a different appearance to the SCSI adapter icons. In Figure 133, the SCSI adapters are the right hand pair in each HA bay. Each adapter contains two selectable ports. You may select only one port at a time for configuration.

Select a port either by clicking on the port icon in the first row or by scrolling the port selector box to show the desired port number. The selected port is highlighted in yellow and the port number is identified in the First Bus Connection field located in the second row.

If a host system is attached to two ports via a common SCSI bus cable, use the Second Bus Connection field to identify the ESS port attached to the other end of the cable. The panel will configure both ports according to a common SCSI Bus Configuration and ensure that the Map of SCSI IDs for the ports are consistent with each other.

From the SCSI Bus Configuration pull-down, select the type for the host you are attaching to the port. Options are:

- AS/400_A
- AS/400_B
- DATA_GENERAL
- HP_800

- IBM_RS/6000
- Sun_Ultra_A
- Sun_Ultra_B
- WINDOWS_NT

A list of previously defined hosts (if any) that match the selected bus configuration appears in the Available Hosts list.

15.2.5.2 Editing a SCSI bus configuration

If you want to examine or modify any of the SCSI bus parameters for your host system, click on the **Edit** button. The SCSI Bus Configuration panel is displayed. Refer to Figure 134.

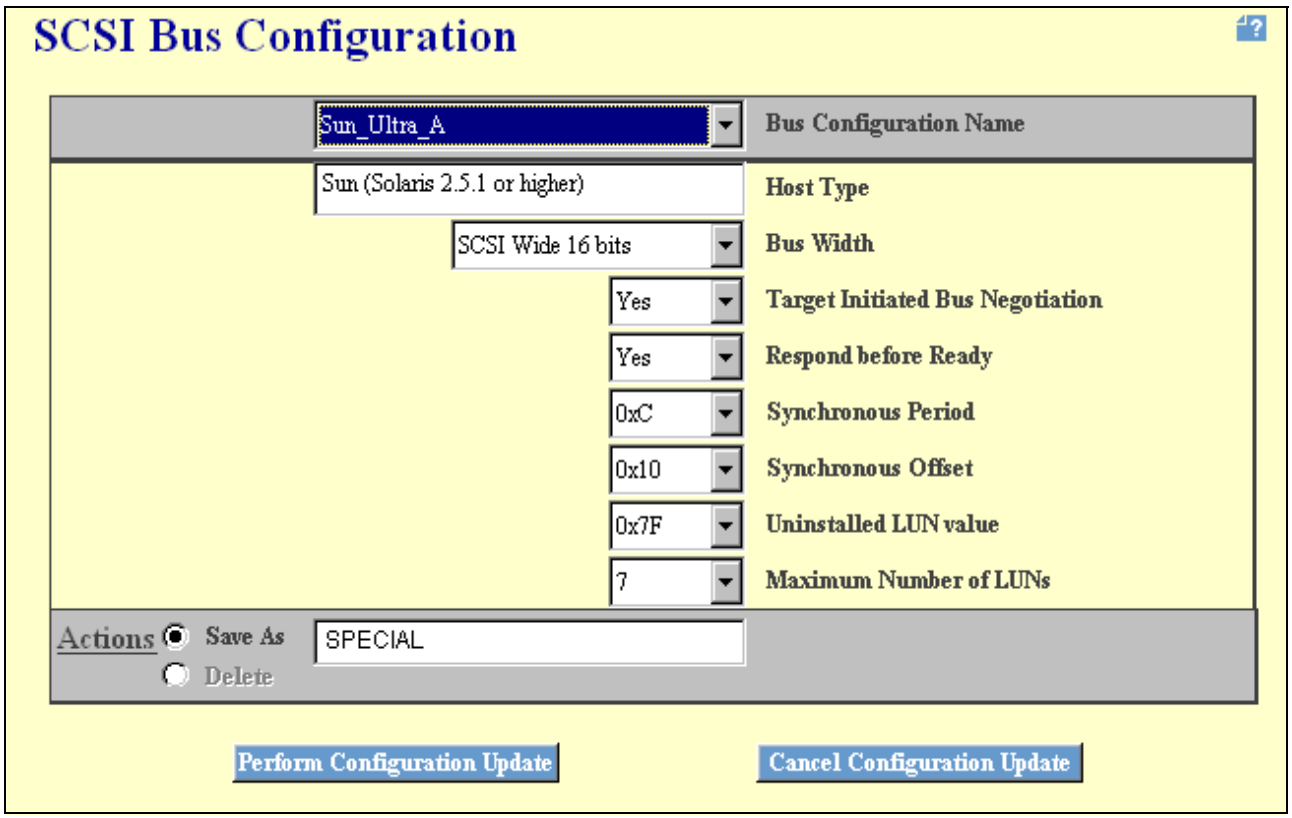

Figure 134. ESS Specialist: SCSI bus configuration panel

In this example, the Sun_Ultra_A bus configuration is displayed. You can select any available configuration from the Bus Configuration Name pull-down list on this panel. Any parameter can be changed by selecting an alternative value from the associated pull-down list. The altered configuration must be saved under a name you specify in the Actions-Save As field. You cannot alter the original bus configuration. Click on **Perform Configuration Update** to save the new bus configuration. Msg 2141 is presented to indicate successful bus configuration. See Figure 135 on page 224.

Figure 135. ESS Specialist: Message 2141

Click on **OK** to return to the Configure SCSI Ports panel. Your newly defined bus configuration should now appear in the SCSI Bus Configuration list. If you select this configuration, any host systems associated with the original bus configuration (in this example, Sun_Ultra_A), should appear under the new configuration (in this case, SPECIAL), provided that you did not alter the bus width parameter from 16 bits to 8 bits. The hosts also continue to be available under the original bus configuration.

15.2.5.3 Adding hosts

After you have selected the correct SCSI Bus Configuration as described above, The Available Hosts lists contains the previously defined hosts that conform to the selected SCSI bus configuration. Another entry, "Unrelated Host or Device" also appears. Select the desired host from the Available Hosts list.

In the Map of SCSI IDs on Bus, ensure that there are no host entries or Unrelated Host or Device entries that are not required. Any such entries restrict the Target IDs available for allocation to hosts you add, and to logical volumes you will add later. Remove unwanted host entries first, then unwanted Unrelated Host or Device entries. You cannot remove the last entry. If the last remaining entry is not required, change the target ID to 8. You can remove it after you have added a host.

Ensure that any remaining host or Unrelated Host to Device entries are set to the correct target IDs. Any entry *x Volumes Attached to SCSI ID* **x** cannot be changed to a different target ID. These entries are previously assigned ESS logical volumes. However, an entry can be deleted if not required. Any data on the volumes is discarded.

Click on **Add**. The selected host in the Available Hosts list appears in the Map of SCSI IDs on Bus list. The first added host defaults to SCSI ID 7. This is the most commonly used target ID for an initiator.

The SCSI ID is the SCSI ID of the adapter residing in the host. The adapter in the host comes with a default address which is usually set to a specific value. Typically for single host SCSI adapters, the SCSI ID is 7, but we recommend consulting the documentation for your system to determine the ID and for instructions on how to change it.

Using the spin buttons, the target ID can be changed to any value 0-15 to match the target ID of the host.

You can attach up to four hosts to any one ESS port, but they must be of the same host type. If you select a different SCSI Bus Configuration to the one currently in use for the port, warning 1610 is presented. See Figure 136 on page 225.

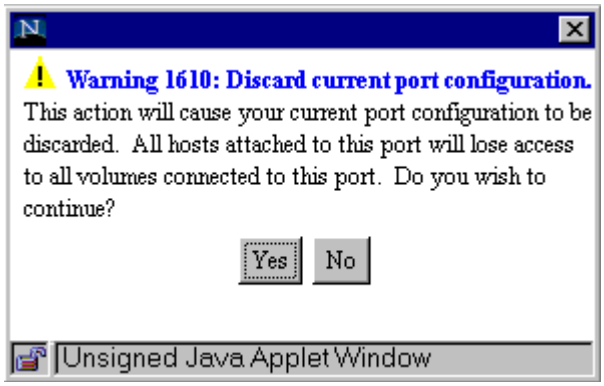

Figure 136. ESS Specialist: Warning 1610

Second, and subsequent added hosts default to target IDs 6,5,4 unless existing entries in the table have these IDs. Although the ESS supports up to four hosts per SCSI port, IBM does not recommend multiple hosts on a bus because it results in lower performance due to increased bus arbitration. Where multiple hosts are configured, we recommend that the target IDs be set to 7,6,5,4. This provide the most efficient arbitration.

The host systems must support attachment of the number of hosts you want to add to the bus.

All SCSI IDs in the Map of SCSI IDs on Bus list must be unique. This is not checked by ESS Specialist until you click on the Perform Configuration Update button. *Error 1519: Duplicate SCSI ID Assigned* is presented if any ID appears more than once in the list.

Figure 137 shows an example Map of SCSI IDs. The first host, TestRS1 has been added at target ID 7. Then two volumes were added to the port using the Add Fixed Block Volumes panel. They had target IDs 6 and 5 assigned automatically by ESS Specialist. The volumes are attached at different IDs because they are not in the same LSS.

