IBM TotalStorage Enterprise Storage Server

Implementing the ESS in Your Environment

Learn how to install, tailor, and configure the ESS

Plan for migrating data and changing configurations

Know the tools for monitoring and managing the ESS

Gustavo Castets
Donald (Chuck) Laing
Jukka Myyryläinen
Diane Williams

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM TotalStorage Enterprise Storage Server: Implementing the ESS in Your Environment

March 2002

SG24-5420-01

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page 299.

Second Edition (March 2002)
This edition applies to the IBM 2105 Enterprise Storage Server model F. For other information, please see the Publications section of the IBM Announcement letter for the IBM Enterprise Storage Server. You can also visit our Web site at: `http://www.storage.ibm.com`

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

Summary of changes

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

Summary of Changes
for SG24-5420-01
for IBM TotalStorage Enterprise Storage Server
as created or updated on March 28, 2002.

March 2002, Second Edition

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

New information

This 2002 version of the redbook is an update of the November 1999 original document, and includes the latest announcements on the IBM TotalStorage Enterprise Storage Server:

- IBM TotalStorage Enterprise Storage Server models F10 and F20
- New Fibre Channel/FICON host adapters (short wave and long wave)
- Cache options: 8, 16, 24 and 32GB
- FICON native host attachment
- ESS Master Console
- Flexible configurations
- Disk drive capacity intermix
- Support for 72.8GB disk drives, scaling the ESS total capacity up to 22TB
- Control Unit Initiated Reconfiguration (CUIR) support
- 32K cylinder Large volume support (LVS)
- Linux support for zSeries servers and Intel-based servers
- PPRC and FlashCopy support for iSeries and AS/400
- ESS Copy Services Command Line Interface (CLI) supporting additional platforms
- Subsystem Device Driver (SDD) supporting additional platforms
- Read from secondary
- INRANGE Channel Extender support
- TPF support for PPRC and FlashCopy

Contents

Figures

Tables

Preface

Since the initial availability of the IBM TotalStorage Enterprise Storage Server (ESS) there have been many modifications and additional functionality. The original ESS Implementation Guide was written in November 1999. The 2002 version is an updated guide that includes all of the changes since then.

This IBM Redbook is a guide for the installation, implementation, and administration activities of the ESS in both the S/390 and open systems environments. It will help you plan and accomplish the installation, tailoring, and configuration of the ESS in your environment. It explains how you can use the functions available for the ESS to efficiently manage your disk storage data as well as the ESS, once it is operative.

We cover the latest announcements on the ESS: disk capacity intermix; 72.8 GB capacity disk drive; flexible configurations; Control Unit Initiated Reconfiguration (CUIR) support; large volume support (LVS); read from secondary; ESS Master Console; Subsystem Device Driver (SDD) and the Command Line Interface (CLI) support for additional operating systems; INRANGE Channel Extender support; TPF support for PPRC and FlashCopy.

We also provide information on the new models F10 and F20; FICON native host attachment; new Fibre Channel/FICON host adapters (short wave and long wave); Linux support for Intel-based servers and zSeries servers; iSeries and AS/400 support for copy services; and new cache options.

The team that wrote this redbook

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

Gustavo Castets is a Project Leader with the IBM International Technical Support Organization, San Jose Center. He writes extensively and teaches IBM classes worldwide on areas of Disk Storage Systems. Before joining the ITSO, Gustavo worked in System Sales as Field Technical Support Specialist, and has worked for more than 22 years in many IT areas, in Buenos Aires, for IBM Argentina.

Donald (Chuck) Laing is a Disk Administrator in the south delivery center in Texas. He has 13 years of experience in the UNIX field, and has worked at IBM for 3 years. Before joining IBM, Chuck was a Hardware CE on UNIX systems and taught basic UNIX at Midland College in Midland Texas. He holds a degree in Computer Science from Midland College. His areas of expertise include Open Systems UNIX System/Disk Administration.

Jukka Myyryläinen is an Advisory IT Specialist with IBM Global Services, Finland. He has 15 years of experience in storage product implementations. He has contributed to several storage related redbooks in the past.

Diane Williams is a Consulting IT Specialist with the San Jose Advanced Technical Support group. She has 23 years of experience in several areas of IT. Her areas of expertise include disk performance.

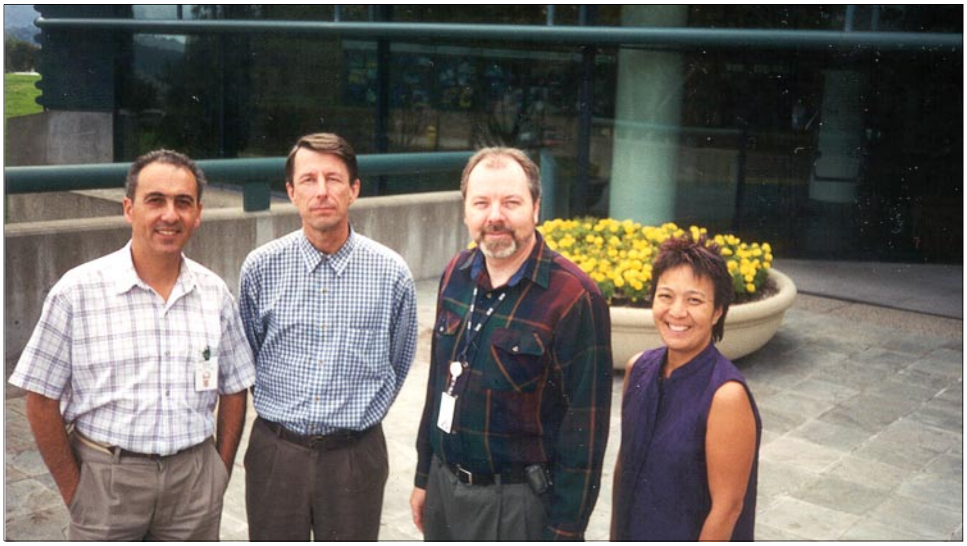

Gustavo Castets, Jukka Myyryläinen, Donald (Chuck) Laing, Diane Williams

The authors of previous editions of this book are:

Mark Blunden, IBM Australia
Ivan Avsic, IBM Slovenia
Ian Black, IBM Australia
Peter Crowhurts, IBM Australia
Carlos Sadao Miyabara, Banco Itau, Brazil