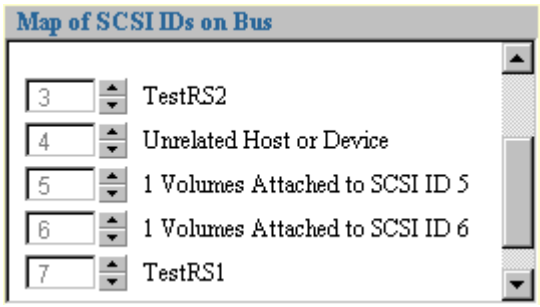

Figure 137. ESS Specialist: SCSI ID map

Subsequently, an unrelated host or device has been added at ID 4, and host TestRS2 added at target ID 3. The two volumes are available to both the TestRS1 and TestRS2 hosts.

You can control the target IDs at which logical volumes are later added on the Add Fixed Block Volumes panel. When all the necessary hosts, and Unrelated Host or Device entries have been added to the port, you can reserve any of the

remaining unassigned IDs by adding an Unrelated Host or Device entry set to that ID. This can be removed later if you want to assign the ID to another host, or to additional logical volumes.

Note: When adding logical volumes to a port, ESS Specialist assigns new target IDs to the port as required, in the order 7 to 0, then 15 to 8, skipping any IDs already assigned to:

- Hosts
- Unrelated hosts or devices
- Existing logical volumes

If two ESS ports are on the bus, each port must have Unrelated Host or Device entries for any target IDs assigned to logical volumes in the other port. This prevents possible duplicate IDs on the bus if further volumes are added to either port from a new LSS. This is managed automatically by ESS Specialist if the Second Bus Connection field is used to specify the second ESS port on the bus.

The number of devices you can add to the port or, more precisely, the number of different Logical Subsystems to which added volumes can belong is limited to the number of unreserved target IDs available for the port.

When the required hosts have been added, and the Map of SCSI IDs on Bus list is correct, click the **Perform Configuration Update** button. Message 2104 indicates successful port configuration. See Figure 138.

Figure 138. ESS Specialist: Message 2104

Once you have defined the SCSI host, you can connect it to the ESS. The physical link is made by connecting a suitable cable between the ESS port and the host SCSI port. If other SCSI hosts or devices are cabled onto the bus, the ESS port must be at one end of the cable. The physical connection should be made after defining the connection through the ESS Specialist Configure SCSI Ports panel.

Each SCSI adapter has two ports: A and B. In the host adapter icons on the panel, port A is on the left, port B on the right. On the adapter, port A is the upper SCSI connector. SCSI adapters in the ESS can be located anywhere in the four host adapter bays.

15.2.6 Add fixed block volumes

In order for hosts to access data located in the ESS, the ESS storage must be partitioned into *logical volumes.* Each FB logical volume is seen by the host as a SCSI target/LUN, even though the logical volume is physically allocated in SSA disk storage in the ESS. A logical volume cannot span ranks, but many logical volumes in different ranks can be accessed through a single host port.

Although logical volumes must initially be defined to a particular host and port, a volume can subsequently be added to one or more additional ports. The volume is accessible by any host configured to any port to which the volume has been added.

Logical volumes are selectable from a range of sizes, from 0.5 GB up to 224 GB. For the AS/400, only 4.19 GB, 8.59 GB, 17.55 and 36 GB are available.

To define logical volumes, and associate them with a SCSI port, on the Storage Allocation -- Graphical View panel, select **Add Fixed Block Volumes**. The Add Fixed Block Volumes (1 of 2) panel is displayed. Refer to Figure 139.

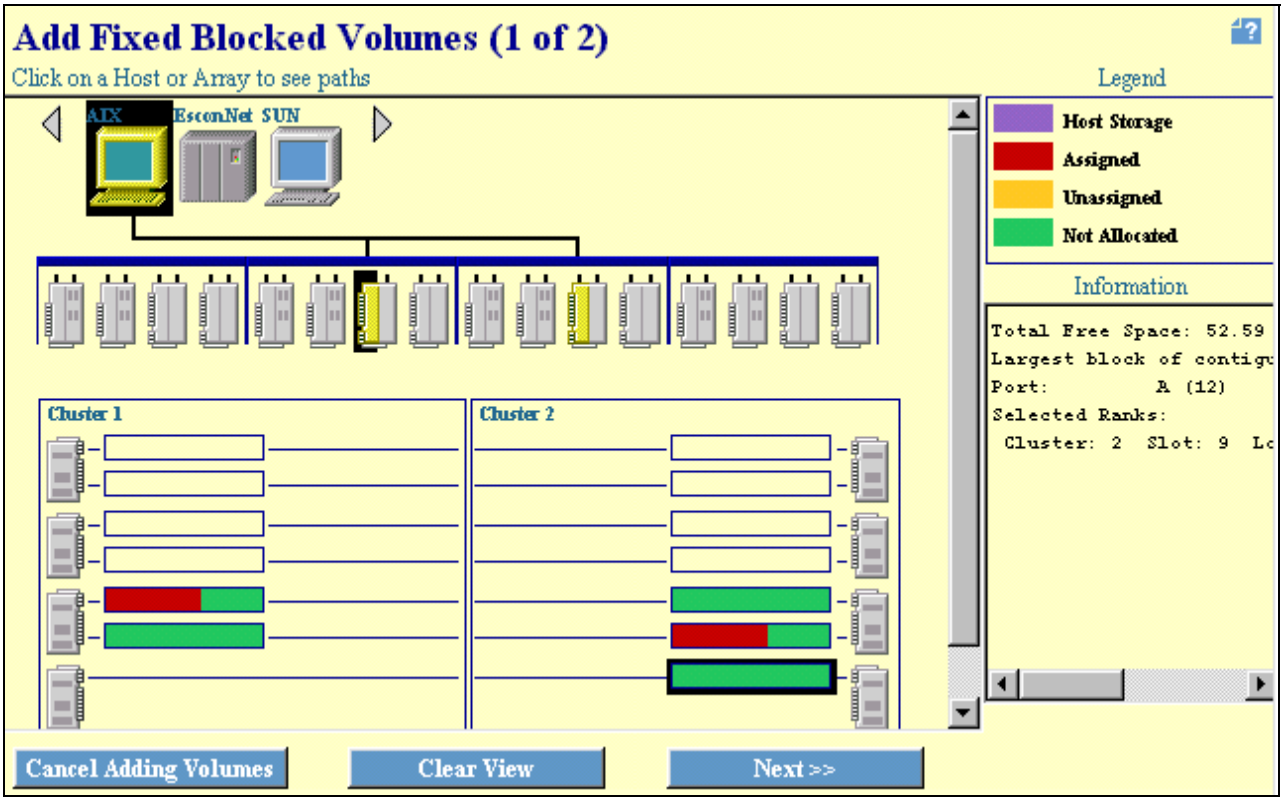

Figure 139. ESS Specialist: Add Fixed Block Volumes (1 of 2) Panel

15.2.6.1 Select host

On this panel, the first row contains icons for the defined host systems. The second row contains icons for the installed host adapters. When you select a host by clicking its icon, the background of the host icon changes to black to indicate selection, and the SCSI ports attached to that system are highlighted in yellow.

Any disk groups that currently have volumes added to any of the highlighted ports are now shown in color. The color code indicates current allocation of the space in the groups. See the legend in the top right corner of the panel.

15.2.6.2 Select port

Select the desired port by clicking on the left side of the host adapter icon for port A, or the right side for port B. The background of the selected port changes to black to indicate selection.

15.2.6.3 Select disk groups

All disk groups available for allocation of fixed block volumes are now shown in color. Any groups shown in outline only are assigned as S/390 storage. Select one or more groups to have volumes allocated by clicking on each group. The background of the groups changes to black to indicate selection. You can allocate volumes to any combination of FB disk groups by selecting them here. You may want to include only RAID groups, or only non-RAID groups in any one selection.

If you select a host or a port and then want to change your selection before clicking the **Next >>** button, you can click the **Clear View** button and start again. You can deselect a disk group simply by clicking on it.

Note: You must select a host, one port, and at least one disk group before proceeding to the Add Fixed Block Volumes (2 of 2) panel.

15.2.6.4 Non-RAID groups

If you select a non-RAID group, the panel shown in Figure 140 is displayed.

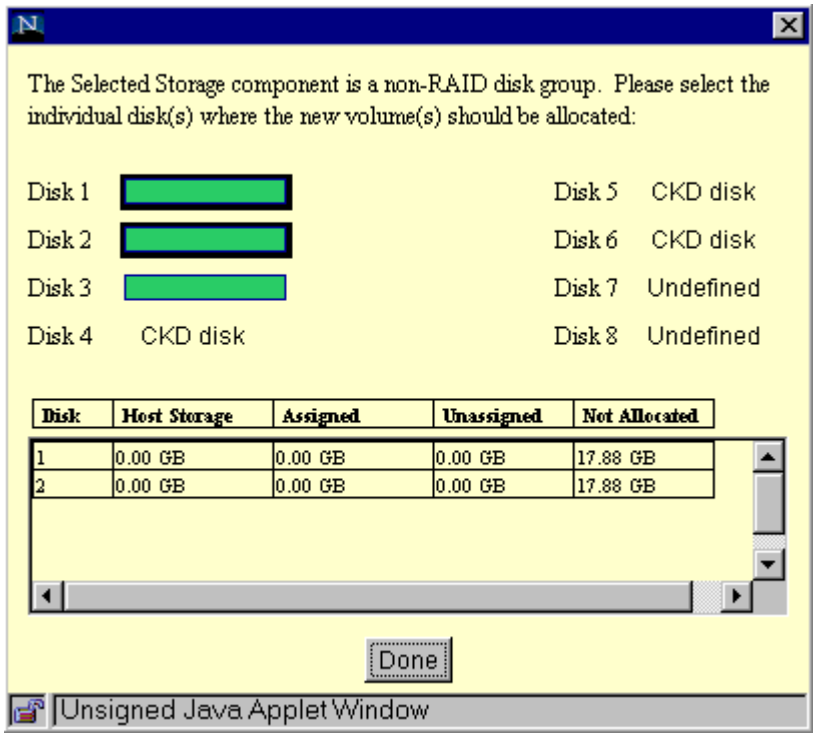

Figure 140. ESS Specialist: Non-RAID disk group

In this example, the non-RAID disk group was created using the S/390 panels, and three of the disks defined as CKD storage. On the Configure Fixed Block Disk Groups panel, three of the disks were defined as FB storage. Two of the disks in the group remain undefined.

Of the three FB disks colored green to indicate that they are available for allocation of FB volumes, we have selected disks 1 and 2. Their background is changed to black to indicate selection.

Once the desired disks in the array are selected, click **Done**.

Disk groups on different DA pairs are in different logical subsystems. Two, or three groups on the same loop may be in one LSS, or two different LSS. Where no expansion frame is installed, there are usually only two disk groups per loop, and they belong to LSSs in alternate clusters. On the Storage Allocation -- Graphical View panel you can click on each group and view its status in the Information box. On an SSA loop, even numbered groups belong to one LSS for the loop; odd numbered groups belong to the alternate LSS.

This is of interest because each LSS that has one or more disk groups represented in your selection needs a new target ID in the associated SCSI port if you add volumes from those groups, and that LSS has no volumes already allocated to the port. If there are no available IDs for the port, you may still be able to add new volumes, provided you do not select disk groups in any LSS that does not already have logical volumes allocated to the port.

When your selection of host, port and disk groups is complete, click on **Next >>** to proceed to the Add Fixed Block Volumes (2 of 2) panel. See Figure 141 on page 229.

Click on **Cancel Adding Volumes** to return to the Storage Allocation panel.

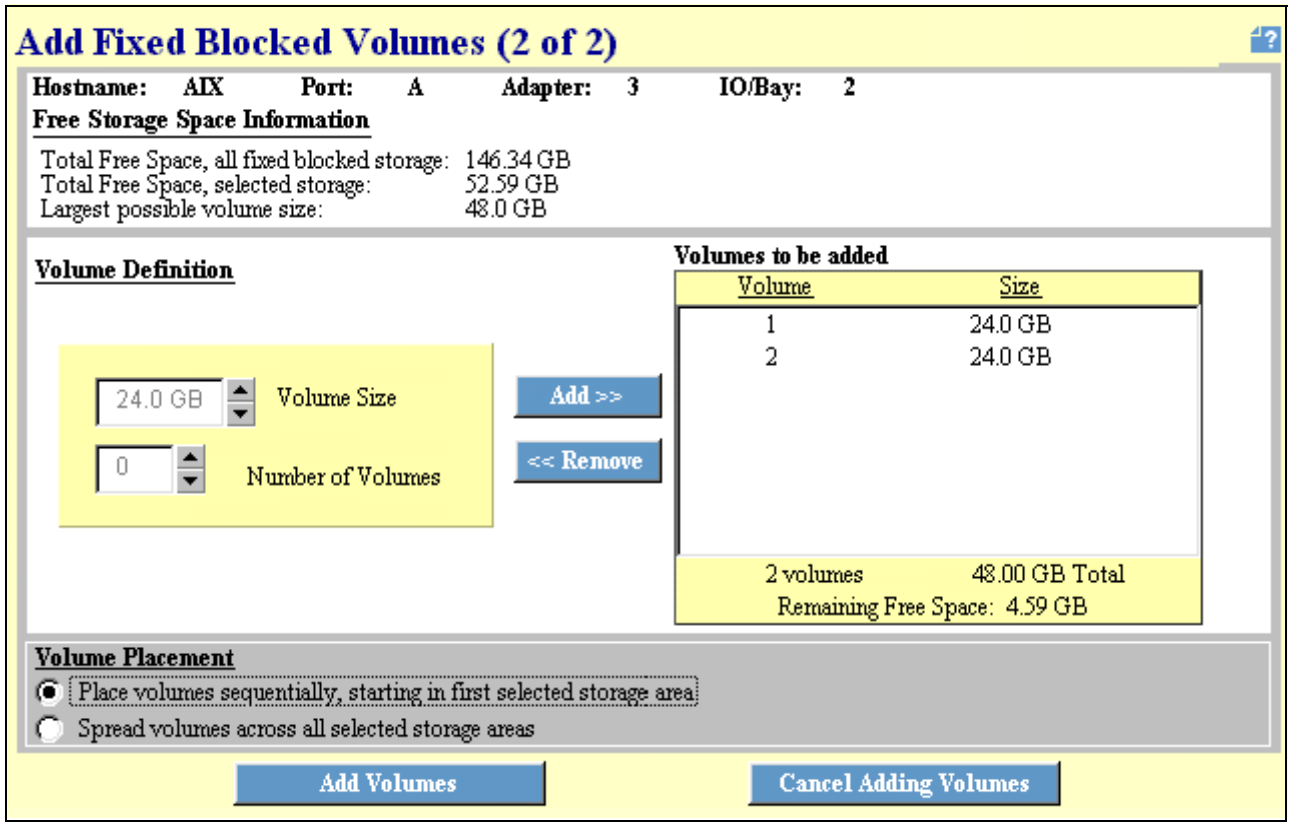

Figure 141. ESS Specialist: Add Fixed Block Volumes (2 of 2) panel

This panel is used to allocate volumes in the selected disk groups, and add them to the selected SCSI port.

The selected host, and port are shown on the top line of the panel. Under Free Storage Space Information are three items:

Total free space, all FB storage	The total unallocated space in fixed block ranks in all disk groups. Additional groups, or non-RAID ranks (disks) may be defined as S/390 space, or undefined.
Total free space, selected storage	The total unallocated space in fixed block ranks in the selected disk groups. Selected non-RAID disk groups may have individual ranks (disks) that are either defined as S/390 or undefined.
Largest possible volume size	The largest contiguous unallocated space in any FB rank in the selected disk groups.

15.2.6.5 Add volumes

To add volumes, follow these steps:

1. In the Volume Size field, use the spin buttons to select the required volume capacity you want, in gigabytes (GB).
2. In the Number of Volumes field, enter the number of volumes of this size you want to create, or use the spin buttons.
3. From Volume Placement, select to either:
   - Place volumes sequentially, starting in first selected storage area.
   - Spread Volumes across all selected storage areas.
4. Click on the **Add >>** button.

Notes:

1. The selector fields do not allow you to select a volume size that exceeds the *Largest possible volume size*, and they do not allow you to select a combined volume size and quantity that cannot be allocated in the *Total Free Space, selected storage*.
2. If there is more than one rank in your selection of disk groups, you can control how the volumes are distributed in the selected storage, using the Volume Placement options. You can fill the selected storage areas sequentially with allocated volumes, in the order the ranks (disk arrays or non-RAID disks) were selected, and in the order the volumes are allocated. Alternatively, you can evenly distribute allocated volumes one per rank in rotation around the selected ranks.

   If both RAID disk groups and non-RAID disk groups are selected, the volumes are allocated among both RAID and non-RAID ranks as described. This is probably not desirable. To avoid this, include only RAID groups, or only non-RAID groups in any one selection of storage.
3. If the selected host is an AS/400, the logical volumes created are AS/400 volumes. These use a 520 byte-per-block logical track format. For other open hosts, the added volumes have the conventional 512 byte per block logical track format. For clarity, we repeat here that the physical SSA disks in the ESS are formatted using a 524 byte-per-block format. This accommodates both AS/400 and Open Systems logical volumes on the same physical disk or array.

Because the AS/400 volumes and open systems volumes are incompatible, you cannot add both types of volume to a single host port on the Modify Fixed Block Volume Assignments panel.

You may add additional volume size and quantity selections until the required configuration for the selected port has been defined. To remove a volume from the list, click on the volume to select it. It highlights. Click the **<< Remove** button. The Remaining Free Space area in the Volumes to be added table is recalculated each time an entry in the list is added or removed.

Clicking on **<< Back** at any time returns you to the Add Fixed Block Volumes (1 of 2) panel.

When the Volumes to be added list is correct, click on **Perform Configuration Update**. Message 2101 indicates successful volume configuration. See Figure 142.

Clicking on **Cancel Configuration Update** returns you to the Storage Allocation - Graphical View panel.

Figure 142. ESS Specialist: Message 2101

Click on **OK** to return to the Add Fixed Block Volumes (1 of 2) panel. On this panel, when you click on the host for which you just added volumes, the color coding in the disk groups shows the modified space allocation. If you have no further volumes to add, click on **Cancel Adding Volumes** to return to the Storage Allocation - Graphical View panel.

Note: The Add Fixed Block Volumes panel allows you only to create new volumes; it does not allow you to remove volumes created in a previous operation. Entries in the Volumes to be added list, but not yet created can be removed. Once the volumes are physically allocated, they cannot be deleted. You can, however, unassign volumes through the Modify Fixed Block Volume Assignments panel. This removes all association between the volumes and any port, but the volumes, and the data on them continue to exist. Undefining a disk group on the Configure Fixed Block Disk Groups panel deletes all volumes residing in the group. Any data on the volumes is discarded.