Thanks to the following people for their diligence in reviewing this manual:

Don Maggert, IBM, Tucson, ESP Management Office
James Cammarata, IBM, Chicago, DFSMS Device Support
Jennifer Eaton, IBM, San Jose, User Technology - Information Development
Jim Garadis, IBM, San Jose, IRT - SSG Brand Support
Leslie Barton, IBM, San Jose, DFSMS Device Support
Nicholas Clayton, IBM, UK, EMEA ATS for Storage
Omar F Escola, IBM, Argentina, Hardware Support Specialist
Pat Blaney, IBM, San Jose, Storage Systems ATS
Paul Coles, IBM, UK, ITS Storage Specialist
Paulus Usong, IBM, San Jose, Storage Systems ATS
Phil Lee, IBM, Canada, ESS Configuration HelpLine
Rainer Wolafka, BM, San Jose, Open Systems Validation Lab
Scott Cardinell, IBM, Tucson, Microcode Tools
Steven Winch, IBM, Rochester, iSeries Hardware Test
Thomas Fiege, IBM, San Jose, SSD RAS Engineering

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

Amine Hajji, IBM, San Jose, Advanced Product Development
Amy L Therrien, IBM, San Jose, ESS Web Development
Brian Kraemer, IBM, San Jose, RAS Microcode Develpment
Cindy Sauer, IBM, San Jose, SSG RFA
Edsel Carson, IBM, San Jose, Staff Information Development
Ellen Vuong, IBM, San Jose, ESS Web Development
Gail Spear, IBM, Tucson, ESS Strategy and Architecture - Microcode
Henry Caudillo, IBM, San Jose, User Technology - Information Development
Hilda Wu, IBM, San Jose, Expert Development
Jack Flynn, IBM, San Jose, Device Support Architecture
Jane Madriaga-Miller, IBM, San Jose, ESS Web Development
Jed Dyreng, IBM, San Jose, Service Panel Development
Jeff Ferrier, IBM, San Jose, SSD RAS Engineering
Jeffrey Steffan, IBM, San Jose, DASD RAS
Joe Writz, IBM, Rochester, Storage IO Subsystem Develpment
John Bernauer, IBM, San Jose, RAS Programming
Kurt Lovrien, IBM, Tucson, Microcode Development
Lynsey V Bishop, IBM, San Jose, User Technology - Information Development
Maria Magana, IBM, San Jose, ESS Web Development
Marilyn Pullen, IBM, San Jose, User Technology - Information Development
Michael Washuk, IBM, Tucson, Shark Project Management
Mike Janini, IBM, San Jose, Open Systems Validation Lab
Richard Heffel, IBM, San Jose, Open Systems Lab
Robert Moon, IBM, San Jose, Open Systems Validation Lab
Sheryl Reed, IBM, San Jose, DFSMS Device Support
Stefan Jaquet, IBM, San Jose, DASD Web Development
Steven Winters, IBM, San Jose, DFSMS Optimizer Development
Walter Weaver, IBM, San Jose, Information Development
William H Turner, IBM, San Jose, User Technology - Information Development

Notice

This publication is intended to help IT professionals who are implementing the IBM TotalStorage Enterprise Storage Server in S/390 and open environments. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM or third parties. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM TotalStorage Enterprise Storage Server, Models F10 and F20, for more information about what publications are considered to be product documentation.

IBM trademarks

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

e (logo)®
Redbooks Logo™

AIX®
AS/400®
CICS®
CUA®
DB2®
DFS™
DFSMS/MVS®
DFSMS/VM®
DFSMSdfp™
DFSMSdss™
DFSMShsm™
DFSORT™
DYNIX®
DYNIX/ptx®
Enterprise Storage Server™
Enterprise Systems Connection®
Architecture®
ESCON®
FICON™
FlashCopy™
IBM®
IMS™
IMS/ESA®
iSeries™
Magstar®
MVS™
Notes®NUMA-Q®
NUMACenter™
OS/390®
OS/400®
Perform™
PR/SM™
Predictive Failure Analysis®
PR/SM™
pSeries™
RACF®
RAMAC®
Redbooks™
RETAIN®
RMF™
RS/6000®
S/390®
Seascape®
Sequent®
SP™
StorWatch™
System/390®
Tivoli®
TotalStorage™
VM/ESA®
VSE/ESA™
xSeries™
z/OS™
z/VM™
zSeries™

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to the address on page ii.

Part 1

Characteristics, planning, and installation

In this part of the book we present the characteristics of the IBM TotalStorage Enterprise Storage Server. We then discuss considerations for the planning and installation activities you will be conducting for your ESS implementation. Finally, we present two ESS specific tools, the ESS Specialist and the ESS Expert. This part of the book is intended for readers from both open systems and S/390 environments.

1

Introduction and positioning

In this chapter we provide an overview of the IBM TotalStorage Enterprise Storage Server and the functions and features it provides. We also offer some positioning information to give you an understanding of when the ESS could be used in your environment. In addition, we present an overview of the terminology and definitions that will be used throughout this book.

This chapter covers the following topics:

- Overview of the IBM TotalStorage Enterprise Storage Server components
- IBM TotalStorage Enterprise Storage Server positioning
- IBM TotalStorage Enterprise Storage Server functions and features
- An explanation of general terms used throughout this book

1.1 Overview

The IBM TotalStorage 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 Storage Architecture (SSA) 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, or a power supply fails.

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 portions of multiple disk drives.

The following host I/O interface attachments are supported:

- SCSI-3 Parallel Interface
- Fibre Channel (FCP)
- ESCON
- FICON

On SCSI interfaces, the ESS emulates a variety of fixed block devices with either 512 or 520 byte blocks. Host systems with SCSI-2 or SCSI-3 interfaces can attach to the ESS. The ESS provides multiple SCSI I/O interfaces (buses), 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 FCP the ESS emulates a variety of fixed block devices with either 512 or 520 byte blocks. It allows 1 target per port, with 4096 LUNs per target in the ESS, and up to 128 concurrent port logins per Fibre Channel port (maximum 512 concurrent port logins per ESS subsystem). FCP ports can be configured as point-to-point or fibre channel arbitrated loop (FC-AL). FCP servers can be directly attached to an ESS, or can be SAN attached using a fibre channel director or switch.