15.2.7 Modify fixed block volume assignments

The Modify Fixed Block Volume Assignments panel is used to associate logical volumes with additional ports, and to remove volumes from ports as needed. The Modify Fixed Block Assignments panel is shown in Figure 143 on page 232.

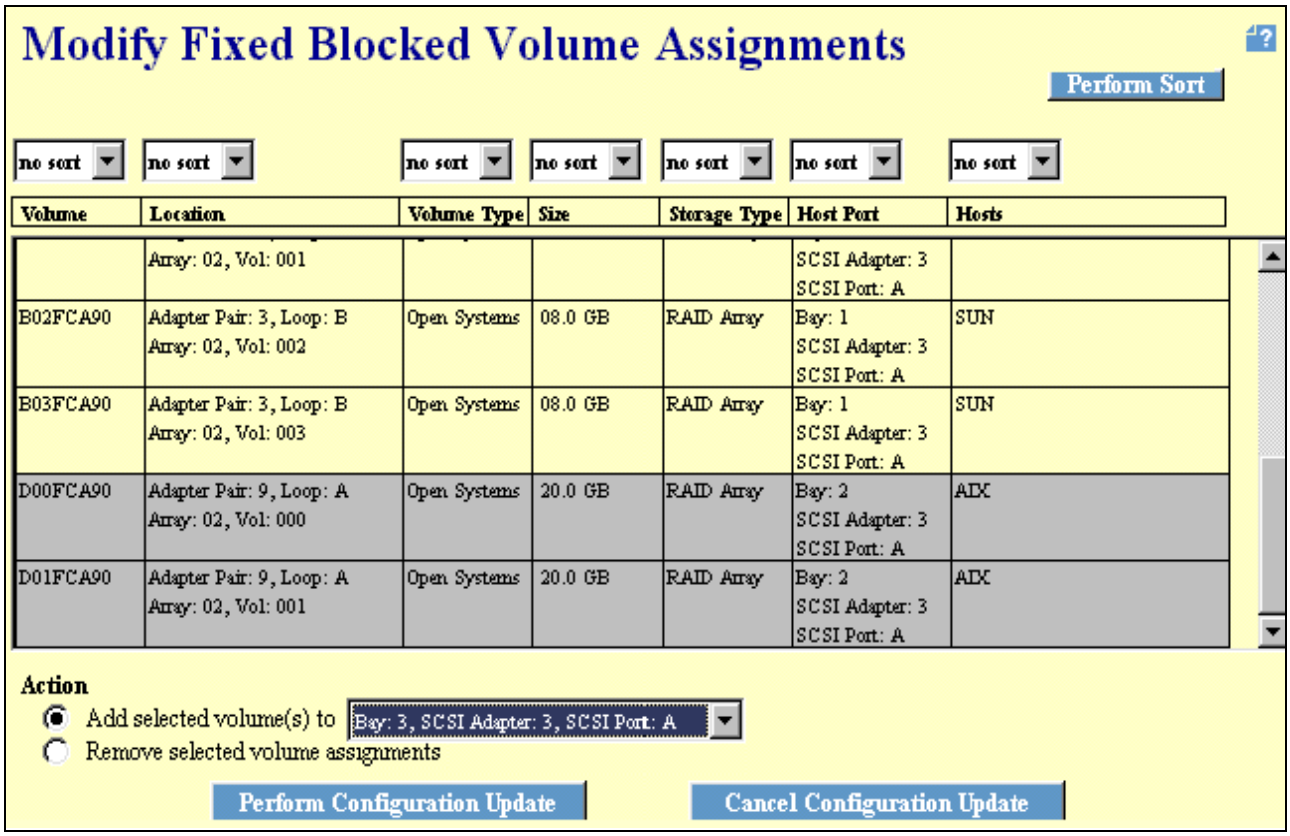

Figure 143. ESS Specialist: Modify fixed volume assignments panel

The Modify Volume Assignments panel is a table of seven columns. A minimum of one row exists for each defined volume. If a volume is assigned to a port then the row identifies the assigned host port and host system. If a volume is assigned to two or more ports, the table contains an additional row for each extra port assignment.

15.2.7.1 Volume table

The Volume table columns are:

Volume Displays the volume serial number if the volume is assigned to a host port; otherwise the volume is unassigned. The volume serial is an ESS-generated eight digit hexadecimal number created by concatenating the logical volume ID and the ESS machine serial number. The logical volume ID is a three digit number assigned in sequence to logical volumes as they are created. The volume serial number is unique among ESS subsystems.

Location Displays the location of the logical volume expressed as DA pair/cluster/loop/group/rank/volume number. Rank is either array or disk, as applicable.

Volume Type Displays the volume type as either AS/400 or Open Systems.

Size Volume size in gigabytes.

Storage Type RAID array or non-RAID.

Host Port Displays the assigned SCSI port expressed as Bay/Adapter/Port. If a volume is assigned to multiple ports, there is a row for each unique volume and port number combination.

Hosts Displays a list of host system names that are associated with the volume and port number combination.

Use the scroll bar to move the table past the viewing window as required.

Note: Do not add a volume or modify a volume assignment on an active port. This may affect operation of some host systems.

15.2.7.2 Sort options

Each column in the table has a pull-down combo-box, containing the entries no sort, first, second, and third. Select these options in sequence for the columns of your choice to establish a sort hierarchy. When a sort choice is selected for one column, it is no longer shown for the other columns. Click on the **Perform Sort** button to rearrange the table according to the sort options you selected.

For example, if you select first for the hosts column, and second for the size column, the table shows the entries sorted by host. The entries for each host are further sorted by volume size.

15.2.7.3 Volume selection

Click on a single row in the table to select it. It highlights in grey. To select additional rows, use CTRL- left click on each row. For a range of rows, select a row at one end of the range, then SHIFT- left click on the row at the other end. To deselect highlighted rows, CTRL- left click on them. To deselect all but one highlighted row, click on it.

15.2.7.4 Select action

Perform one of the following actions for a selected volume(s) by clicking one of the action buttons at the bottom of the panel:

Add Volumes(s) to Adds selected volumes to the port you select from the pull-down in the associated port selector box

Remove selected volume assignments Removes selected volume(s) from a port

Click on the **Perform Configuration Update** button. Message 2106 indicates successful assignment. See Figure 144.

Figure 144. ESS Specialist: Message 2106

Click **OK**. You are returned to the same panel. To return to the Storage Allocation panel, click on **Cancel Configuration Update**, or on **Storage Allocation.**

Notes:

1. To add a volume, you must select a port in the selector box.

   You cannot add a volume to a port if:

   - The volume type (AS/400 or open systems) is not compatible with hosts assigned to that port.
   - Addition of the volume requires that a new SCSI ID be allocated in the port and one is not available (all 16 IDs on the bus are reserved or allocated). Addition of a volume does not require a new SCSI ID to be allocated if the volume is in the same logical subsystem as any existing device on the port.
   - The volume is in an LSS whose target ID in the port has the maximum number of LUNs assigned for the port's host type.

   In each case, the port is not available in the selector box.

2. Using the Add and Remove functions you can:

   - Assign a volume to a second or subsequent host port. This allows hosts on different buses to address the volume, or a single host to address the volume over two or more paths.
   - Assign a currently unassigned volume to a host port.
   - Remove a volume from a port.
   - Move a volume from one port to another.
   - Leave the volume unassigned by removing its last port assignment.

3. Leaving a volume unassigned does not affect the volume definition. The data remains on the volume and will be accessible when the volume is next assigned to a port.

4. A volume, once created, cannot be individually deleted. To remove logical volumes from a rank (array or non-RAID disk), you have to undefine the entire disk group. Once you undefine the group, all logical volumes and data on that group are lost. Refer to 15.2.3.1, "Deleting a disk group" on page 216. You can then redefine the group and add new volumes.

5. If a volume is accessed by more than one open system host, the hosts should be architecturally compatible. The format used to store data on disk (that is, the file system) should be the same. Data corruption may occur if hosts are not compatible. Compatibility of hosts is partially enforced by allowing only one Host Type to be associated with a SCSI port. However, this does not prevent a volume from being accessed by incompatible host systems through different ports.

15.2.7.5 Example volume assignments

The rather busy diagram in Figure 145 on page 235 illustrates some examples of volume addition to SCSI ports in an ESS.

On the left of the diagram are two disk groups; one a 6+P RAID array, the other a non-RAID group of eight drives. They belong to different logical subsystems, determined by their location, as specified in the figure.

On the right of the diagram are two SCSI ports; bay 1, adapter 3, port A and bay 2, adapter 3, port A. The boxes in the figure are the map of SCSI IDs on bus for the ports as shown in the Configure SCSI Port panel of ESS Specialist.

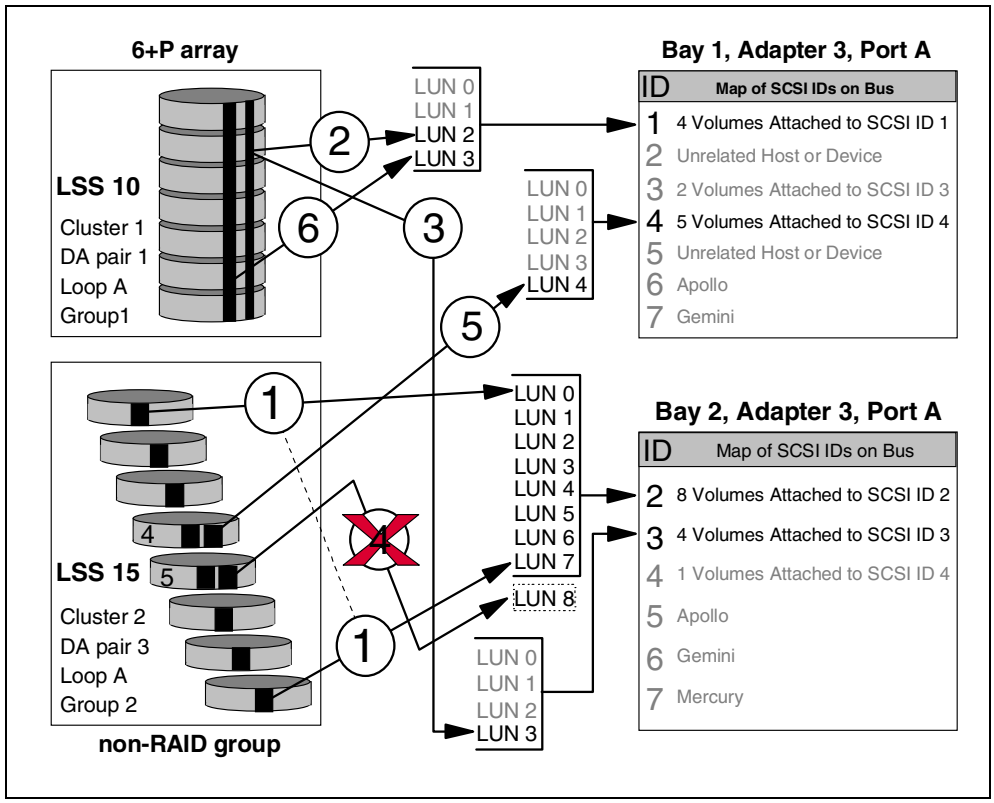

Figure 145. Volume assignments to a SCSI port

To the left of the SCSI ID maps are the LUNs assigned to the indicated SCSI IDs for the port.

The IDs and LUNs shown in grey were already assigned before the volume additions described in this example. The IDs and LUNs shown in black represent the state of the ports after the volume additions.

Summary of rules:

- Each FB LSS can appear at any SCSI port at only one SCSI ID.
- A volume can appear at any SCSI port at only one ID/LUN combination.
- The number of volumes in an LSS that can be added to any one port is limited to the maximum number of LUNs per ID in the SCSI bus configuration for the port.
- The maximum number of IDs available for volume addition at a SCSI port is 15 (when only one host is defined, and no IDs are reserved).
- FB ranks configured as AS/400 storage can be added only to SCSI ports with the AS/400 SCSI bus configuration. FB ranks configured as Open storage (host other than AS/400) can be added only to SCSI ports with other than one of the AS/400 bus configurations.

ESS Specialist assigns SCSI IDs within a port in the order 7 through 0, then 15 through 8, skipping IDs already reserved or assigned to hosts. It assigns LUNs to an ID in increasing order, beginning with 0.

Example additions made as described here:

1. A group of eight volumes, one per disk, is assigned from the non-RAID group in LSS 15 to bay 2, adapter 3, port A. Because these are the first additions from LSS 15 to this port, a new ID is assigned in the port. The volumes appear at ID 2, LUNs 0 through 7.

   ID 3 is already in use because of previous volume additions from LSS 10 (shown in grey as LUNs 0,1,2). ID 4 is assigned to a volume from another LSS. IDs 5,6,7 are assigned to three hosts, Apollo, Gemini, and Mercury.

2. A volume striped across the RAID array in LSS 10 is added to bay 1, adapter 3, port A. Previous additions have been made to this port from LSS 10, at ID 1, LUNs 0,1. The new volume is added at ID 1, LUN 2.

3. The same volume is added to bay 2, adapter 3, port A. Previous additions to this port from LSS 10 appear at ID 3, LUNs 0,1,2. This addition is at LUN 3. The volume is now visible to the Apollo and Gemini hosts through two ports: bay 1 adapter 3, port A (ID 1, LUN 2), and bay 2, adapter 3, port A (ID 3, LUN 3).

4. An attempt is made to add a second volume from disk 5 in the non-RAID group in LSS 15, to bay 2, adapter 3, port A. This is unsuccessful because the SCSI bus configuration chosen for the two ports in this example allow eight LUNs per SCSI ID. All eight LUNs for ID 2 in this port are now assigned. No further volume additions can be made from LSS 15 to this port. However, this device can still be added to other SCSI ports, including bay 1, adapter 3, port A. Note that there may be one or two groups in LSS 15 in addition to the non-RAID group shown, if the expansion frame is installed.

5. A second volume is added from disk 4 in the non-RAID group in LSS 15, to bay 1, adapter 3, port A. Four previous additions from LSS 15 appear at ID 4, LUNs 0,1,2,3. This volume is added to ID 4 at LUN 4.

6. A second volume from the RAID array in LSS 10 (the fourth in total from LSS 10) is added to bay 1, adapter 3, port A, at ID 1, LUN 3.

Chapter 16. Data Migration in the Open environment

This chapter describes migrating data to the ESS in the OS/400, UNIX and Windows NT operating system environments.

16.1 General considerations

This section describes the general data migration considerations for the ESS.

16.1.1 Method selection

Most methods of data migration affect everyday operation of a computer system. When data is moved, it must be in known state, typically requiring updates or changes to cease while the movement occurs. Depending on the amount of data you are moving and your methods, data could be unavailable for an extended period of time, perhaps several hours.

Factors like creation of new logical volumes or file systems, modification of configuration files, and data integrity checks, contribute to the unavailability of data that you are migrating. This section describes some of the aspects which you must consider when selecting a method to use for migrating data.

Select the method that best meets your criteria for:

- The amount of data to be migrated
- The amount of time available
- The availability of tape devices or spare disks for temporary storage
- The format of the data itself

You should select the method that is the best compromise between effiency and the least impact on the users of the system

16.1.2 Replacing existing storage

When the ESS replaces existing storage, you must partition it so that its logical disks are similar in configuration to the drives that it is replacing. New configurations should be large enough to accommodate the existing data.

The ESS supports existing 7133-020 or 7133-D40 drawers which can be part of the versatile storage server (VSS) or can be installed in expansion rack of the ESS. If VSS is attached to an ESS, the VSS control unit function is not active, and only rack and disks in the drawers are used.

Whenever those units are attached to an ESS, the drives must be reformatted for ESS usage. The ESS uses a different format and internal drive information. Therefore, you need to backup all data on drives that you are migrating before commencing movement of the drawers.

16.2 Data migration for AS/400 systems

The ESS emulates a 9337 subsystem when you attach it to an AS/400. As an AS/400 requires 9337 subsystems to have groups of four to eight drives, you must also configure the ESS to appear as four to eight drives.

The AS/400 expects to see a separate device address for each drive in the subsystem. The ESS meets this requirement by reporting unique addresses for each virtual drive defined to the AS/400.

Since the AS/400 has not supported 7133 drawers, you will have no existing 7133 drawers to reformat. If you have an existing disk subsystem that you are replacing with a ESS, you need to migrate the data to the ESS. Use your existing host utilities for the data migration.

You can select from several methods to migrate data to the ESS:

- Logical ADD and REMOVE functions
- Save and restore methods with tape devices

  You may also use these methods if you remove an existing disk subsystem before you install the ESS.

Note: Since the ESS always reports to the AS/400 as a Redundant Array of Independent Disks (RAID) protected 9337, any reference to starting or stopping device parity does not apply.

For more information on these procedures, see the AS/400 Backup and Recovery Guides for the release of the operating system you have on your host system. IBM also offers data migration services.

16.3 Data Migration for UNIX systems

For UNIX systems, you can use a variety of methods for copying or moving data from one disk drive to another. The host sees the ESS as one or more generic DDMs.

16.3.1 Migration methods

16.3.1.1 Volume management software

Volume management software provides specific tools for wholesale movement of data. Management software provides simple robust methods that you can generally use during production without disturbing users.

AIX, Solaris, and HP-UX all have volume management software that directly controls the disks and storage devices attached to the system. It provides the system administrator with tools for configuring and managing disk drives, file systems, paging, and exchanging spaces. The software also provides the operating system interface to the data.

16.3.1.2 Direct copy

If the data you are migrating resides as individual files on UNIX file systems, and no volume management software is available, use a utility. A utility which supports a direct copy feature, such as *cpio* with the *-p* (pass) option provides the next easiest method of moving the data. The *cpio* feature is available on all of the UNIX operating systems which support the ESS.

16.3.1.3 Backup and restore

In some cases, the only method available to transfer data is to back it up to a tape device and restore it to the new disk. This method is obviously slower, because

tape devices are essentially slow devices. However, if disk are being removed before you install the ESS, the only way to move the data is with a tape device.

Backup and restore procedures will generally have the most impact on the system usage. They require that databases and file systems be in quiescent states to ensure a valid snapshot of the data.

16.3.1.4 Dump and restore

These commands are similar in function to the Backup and Restore command; you find them on almost all forms of UNIX. They too require an intermediate device.

16.3.1.5 Other commands

You can find other commands on UNIX systems for archiving the data. Again, these commands require that you create an intermediate archive, usually on a tape drive or spare disk drive. For more see Chapter 18, “ESS Specialist Copy Services for Open systems” on page 249.