On ESCON and FICON interfaces, the ESS runs as a 2105 control unit with one or more S/390 logical control units (LCUs) attaching variable size IBM 3390 devices in either 3390 or 3380 track format. The ESS provides multiple ESCON and/or FICON interfaces and a set of control unit images, each with multiple disk devices. The storage provided by the ESS for ESCON and FICON interfaces is configured so that it is accessible from any of the ESCON or FICON ports.

Fibre Channel/FICON host adapters for SCSI-FCP or FICON attachment are only available on the ESS models F10 and F20 (FC 3021 for long wave and FC 3023 for short wave). The original short-wave Fibre Channel host adapters (FC 3022) are available on both the previous “E” models and the current “F” models of the ESS.

The ESS is composed of the following major components:

- **Host adapters:** The ESS can have up to sixteen host adapters (HA). Each host adapter provides one or two host interfaces. A host adapter can communicate with either cluster complex.

- **Clusters:** 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:
  - **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 (CPs) configured as symmetrical multiprocessors (SMPs). The cluster processors execute the licensed internal code (LIC) that controls operation of the cluster.
    - **Cluster memory/cache:** The cluster memory or cluster cache is used to store instructions and data for the cluster processors. The cache memory is used to store the data read from, or to be written to, the disk drives. The cache memory is accessible by the cluster complex, by the device adapters in the cluster, and by the host adapters.
    - **Nonvolatile storage:** The nonvolatile storage (NVS) is used to store a nonvolatile copy of active written data. In normal operation, the NVS in one cluster stores active write data belonging to logical disks managed by the other cluster. The ESS always has two copies of write data — one in volatile cache on one cluster and one in NVS on the other cluster — so that data integrity is ensured if a cluster fails before data is destaged to disk.
  - **Device adapters:** Each cluster has associated four device adapters (DAs). Each device adapter provides two serial storage architecture loop 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.
  - **Disk drives:** These provide the primary nonvolatile storage medium for any host data stored within the ESS. The disk drives 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 Storage Area Network (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 demands of current and future applications.

Following are some of the many factors that make the IBM ESS an ideal SAN solution:

- Support for all major server platforms, including zSeries, iSeries, pSeries, xSeries, other Windows NT and Windows 2000 servers, other Netware servers, and many varieties of UNIX
- Fibre Channel attachment capability
- Extensive management capabilities through a Web interface
- Excellent scalability:
  - From 420 GB to over 22 TB
  - Drive capacity intermix (the ability to mix drive sizes), thus giving you an increased number of configuration options
- Performance optimized to your heterogeneous environment needs
  - High bandwidth and efficient algorithms that provide solutions for both online and batch applications

- Innovations such as Parallel Access Volumes (PAV) to reduce resource contention and dramatically improve performance in the z/OS environment

- Availability required to support e-business applications:
  - Non-disruptive access to data while making a copy using Concurrent Copy in the zSeries environment
  - Business continuity through remote copy services — Peer-to-Peer Remote Copy (PPRC) for all platforms, and Extended Remote Copy (XRC) for the z/OS users
  - Rapid data duplication through FlashCopy providing extensive capabilities to exploit, manage, and protect your information in a 24 x 7 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. These functions include the following:

- **ESS Copy Services:** These optional functions include the following:
  - **Flashcopy:** This offers the ability to create time zero (T0) copies of logical volumes.
  - **Peer-to-Peer Remote Copy (PPRC):** This gives the ability to create synchronous volume copies over ESCON links.
  - **Extended Remote Copy (XRC):** The ability to create asynchronous volume copies over long distances for the z/OS users
  - **Concurrent Copy (CC):** For the zSeries environment, this is the ability to create volume or data set copies locally.
- **IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist):** This is the Web interface tool used to manage the ESS logical configuration.
- **High performance features:** These are hardware or software features which include the following:
  - **Parallel Access Volumes (PAV):** This optional feature allows you to dramatically reduce or eliminate IOSQ time on single z/OS or OS/390 images.
  - **Multiple Allegiance:** This feature allows you to dramatically reduce or eliminate PEND time for a multiple image z/OS or OS/390 environment.
  - **Fast write with full data integrity:** Data is written to cache and NVS before I/O completion is sent to the server issuing the I/O. Applications with sensitivity to write I/O service times, such as database logging, benefit significantly with this feature.
  - **SCSI command tag queueing:** For open systems logical disks, the ESS can perform multiple I/Os in parallel, similar to the PAV in the zSeries processors. This improves throughput to the logical disks since cache hits and I/Os involving different RAID disk groups can be processed in parallel.
  - **Performance Enhanced CCW commands:** These are provided for the zSeries users.
  - **I/O Priority Queuing:** This allows z/OS users to define priority of application workloads.
  - **Adaptive caching:** The ESS uses adaptive caching algorithms to determine how much data to stage from disk to cache in response to a cache miss.

- **Sequential detect:** The ESS detects a sequential read pattern and begins prestaging data from disk to cache. This allows parallel access to the disks in an array, increasing the sustained throughput the ESS can deliver for a single sequential read stream.
- **Full-stride RAID-5 writes:** The ESS can accumulate full RAID-5 stripes in cache before initiating a destage operation from cache to disk. For sequential write streams, writing full RAID-5 stripes in a single operation avoids the traditional RAID-5 write penalty and allows significantly higher sequential bandwidth. Server applications need not be RAID-5 aware in order to benefit from full-stride RAID-5 writes in the ESS.
- **Custom Volumes:** For zSeries users, this is the ability to create your own custom-sized logical volumes.
- **Large Volume support (LVS):** This is the ability to define a 32,760 cylinder 3390-9 volumes for the z/OS and z/VM users.

- **Scalability features:**
  - Disk capacity is scalable from 420 GB to over 22 TB.
  - Disk drive capacity intermix offers flexible configuration options.
- **Availability feature for S/390 servers:**
  - Control Unit Initiated Reconfiguration (CUIR): This helps by automating channel path quiesce and resume actions thereby reducing the manual actions required during selected ESS service actions.
- **Serviceability feature:**
  - Machine reported product data (MRPD): This is an aspect of vital product data (VPD) which facilitates system management of the ESS. This provides the ESS with the capability to report on its current hardware and software configuration. This information can be transmitted back to IBM electronically, through fax, hardcopy or diskettes.