Note: You may not be able to use the volume management methods of migrating data for the following:

- For databases that use raw file systems
- For logical volumes or methods other than a UNIX file system

  This is most probable if the database uses volume serial numbers in its licensing code, or validity checking.

  If databases use licensing methods or validity checking, you may only be able to:

  - Export the database from its old locations
  - Import the database to its new location

  It is up to the database software to provide some mechanism to move the data. This may take the form of a standard database backup and restore if it does not have any specific tools for movement of data.

16.3.1.6 The AIX logical volume manager (LVM)

The AIX LVM provides methods that you can use at any time without disrupting access to the data by users. You may notice a small performance degradation, but this is better than having to shut down database or require users to log off the system.

The AIX LVM provides useful tools and utilities for migration of data as part of the AIX base operating system release. You can use these tools to move data to and from the ESS as you would on any other disk drive connected to an AIX system.

These methods for data migration work below the file system level. They do not care what sort of data resides on the logical volume, whether it is a UNIX file system or a raw database.

Instruction for data migration using AIX LVM mirroring:

- Determine which logical volumes need to be migrated.

  Most of you have existing disk drives which will be replaced by new ESS storage. The old storage might be reused at a later time for new data, after the current data is migrated to ESS storage. For example, you may want all of the

data residing on the SSA disks to be moved to the ESS. After this migration is complete, the old SSA disks might be used on a different system or installed in the ESS for future use. In this case, in order to move the data while keeping it available to users, you will have to mirror the logical volumes in which the data is stored. The following steps describe how to identify those logical volumes.

- Make a list of the physical disks (hdisks) that need to be migrated.

  An hdisk may represent a RAID 5 array and, therefore, be associated with several pdisks (physical volumes). In this case you will still have to mirror all the logical volumes contained in the array, even if some of the pdisks in the array will not be replaced by ESS storage.

- Make a list of the logical volumes that need to be migrated.

  For each hdisk, run the command `lspv -l hdiskx` where `x` is the hdisk number. This output will give you the name of all the logical volumes that reside on the hdisk. If there is a filesystem mounted on the logical volume, this output will also give you the mount point for that filesystem.

  Based on the outputs of your `lspv -l hdiskx` commands, create a list of all the logical volumes that will need to be mirrored.

  Note: There may be logical volumes that span more than one hdisk. In this case, the logical volume will be indicated in the output of more than one `lspv` command that you ran. You should take care not to confuse the number of PPs that a Logical Volume uses on a single disk with the number of PPs that a logical volume uses on the entire system. In other words, do not rely on the output of this command to indicate whether or not the logical volume is mirrored.

- Mirror the Logical Volumes.

  Mirroring does degrade system performance slightly. However, most customers will prefer performance degradation to downtime. On the other hand, downtime may be scheduled in order to facilitate faster mirroring and prevent users from having to put up with degraded performance. Mirroring of individual logical volumes should be done in accordance with the migration schedule that was agreed upon by the customer. The following steps describe how to mirror logical volumes.

  - Mirroring individual logical volumes

    Individual logical volumes can be mirrored with the command:

    ```
    mklvcopy -e m -s y -k lvname # hdiskx hdisky
    ```

    Where

    `lvname` = the name of the logical volume without the /dev/ prefix

    `#` = the total number of copies desired of the logical volume (in most cases, this will be 2)

    `hdiskx` = the source hdisk

    `hdisky` = the target hdisk

    Note: This command will create a synced copy of the logical volume. The amount of time for the system to produce the mirrored and synced copy will depend on the amount of I/O operations performed on the logical volume as well as the amount of data in the logical volume. We have found that our AIX system can mirror 6 GB in 1 hour.

- Mirroring all the logical volumes in a volume group.

  An entire volume group can be mirrored and synced with the command:

  `mirrorvg vgname hdiska hdiskb hdiskc`

  Where

  `vgname` = volume group name

  `hdiska, hdiskb, hdiskc` = the physical disks to use for each copy of the volume group

  Note: This command inherently uses strict allocation, which keeps each copy of a logical volume on a different disk. This command will not mirror a dump device if the dump device is not paging space. Use the `mklvcopy` command instead.

- Split the logical volume copies.

  Once you have made duplicate copies of each of the logical volumes that need to be migrated, the next step is to separate the copies. This will allow you to test the new copies of the logical volumes to ensure that they contain all the customers' information before you remove the old copies.

  - Split each logical volume.

    The following process should be used to split each logical volume

    1. Close the logical volume by either unmounting the filesystem that is mounted on the logical volume or by killing the process that is using the logical volume to be split. You can see if any processes are using a logical volume by using the `fuser` command. You can verify that a logical volume is closed by using the `lsvg -l vgname` command.
    2. Use the following command to remove a copy of the logical volume from the mirrored copies. This command will separate a copy from the mirrored copy or copies and rename that copy, thereby creating a separate, new logical volume on the specified hdisk.

       `splitvcopy -y newlv -c # hdiskx`

       Where

       `-y` = indicates that you will supply the name of the logical volume that will be created

       `newlv` = the name of the new logical volume to be created

       Note: It is recommended that you use a naming convention that is consistent for all the new logical volumes that will be created by the `splitlvcopy` command. This way, it will be easy to recall the names of the new logical volumes if you need to remove them after testing. For example, if splitting a logical volume called lv00 or hd6, use the names xlv00 and xhd6, respectively.

       `-c` = indicates that you will designate the maximum number of copies that will be left in the logical volume that you are copying from

       `#` = the number of copies that will be left in the logical volume that you are copying from

       `hdiskx` = the hdisk that you specify here will contain the renamed copy of the logical volume

    - Remount the filesystem.

The following process should be used to remount each filesystem, thereby making the data on the logical volume available to users and applications

1. Verify that you have the original logical volume (the same name that is in etc filesystems, not the name of the new logical volume that was created by the `splitlvcopy` command) on the new, ESS disk.

2. Remount the filesystem by using the mount command or using smit

- Test the newly migrated data.

  Now that you have a new copy of the data on ESS disk, you can go ahead and test all data according to the test plan designed by the customer.

- Remove the old copies of data.

  You may or may not want to remove old logical volumes depending on whether the old storage will be removed from the system.

  - Remove each logical volume copy.

    You will need to recall the names you gave to the new logical volumes created by the `splitlvcopy` command.

16.4 Data migration for Windows NT

You can migrate data to an ESS by using standard Windows NT commands. No special tools or methods are required for transferring data to ESS. Windows NT recognizes the logical volumes in an ESS as normal physical SCSI disks.

16.4.1 Migration methods

There are several different methods you can use for migration data in Windows NT environments. You can use for example basic MS-DOS commands, such as Copy, Backup, Restore or you can use drag and drop technique to achieve the same result. You can find on the market also many products from different vendors who provides the storage management software for migration and copying. Widely accepted such a product is IBM ADSTAR Distributed Storage Manager (ADSM).

16.5 IBM migration services

The IBM Migration Services for the ESS provides an IBM Services Specialist at the customer's location to plan and migrate data and applications from IBM AIX, HP/UX, Sun Solaris, and Windows NT servers to the ESS. This migration is accomplished with minimum interruption to service. In addition, IBM will provide a migration control book which specifies the activities performed during these Services. For data migration the following methods will be used

- NSI (Network Specialists Incorporated) Doubletake (for Solaris and Windows NT)
- Logical Volume Mirroring (for UNIX)
- Other

IBM will perform the tasks explained in the following sections.

16.5.1 Migration Planning

IBM will conduct up to a two day planning session with you to

- Assess your migration readiness
- Review for completeness the IBM Storage Server (ISS) configuration worksheets, which identify your host names, host types and storage attached to your current server
- Identify the specific applications and data to be migrated to the ESS, including identifying the volumes specified for migration
- Develop a test plan with your representative that will be used to verify the integrity of the data upon completion of the migration
- Establish sequence of events, including scheduling migration activities and identifying your responsibilities
- Develop a data migration plan that includes a complete description of the data and applications to be migrated, target locations for migration within the ESS, and timing for the migration

16.5.2 Data migration

IBM will:

- Review the data migration plan with you
- Verify you are at the required system software levels
- Perform migration activities, as defined in the data migration plan
- Perform post migration testing and verification
- Create and deliver a Migration Control Book Material which contains:
  - The ISS configuration worksheets
  - The test plan
  - Data migration plan
  - A record of the volumes which have been migrated

Chapter 17. Managing and monitoring the ESS

In this chapter we discuss the tools that are available for managing and monitoring the ESS. Each operating system has some form of storage management capability. In addition to this, you can use IBM StorWatch family of products.

17.1 IBM ESS Expert

The IBM ESS Expert is a software tool designed to help storage administrators manage and monitor IBM storage servers. It is an optional product. For more information and sample reports see 6.1, “Using ESS Expert” on page 93.

17.2 AIX

There are the following main methods of controlling SCSI disks on AIX

17.2.1 AIX System Management Interface Tool (SMIT)

You can use SMIT to control Adapters, Disks, and Arrays.

17.2.2 AIX diagnostics

With the `diag` command you can invoke a standard AIX diagnostics menu. This menu lists these aids:

- Set service mode
- Link verification
- Configuration verification
- Format disk
- Certify disk
- Identify

17.2.3 Special tools

A variety of tools have been developed for resolving a specific situation. Most of these commands are an integral part of the latest AIX releases and of the SMIT panels. The script files are useful and the description of how to use the commands in the tools README file produce a good understanding of exactly what is happening.The main ones are described in the following sections.

17.2.4 Other tools

AIX incorporates some standard UNIX commands that are very useful for measuring disk performance.

There are several tools and commands on AS/400 you can use for managing and monitoring disks.

17.2.5 Dedicated Service Tools (DST)

The DST is part of the service function used to service the system when the operating system is not running. With this function you can add new disks, move the disks between auxiliary storage pools, and perform other functions.

17.2.6 System Service Tools (SST)

The SST is part of the service function that can be used when the system is up and running. When the command `STRSST` base menu is displayed, you can select option *Work with disk units*. After that you can select:

- Display disk configuration
- Work with disk configuration
- Work with disk unit recovery

All of the following menus will allow you to access and start additional functions.

17.2.7 WRKDSKSTS command

The command `WRKDSKSTS` will show you the status of all installed disks, occupancy percentage, the number of I/O operations, and more.

17.2.8 Go `DISKTASKS` Command

This command displays menus that allow you to select these options:

- Collect disk space information
- Print disk space information

17.2.9 AS/400 operations navigator

The AS/400 operations navigator provides a Windows interface to common AS/400 management functions. AS/400 Operations Navigator is supported on Windows 95, Windows 98, and Windows NT 4.0. Some of the common management functions include Basic Operations, TCP/IP Configuration, Job Management, Users and Groups, and Database Management. Operators configure AS/400 systems to be managed. Operations Navigator will connect directly to the AS/400 from the Windows workstation using TCP/IP to provide management of that single AS/400 system. Management Central is an extension of Operations Navigator that is accessed through a toolbar button on the Main Operations Navigator window. Management Central allows operators to monitor multiple AS/400 systems with the same ease that Operations Navigator allows them to perform other basic management functions on a single system.

You can define, for example, one monitor to watch several metrics that correspond to response time, and another monitor could keep track of communications lines. Metrics include:

- CPU utilization
- Interactive response time
- Transaction rate
- Batch logical database I/O
- Disk arm utilization
- Disk storage

- Disk I/O processor (IOP) utilization
- Communications IOP utilization
- Machine pool faults
- User pool faults
- Communications line utilization
- LAN utilization

Once you have the basics down, you can explore the advanced features. Some of these include: changing the collection interval to collect metric data less frequent, ignoring data spikes by extending the threshold duration, setting threshold triggers and reset values, and performing automation.

After you have your monitors defined, you simply press the start button and select one or more system groups or individual endpoint systems to run your monitor. The monitors window will give you up-to-date status information for each of your defined monitors. You will be able to see which monitors have started, any monitored systems that have failed, and which monitors have reached a threshold. A monitor may run for minutes, hours, days, or until you stop it. Even though a monitor is running, you are free to perform other functions in Management Central, Operations Navigator, or on your PC. You may also quit your Operations Navigator session or even shut down your PC and your monitors will continue to run. The next time you enter Management Central you can check the status of the monitors you are running.

17.2.9.1 Monitor graph

When you open a monitor that you started, you will see the monitor graph window. The monitor graph window shows a separate real-time data collection graph for each metric defined in the monitor. If the monitor is collecting metric data from multiple endpoint systems or system groups, all systems will be displayed in each metric graph.

17.2.9.2 Events and automation

If you defined your monitor to watch for thresholds, you can also set actions for when the threshold is triggered and reset. Actions on the PC include: opening the graph window, opening the event log window, sounding an alarm, and performing automation. Actions on the AS/400 include logging an event and performing automation.

If log an event is set, you will be able to see these in the event log. You can filter the event log to see all events for a particular system or group, for a particular metric, or for all the events associated with a specific monitor. Each event contains helpful details like the monitor threshold value, the actual metric data collected value, system that sent the event, and whether any automation occurred.

17.3 Windows NT

17.3.1 Disk administrator

The disk administrator is a graphical tool for managing disks in a Windows NT environment. This tool encompasses and extends the functionality of character-based disk management tools.

With the disk administrator you can:

- Create and delete partitions on a hard disk and logical drives within an extended partition.
- Format and label volumes.
- Read status information about disks such as the partition sizes and the amount of free space that is available for creating additional partitions.
- Read status information about Windows NT volumes such as the drive-letter assignment, volume label, file system type, and size.
- Make and change drive-letter assignments for hard disk volumes.
- Create and delete volume sets.
- Extend volumes and volume sets.
- Create and delete stripe sets.

Chapter 18. ESS Specialist Copy Services for Open systems

Copy Services provides a Web-based interface for setting up and managing Peer-to-Peer Remote Copy (PPRC) and FlashCopy. PPRC operates at the volume level from one SCSI bus target/LUN group (LSS), to the volume level within another LSS. FlashCopy also operates at the volume level, however the source and target volumes are within the same SCSI bus target/LUN group (LSS).

PPRC and Flashcopy are used as a disaster recovery, data backup tool, creation of test data or for data migration.

18.1 Copy Services considerations

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

There are a few restrictions to consider:

- Only one FlashCopy at a time can be active on a volume, however, you can perform a PPRC concurrently with FlashCopy on the same volume.
- The primary and secondary volumes must reside within an ESS, for both PPRC and FlashCopy.
- You need to manage PPRC using the Copy Services component of the ESS Specialist, therefore, Ethernet, and TCP/IP connectivity is needed between the two participating ESS subsystems, primary and secondary, and the Web browser initiating and managing the PPRC activities.
- When using Copy Services you are required to nominate via the CONFIGURATION panel, one ESS as the Copy Services Primary Server that is the central place for collection of all information. A second ESS can be nominated as the backup server. You will need to carefully plan which ESS you will nominate in which role to ensure maximum availability.
- The source and target logical volumes must be the same, or larger in size.

You will also need to check that the CRIT=YES - Light or Heavy options have been set correctly in the ESS VPD by your hardware specialist. See 2.3.1.4, “VPD settings” on page 19.

You will need to have the hardware engineer check that all the ESS cluster hostnames were added to the hostname list during installation. Otherwise the cluster hostname will not be found when defining the Copy Services primary.

18.2 PPRC

Peer-to-Peer Remote Copy (PPRC) provides a synchronous copy, mirroring (RAID 1) of a source LUN to a secondary LUN. Data updates to the source device are also sent to the secondary device. The host is notified that the update has been completed when both primary and secondary ESS have the update safely in both NVS and cache.

PPRC will only operate over ESCON connections, and therefore, even if it is being used in an Open Systems environment, ESCON adapters and cabling are required.

18.2.1 Rules for configuring PPRC links

The following are the PPRC configuration rules for establishing Escon connections between two ESS subsystems:

- Up to eight Escon links are supported between two ESS subsystems.
- A primary ESS can be connected via Escon links to up to four secondary ESS subsystems.
- A secondary ESS subsystem can be connected to any number of primary ESS subsystems, limited by the number of Escon links available.
- PPRC links are unidirectional, because the ESCON port at the primary ESS is reconfigured to act like an ESCON channel in a host S/390 processor. The primary ESCON port is dedicated to PPRC.
- The ESCON protocol has been streamlined with less handshaking and larger frames transmitted between ESS. This has enabled the PPRC Escon links to be extended to up to 103 km.

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

18.2.2 How to invoke PPRC

The ESS Specialist shows volumes (LUNs) by their ESS internal serial numbers. You will first have to find out what the serial number of the volumes you intend to copy (source and target) from your operating system is, in order to identify the volume on the Copy Services Volume panels.

To invoke PPRC you need to perform the following tasks from the Copy Services main menu:

1. With the PATHS panel, establish a logical path between the primary ESS LSS and host adapter, and the target LSS and its host adapter.
2. Use the VOLUMES panel to find and select the source and then target volume PPRC pairs and Establish, Suspend or Terminate the data transfer. You can optionally establish both paths and volume pairs from this panel.
3. Alternatively, you can select the CONTROL UNIT panel to initiate or remove PPRC relationships between all volumes on a LCU.
4. You can save previously defined PPRC path and pair definitions as tasks. Using the TASKS panel, you can select and run the pre-saved set of tasks.

During steps 1, 2 and 3, once you have selected the resources that you are working with, a right mouse click on the target resource starts a wizard to guide you through the selection of the appropriate PPRC functions.

To assist you in removing unrequired resource information from the panels, a FILTER button is provided to enable display of selected resources. For example, show only:

- S/390 or Open volumes
- Source or target volumes
- Physical or logical ESS

Each panel also has an INFORMATION button that will display PPRC status and other general information about the selected LCUs, volumes or paths.

Detailed information regarding each panel described above can be found in the *Web Interface Users Guide*, SC26-7346.

18.3 FlashCopy

FlashCopy makes a single point in time (T0) copy of a LUN. The target or copy LUN is immediately available after the command has been processed.

18.3.1 How to invoke FlashCopy

The first step requires you to identify the internal serial number of the logical source and target disks.

To invoke FlashCopy, use the VOLUMES panel to find and select the source and target volumes. Once presented with the wizard window, select Establish or Withdraw the FlashCopy pairs.

As with PPRC, you can save your FlashCopy definitions as tasks and run them at any time by using the TASKS panel. You can also FILTER your displays and use the INFORMATION button.

Once a relationship has been established between a FlashCopy source and target volume, a background task commences that copies the entire source volume to the target. Using the ESS Specialist, you can suppress this copy task by specifying a NOCOPY option.

If NOCOPY has been specified, any data about to be updated on the source volume is first copied to the target volume. Hence, the target volume only contains pre-updated data, not a complete volume copy. The T0 copy of the source is still available for use as long as the source target relationship exists.