1.3 Explanation of terms

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

- One decimal gigabyte (GB): 1,000,000,000 bytes. Throughout this book, all disk capacities are expressed in decimal form
- One binary GB: 1,073,741,824 bytes. Cache, NVS and memory capacities are expressed in binary terms
- S/390: This refers to the System 390, zSeries, 9672 G5 and G6, or previous S/390 processors.
- Open: This refers to SCSI or Fibre Channel attached systems, such as UNIX, Intel based, xSeries, iSeries, or pSeries environments
- Device end: Task complete or I/O complete (in the zSeries I/O protocol)
- Task end: Task complete or I/O complete (in the SCSI command-layer protocol)
- LUN (Logical Unit Number): ESS logical unit number that is an open systems unit which is created on an ESS array
- hdisk: Device label referred to by open systems operating systems such as pSeries running AIX
- vpath: A virtual path to a LUN which has multiple SCSI or fibre paths to it. It is the host device label used in place of a LUN that has only one path or connection to it.

- ESCON: Enterprise Systems Connection Architecture which is 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 FB or SCSI devices.
- CKD (Count Key Data): Track format used for emulated 3390 logical devices used by S/390 systems and accessed through ESCON or FICON ports. CKD devices in the ESS are often referred to as S/390 devices, ESCON or FICON 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.
- Host adapter: To those familiar with UNIX, a host adapter is an adapter on the host. Throughout this book, a host adapter (HA) is a SCSI, ESCON, FICON or Fibre Channel adapter which resides in the ESS and is used to connect the ESS to the host server bus.
- LSS (Logical Subsystem): This equates to a S/390 logical control unit (LCU). 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 associated with an individual device adapter. The LSS is made up of ranks managed by a single device adapter. Ranks on different loops both managed by a single device adapter can be part of the same LSS
- Disk group: This refers to a group of eight drives that can be defined as either a RAID array or non-RAID ranks.
- SSR: This refers to the IBM System Support Representative who is responsible for the physical installation and maintenance of the ESS.
- FICON (FIber CONnection): This refers to an ESA/390 and zSeries computer peripheral interface. The I/O interface uses ESA/390 and zSeries FICON protocols (FC-FS and FC-SB-2) over a Fibre Channel serial interface that configures attached units to a FICON supported Fibre Channel communication fabric.
- CUIR (Control Unit Initiated Reconfiguration): This refers to an S/390 feature which automates channel path quiesce and resume actions in support of selected ESS service actions.

1.4 Platform support

The ESS provides support for many different platform connections. These include most forms of UNIX, iSeries, Intel based servers running NT, Windows 2000, Novell Netware, or Linux, and the zSeries processors. For a detailed description of the latest server support for both hardware and software levels please refer to the following Web address:

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

1.5 Getting ESS-related information

Along the different steps needed to implement the ESS in your environment, the people involved in the installation activities will be referencing other ESS-related manuals. In this section we summarize the ESS documentation and manuals that are referenced throughout this redbook, and explain where to get that information.

1.5.1 ESS Web site

The following URL takes you to an ESS Web site where you can start most of your searches for additional information on the IBM TotalStorage Enterprise Storage Server:

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

The Web page you will see at this URL is shown in Figure 1-1.

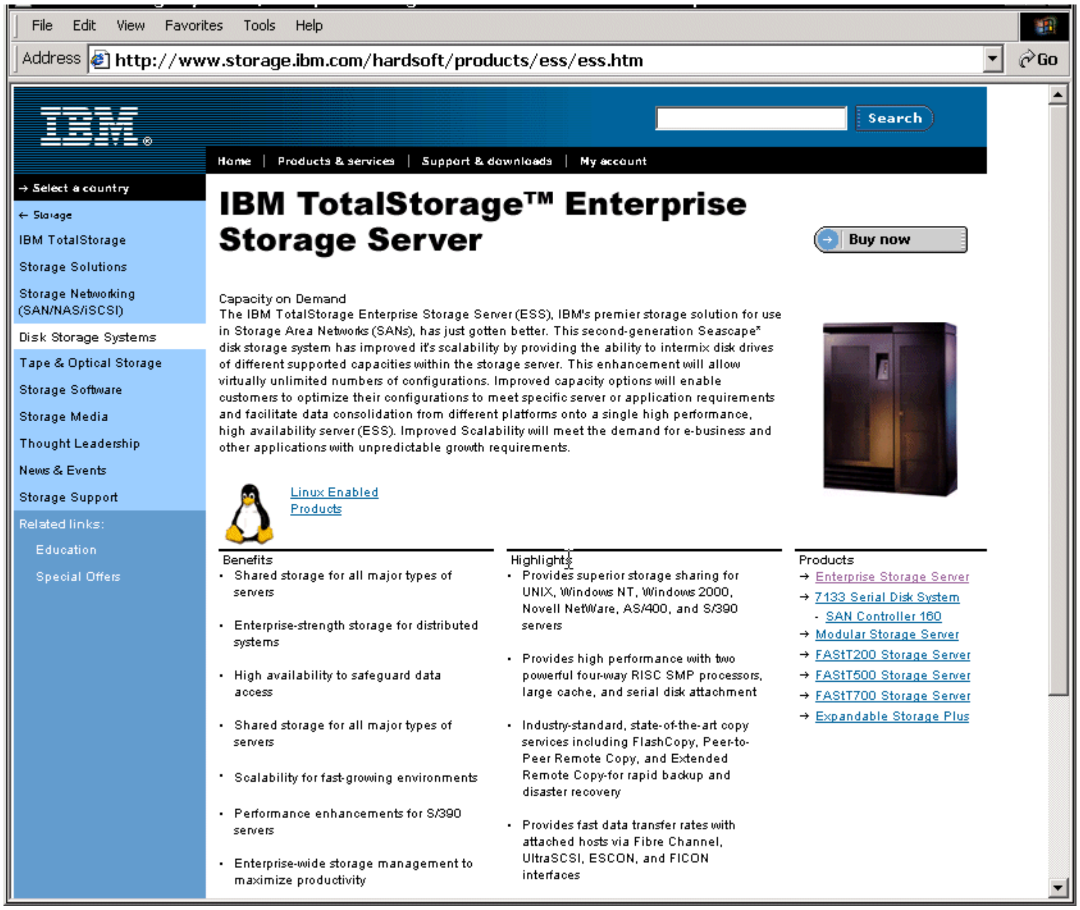

Figure 1-1 ESS: Starting Web site

If you move down the scroll bar to see the bottom of this Web page, you will see a *More Information* area in the page (see Figure 1-2).

Figure 1-2 The More Information area at the ESS Web page

This *More Information* area of the ESS-related Web pages will be the starting point for much of the information and material you will be needing during your implementation activities.

1.5.2 List of supported servers

Click the option *Supported Servers* (see Figure 1-2) and you will reach the following URL:

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

From this ESS Web page, you can make the following selections to get the corresponding material:

- *ESS Open Systems Support Summary*
- *ESS with IBM zSeries and S/390 Systems Support Summary*
- *ESS Supported Servers*. This is a PDF download with the most complete and detailed support information for server attachment
- *SDD (Subsystem Device Driver)*. If you click this option, it will take you to the SDD Web site at:

  http://ssddom01.storage.ibm.com/techsup/swtechsup.nsf/support/sddupdates

  At this site you will find SDD installation support material (downloads, fixes). From here you can also get, whether viewing online or download as PDF, the *IBM TotalStorage Enterprise Storage Server Subsystem Device Driver Installation and User's Guide,* GC26-7442.
- *ESS HBA FC Supported Versions*. This is the Host Bus Adapter Supported Fibre Versions Matrix. This is a matrix which correlates the different models of Fibre Channel I/O adapters to the drivers and to the server operating system level requirements.
- *ESS Host Systems Attachment Guide*. When clicking this option, you either view online or you download the PDF of the *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide 2105 Models E10, E20, F10 and F20,* SC26-7296 manual.
- *ESS Copy Services Command Line Interface User's Guide*. When clicking this option, you either view online or you download the PDF of the *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface Reference,* SC26-7434 manual.

1.5.3 ESS documentation Web site

If when at any of the ESS-related Web pages you click the option *Reference Information* (see Figure 1-2) you will reach the following URL:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

The page you will see is shown in Figure 1-3.

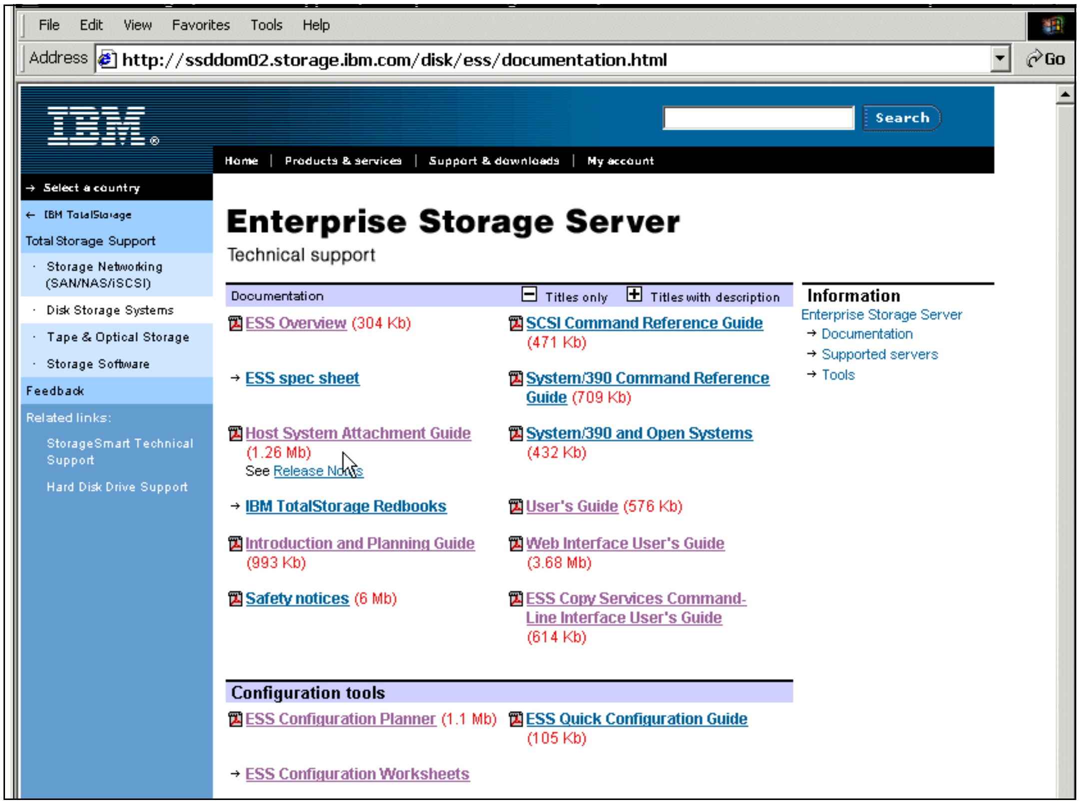

Figure 1-3 ESS documentation Web site

From this Web site, you can either view online (left click) or download the corresponding PDFs (*Save Target As* option after right clicking) for the following ESS documentation:

- An ESS overview of the ESS characteristics
- The ESS specifications sheet
- The *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide 2105 Models E10, E20, F10 and F20,* SC26-7296
- The *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294
- The *IBM Enterprise Storage Server SCSI Command Reference 2105 Models E10, E20, F10 and F20,* SC26-7434
- *IBM Enterprise Storge Server System/390 Command Reference 2105 Models E10, E20, F10 and F20,* SG26-7298
- The *IBM TotalStorage Enterprise Storage Server User's Guide 2105 Models E10, E20, F10 and F20,* SC26-7296
- The *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7346
- The *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface Reference,* SC26-7434
- The *IBM TotalStorage Enterprise Storage Server Configuration Planner,* SC26-7353
- The *IBM TotalStorage Enterprise Storage Server Quick Configuration Guide,* SC26-7354
- The *ESS Configuration Worksheets*

1.5.4 Redbooks

Besides this redbook, for your ESS implementation planning and hands-on activities you will be referencing other ESS-related redbooks that go into more detail on specific tasks involved in the implementation process.

To obtain these related redbooks, you must go to the IBM Redbooks site at:

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

Alternatively, you can enter:

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

Once here, you click the *Redbooks Online* option (on the left navigation bar), and you can get to:

`http://publib-b.boulder.ibm.com/Redbooks.nsf/portals/`

From this site, you can make the appropriate selections to get the redbooks you may be needing to reference.

2

Configuration planning

This chapter overviews 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 configuration planning process will most likely be an iterative process until all parties involved in the planning agree that a valid configuration has been reached.