This relationship must be terminated (WITHDRAW) using the ESS Specialist. If NOCOPY was not specified, the relationship ends automatically once the background source copy has been completed.

Appendix A. ESS Configuration Planner

The *ESS Configuration Planner* - SC26-7353 is a document which defines the process for planning and implementing an initial logical configuration for the IBM Enterprise Storage server (ESS). Your system administrators are asked to complete a set of initial worksheets describing the desired configuration. The IBM CE uses these worksheets, and the instructions in this document to prepare the detailed information needed to implement the required logical configuration.

The ESS product is delivered in a series of standard physical configurations that support most performance requirements. These configurations are then logically configured to meet your specific storage needs. A given ESS can be configured to support multiple operating system environments when required by a user. In practice, many users will configure the ESS to support a single environment, such as a multi-platform UNIX storage server, or an OS/390 storage server.

The ESS provides a set of standard logical configurations to simplify the process for many of your applications. In some cases, if the standard logical configurations do not satisfy the specific needs of an application, the full flexibility of the product is available using the Storwatch ESS Specialist to provide a customized product configuration.

A set of worksheets are used to organize the information needed by IBM service personnel to configure the ESS. Completion of these worksheets provides the following:

- Names and types of systems, S/390, AS/400, UNIX (AIX)
- Amount of storage capacity each host data base requires
- Data bases, if any, that are shared among multiple hosts
- General information about the customer environment

For a copy of this document, and a copy of the worksheets, please go to the IBM website:

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

Appendix B. Special Notices

This publication is intended to help storage administrators who are involved in the installation and maintenance of storage products and applications in either the S/390, UNIX, AS/400 or NT environments. Many customers will have multiples of these platforms, and this book will help in the implementation process of the ESS. The information in this publication is not intended as the specification of any programming interfaces that are provided by the ESS. See the PUBLICATIONS section of the IBM Programming Announcement for the ESS for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating

environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

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

IBM ®
ESA/390

AIX
ESCON
RS/6000
AS/400
S/390
Versatile Storage Server
VSE/ESA
OS/390
DFSORT

ESS
Enterprise Systems Connection Architecture
StorWatch
400
DFSMS
Netfinity
Parallel Sysplex
VSS
MVS/ESA
System/390

The following terms are trademarks of other companies:

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

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

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

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

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

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

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

Appendix C. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 259.

- *IBM Enterprise Storage Server*, SG24-5465
- *DFSMShsm ABARS and Mainstar Solutions,* SG24-5089

C.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

C.3 Other Publications

These publications are also relevant as further information sources:

- *Introduction and Planning Guide 2105 Models E10 and E20*, GC26-7294
- *IBM Data Path Optimizer*, SC26-7291
- *The IBM Enterprise Storage Server Web Interface User’s Guide*, SC26-7346
- *DFSMS/MVS V1R5 DFSMShsm Storage Administration Guide,* SH21-1076
- *DFSMS/MVS V1R5 DFSMShsm Storage Administration Reference,* SH21-1075
- *Planning for IBM Remote Copy,* SG24-2595
- *Host Systems Attachment Guide 2105 Models E10 and E20,* SC26-7296
- *ESS Configuration Planner,* SC26-7353

C.4 Web addresses

The following Web sites contain other relevant information:

http://www.storage.ibm.com

http://home.netscape.com/computing/download/index.html

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** `http://www.redbooks.ibm.com/`

  Search for, view, download, or order hardcopy/CD-ROM redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at `http://w3.itso.ibm.com/` and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access `MyNews` at `http://w3.ibm.com/` for redbook, residency, and workshop announcements.

IBM Redbook Fax Order Form

Please send me the following:

Title	Order Number	Quantity

First name	Last name	
Company		
Address		
City	Postal code	Country
Telephone number	Telefax number	VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date	Card issued to	Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Glossary

array An arrangement of related disk drive modules(DDM) that you have assigned to a group.

bottleneck The narrowest, or slowest, section in the I/O path.

cache A random-access electronic storage in subsystems that you use to keep frequently used data for faster access by a host system.

cluster A partition of a storage server that is capable of performing all functions of a storage server. When a cluster fails in a multiple-cluster storage server, any remaining clusters in the configuration can take over the processes of the cluster that fails.

controller image The System/390 term for a logical subsystem that you access with an ESCON I/O Interface. One or more control-unit images exist in each controller. Each image appears to be an independent controller, but all images share a common set of hardware facilities. The ESS can emulate 3990–3, 3990–6, or 3390–3 TPF control unit images.

destage The process of reading data from cache. Also, the action of storing a logical data unit in cache memory with active write data to the storage device. As a result, the logical data unit changes from cached active write data to cached read data.

disk drive module (DDM) The primary nonvolatile storage medium that you use for any host data that is stored within a storage server. Number and type of storage devices within a storage server may vary.

ESCON Enterprise Systems Connection Architecture. AnESA/390 computer peripheral interface. The I/O interface utilizes ESA/390 logical protocols over a serial interface.

ESCON host systems This term is used to designate System/390 hosts that you attach to the ESS with an ESCON adapter.Such host systems have MVS, VM, or VSE operating systems.

host A computer system that is attached to an IBM storage server. The host send data to and retrieves it from the storage server.IBM Enterprise Storage Server (ESS)A storage server and its attached disk storage devices.

logical device A pointer to a logical volume on a storage server. The pointer enables a host system to access the logical volume.

logical unit The Small Computer System Interface (SCSI) term for a logical disk drive.

logical unit number (LUN) The SCSI term for the field in an identifying message that is used to select a logical unit on a given target.

logical volume The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. For the ESS, you define this unit of storage; the logical volume resides on a RAID–5 array, and is spread over 6 +P or 7 +P drives. A logical volume can also reside on a JBOD (non-RAID array) on one storage device. For CKD the logical volume size is defined by the device emulation mode (3390–2, –3, –6). For UNIX and Windows NT FB, the size is 0.5 GB to 224 GB. For OS/400 FB the size is 2, 4, 6, 8, 16, or 32 GB.

Note: AIX views a logical volume as an hdisk or hard disk. Sun, HP, and Windows NT also view it as a hard disk.

SCSI Small Computer System Interface. An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface uses a SCSI logical protocol.

SCSI host systems This term is used in this publication to designate host systems that you attach to the ESS with an SCSI interface. Such host systems have UNIX, OS/400, or Windows NT operating systems.

SCSI ID A unique identifier (ID) assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs. A SCSI device is either an initiator or a target.

stage The process of reading data into cache from a disk drive module.

volume Refers to a logical volume.

List of Abbreviations

AMS	Access Method Service
BIOS	Basic Input/Output Storage
CC	Concurrent Copy
CE	Custom Engineering
CKD	Count Key Data
CP	Cluster Processor
CPC	Cluster Processing Complex
CPI	Common Parts Interconnect
CSC	Customer Steering Committee
DA	Device Adapter
DDM	Disk Drive Module
DPO	Data Path Optimizer
DS	Deserv
DST	Dedicated Service Tools
EPO	Emergency Power Off
ECC	Error Correcting Code
ESC	Exception System Code
ESCON	Enterprise Systems Connection Architecture
ESS	Enterprise Storage Server
EREP	Environmental Recording and Editing Program
Expert	ESS Expert
FB	Fixed Block
FBA	Fixed Block Architecture
FC-AL	Fibre Channel Arbitrated Loop
GB	Gigabyte
HA	Host Adapter
HCD	Hardware Configuration Definition
IBM	International Business Machines Corporation
IODF	Input Output Definition File
ISV	Independent Software Vendors
ITSO	International Technical Support Organization
JBOD	Just a Bunch of Disks
LCU	Logical Control Unit
LIC	Licensed Internal Code
LSS	Logical Subsystem
LUN	Logical Unit Number
LVM	Logical Volume Manager
MADS	Multi-Area Dataset
MIB	Management Information Base
MosT	Mobile Service Terminal
NVS	Non-Volatile Storage
PAV	Parallel Access Volumes
PE	Product Engineering
PFA	Predictable Failure Analysis
PPRC	Peer-to-Peer Remote Copy
PTF	Program Temporary Fixes
RAID	Redundant Array of Independent Disks
RPC	Rack Power Control
RVA	Rarac Virtual Array
SCSI	Small Computer System Interface
SDM	System Data Mover
SDRAM	Synchronous Dynamic RAM
SFOI	Storage Function Operation Interface
SIC	Serial Interface Chip
SIM	Service Information Message
SMP	Symmetrical Multiprocessors
SMS	Systems-Managed Storage
SNMP	Simple Network Management Protocol
SSA	Serial Storage Architecture
SSID	Subsystem ID
SSL	Secure Sockets Layer
SST	System Service Tools
SODS	Statement of Direction
TDMF	Transparent Data Migration Facility
URL	Universal Resource Locator
VPD	Vital Product Data
VSS	Versatile Storage Server
WLM	Workload Manager
XRC	Extended Remote Copy

Index

T

U

V

W

X

ITSO Redbook Evaluation

Implementing the Enterprise Storage Server in Your Environment
SG24-5420-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to `redbook@us.ibm.com`

Which of the following best describes you?
_ Customer _ Business Partner _ Solution Developer _ IBM employee
_ None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

_____

_____

_____

_____

_____

SG24-5420-00

Printed in the U.S.A.

Implementing the Enterprise Storage Server in Your Environment

SG24-5420-00