2.1 Introduction

The Enterprise Storage Server (ESS) is designed for sharing its storage capacity between a mix of both S/390 and Open 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 2-1.

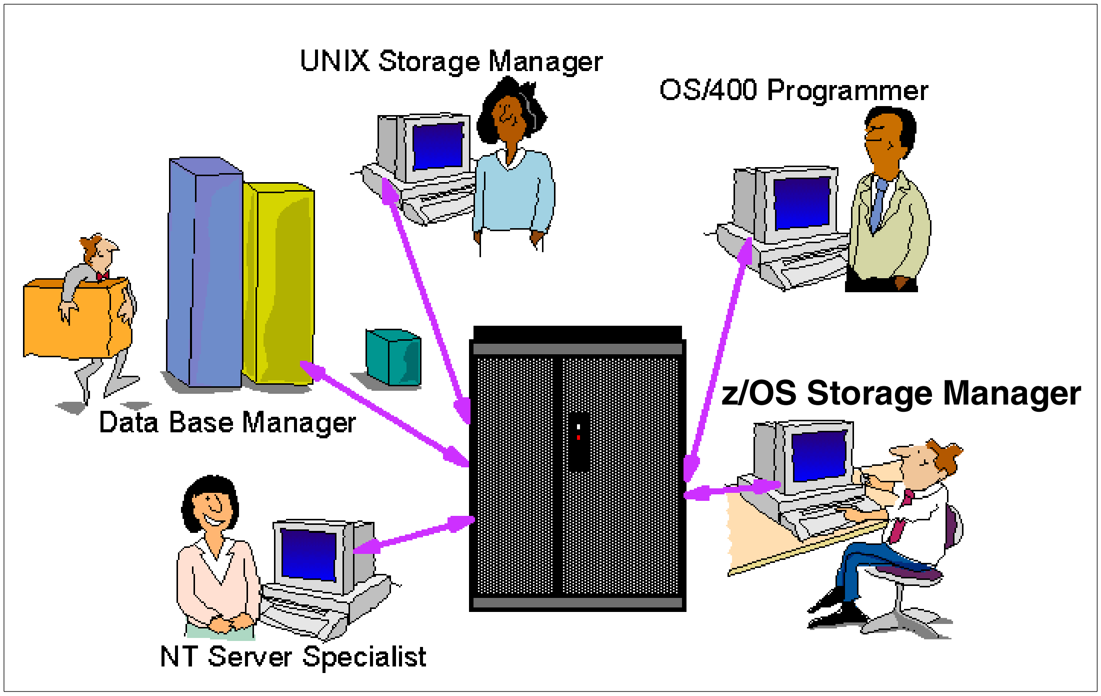

Figure 2-1 The expanded planning team

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's view of the storage allocation will be, are used to build the ESS logical storage configuration.

This logical plan is then used to configure the ESS internally using the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) as part of the installation process. For further information on planning the ESS configuration refer to the IBM publication *IBM TotalStorage Enterprise Storage Server Configuration Planner,* SC26-7353 (see Section 1.5, "Getting ESS-related information" on page 8 for information on how to get this IBM manual).

2.2 Capacity planning

This section is intended to aid the sizing of the physical storage capacity required for your IBM TotalStorage 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 and /or FCP 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 IBM Redbook entitled *IBM TotalStorage Enterprise Storage Server,* SG24-5465.

2.2.1 Capacity and configuration rules

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

Physical features

Physical features of the ESS include:

- Maximum 16 host adapter (HA) cards, ESCON, SCSI, Fibre Channel/FICON, or combination of those HAs. ESCON and SCSI adapter cards have two ports, therefore, you can have a maximum of 32 physical host connections when you use ESCON, SCSI or a combination of only ESCON and SCSI HA. Fibre Channel adapter cards, for FICON and SCSI over Fibre Channel (FCP), have a single port, thereby giving you a maximum of 16 host connections when using only Fibre Channel HA cards.(Fibre Channel host adapter cards do not support the simultaneous transmission of both protocols, they must be configured for one or the other.)
- Four Device Adapter (DA) pairs with two loops per DA pair for a total of 8 loops on the ESS.
- Two to six 8-pack arrays per loop.
- All the drives within an 8-pack are of the same capacity - 9.1 GB, 18.2 GB, 36.4 GB or 72.8 GB.
- When the disk packs on the loops are of the same capacity, 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.
- Multiple drive sizes are supported on the same loop.The first two RAID-5 arrays configured on each loop for each drive size are configured as 6+P+S, providing two spares for each drive size.

Logical features

Logical features of the ESS include:

- The arrays in a loop (maximum 6) may be assigned to any of the 8 LSSs on the DA pair.
- You can have RAID and JBOD ranks on the same LSS or loop.
- You can have CKD and FB formatted arrays/ranks on the same loop, but not on the same LSS.
- A rank can belong to only one of the LSSs on a DA pair.
- The CKD ranks on both loops accessed by one device adapter can be assigned to 1 or 2 of the LSSs of the device adapter.
- The FB ranks on both loops accessed by one device adapter can be assigned to 1 or 2 of the LSSs of the device adapter.

0/8/16 CKD LSSs per ESS

- S/390 ESCON host:
  - Maximum 4096 devices per ESS
  - Maximum 256 devices per LCU
  - Maximum 1024 devices per ESCON channel (CHPID)
  - 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
- S/390 FICON host:
  - Maximum 4096 devices per ESS
  - Maximum 256 devices per LCU
  - Maximum 4096 devices per FICON channel (ESS maximum devices)
  - Maximum 256 logical paths per FICON port
  - Maximum 128 logical paths per LCU
  - Maximum 2048 logical paths per ESS

0/8/16 Fixed Block LSSs per ESS

- SCSI host:
  - Maximum 960 LUNs (15 target x 64 LUN) per ESS host adapter port. This maximum is supported by the ESS but may be further limited by the host operating system support.
  - Maximum 256 devices per LSS, and 4096 FB devices per ESS.
  - 1 to 15 SCSI targets per SCSI bus/port.
  - Either up to 8 or up to 32 LUNs per target, depending on the operating system support.
  - Up to four initiators (hosts) per SCSI port, for selected operating system platforms.
  - LUNs may be shared between hosts, with hosts controlling contention.
  - The Subsystem Device Driver (SDD) manages all the LUNs configured on the ports with which it works.
- FCP host:
  - Maximum 256 LUNs per target (port login) for host that do not support the SCSI command “Report LUNs”, and 4096 LUNs per target (port login) for hosts that support the SCSI command “Report LUNs”.
  - One target per Fibre Channel port login.
  - Up to 128 port logins per ESS Fibre Channel host adapter port. Each of the up to 128 server host I/O adapters performing a port login to an ESS Fibre Channel host adapter port will all see LUNs on a single SCSI target, but they’ll potentially see different LUNs.
  - Maximum of 512 port logins per ESS.
  - Maximum 256 devices per LSS, and 4096 FB devices per ESS.
  - LUNs have affinity to the server Fibre Channel I/O adapter, independent of the ESS Fibre Channel host adapter port the server is attached to. Three implementations can be used to avoid integrity problems: install Subsystem Device Driver (SDD) in the server; create zones in the switch fabric; use the ESS function “Access Control by Host by Port” to restrict the host’s access to a single port.

S/390 host view

Figure 2-2 shows an example of the S/390 logical path limits and considerations for an ESCON implementation.

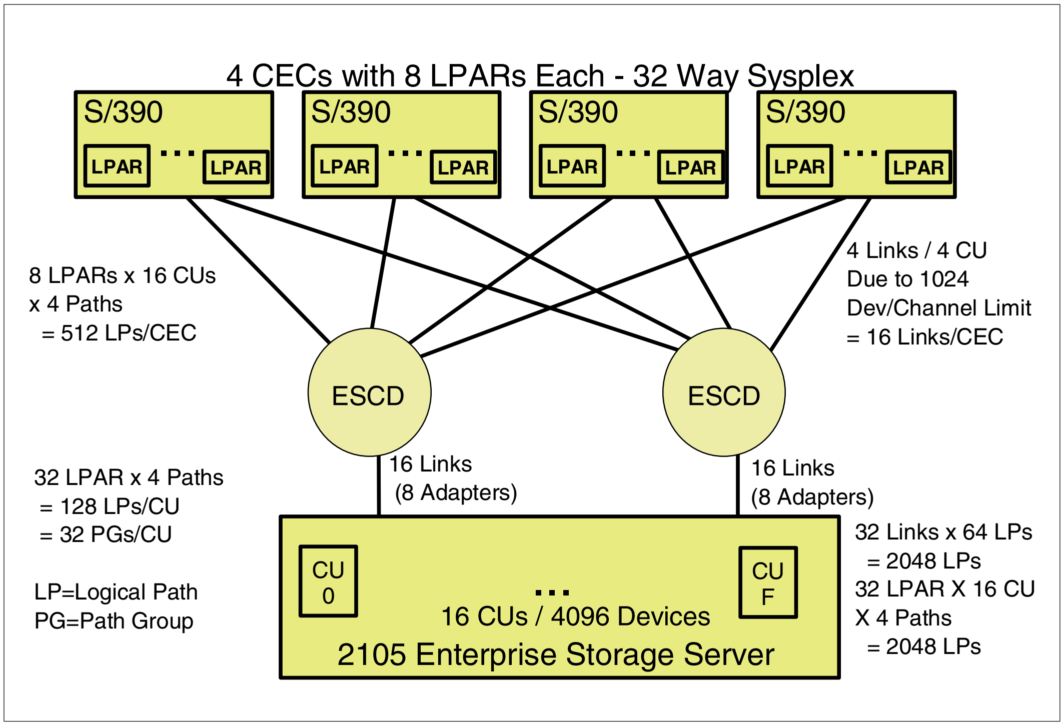

Figure 2-2 S/390 logical paths using ESCON

SCSI host view

Figure 2-3 shows an example of the logical SCSI bus view of the ESS.

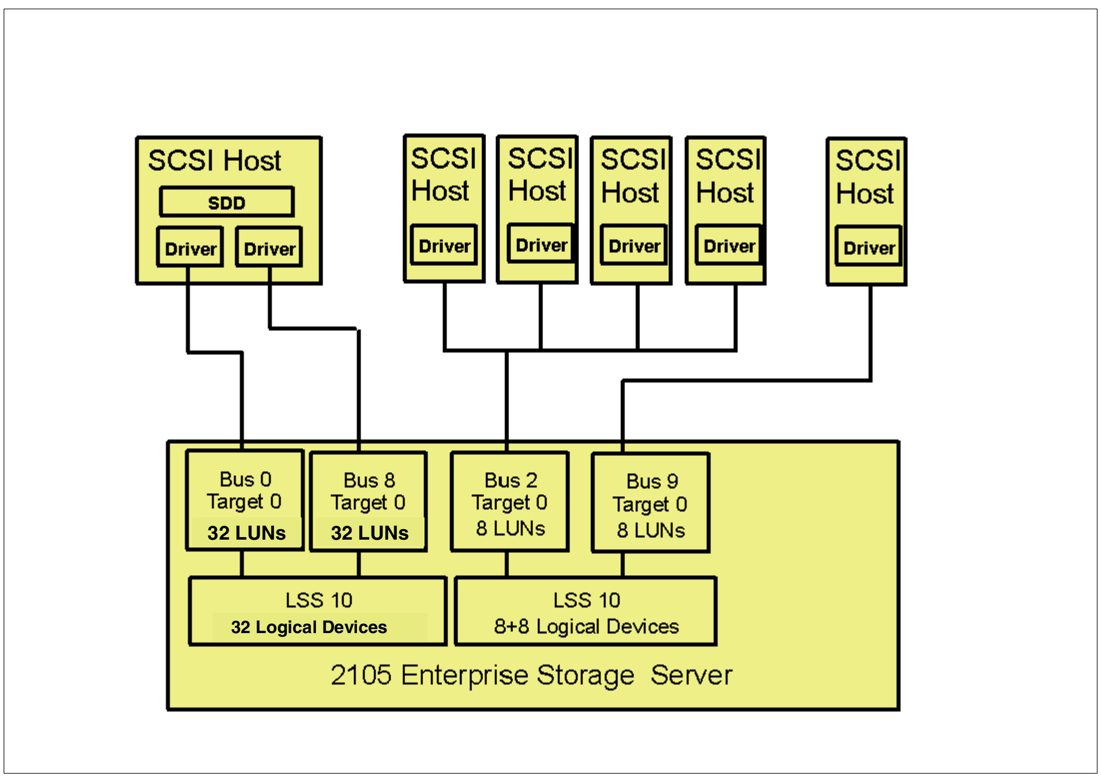

Figure 2-3 SCSI logical configuration

Figure 2-4 shows an example of the Fibre Channel logical configuration of the ESS.

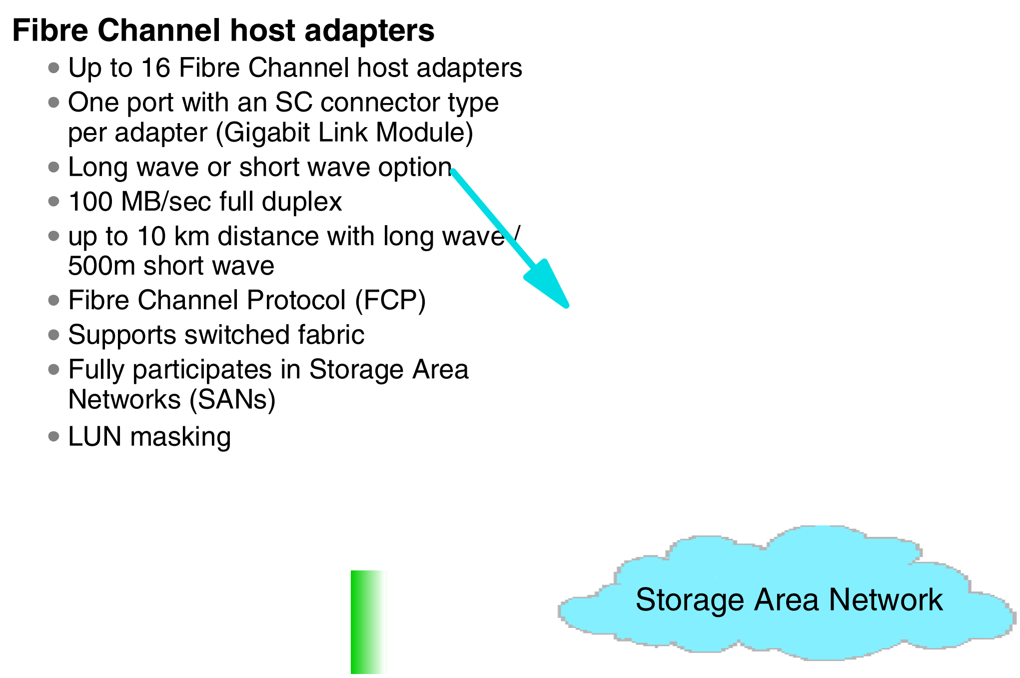

Figure 2-4 Fibre logical configuration

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 9.1 GB, 18.2 GB, 36.4 or 72.8 GB, due to the following:

- 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 and/or Fibre Channel FB devices or S/390 CKD devices

ESS model capacity

Table 2-1 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 2-1 ESS RAID-5 disk storage capacity by model (rounded)

ESS	9.1 GB	18.2 GB	36.4 GB	72.8 GB
Minimum F10/F20	420 GB	420 GB	841 GB	1683 GB
Maximum F10	420 GB	841 GB	1683 GB	3367 GB
Maximum F20 Base	841 GB	1683 GB	3367 GB	6734 GB
Maximum F20 w/Expansion	2804 GB	5610 GB	11224 GB	22449 GB

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 and/or Fibre Channel 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 in planning, the RAID array capacity data is shown separately for ESS arrays 1 to 8 for the Model F10 and 1 to 16 (base frame), and 17 to 48 (expansion frame) for the Model F20. This is due to the fact that the first two RAID-5 arrays on a loop contain a spare disk (the exception being JBODs). Therefore all the RAID-5 arrays in the base ESS will contain a spare disk, 6+P+S. The RAID-5 arrays in the expansion frame will all be 7+P arrays if the configuration consists of similar capacity disk drives only. But if you choose to intermix drive capacities on the same loop, because the first two RAID-5 arrays for each capacity must be 6+P+S, then the expansion frame can also start to become populated with 6+P+S arrays.

For your existing ESSs, when planning for the capacity requirements, you may want to consider the possibility of upgrading the installed eight-packs to the higher capacity disk drives of 36.4 GB or 72.8 GB.

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

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 36.4GB disks the maximum available capacity for:

- JBOD is 280 GB.
- RAID-5 is 245 GB (if a 7+P disk group).

The effective data capacity of a 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.

JBOD 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 72 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. For JBOD array capacities based on disk capacity, see Table 2-2.

Note: You must remember that JBOD ranks are not fault tolerant, so they should be used in conjunction with operating system mirroring techniques to provide fault tolerance.

Table 2-2 JBOD array capacity

Disk drive capacity	JBOD array capacity (8 + no parity or spares)
9.1 GB	69 GB
18.2 GB	140 GB
36.4 GB	280 GB
72.8 GB	561 GB

SCSI or Fibre Channel fixed block data capacity

A SCSI FB RAID-5 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 2-3. UNIX, AIX or Windows NT, can allocate LVs from 0.5 to 491 GB. Multiple LVs can be placed in an array/rank.

OS/400 supports 9337 devices with SCSI and 2105 devices for Fibre Channel Logical Volume sizes of 4.19 GB, 8.59 GB, 17.54 GB, 35.16 GB, 36 GB and 70.56 GB. Multiples of these can be placed into an array. OS/400 volumes can be intermixed on an ESS array with other open systems LUNs.

Table 2-3 RAID-5 array capacity - FB

Disk Drive Capacity	Arrays 1 - 16 (F20) Arrays 1 - 8 (F10) (6 + P + S)	Arrays 17 - 48 (F20) (7 + P) (see Note)
9.1 GB	52 GB	61GB
18.2 GB	105 GB	122 GB
36.4 GB	210 GB	245 GB
72.8 GB	420 GB	491 GB
Note: Arrays 17 to 48 will all be 7+P only if there is no disk drive capacity intermix. If the configuration consists of intermixed capacity disk drives, then arrays 17 to 48, some of them, will be 6+P+S. See explanation in 2.2.3, “Capacity planning by storage format” on page 19.		

S/390 CKD data capacity

S/390 effective data capacity will differ from the FB capacity in a rank due to the 3390 standard device formatting required. Also CKD has additional overheads, for example the count and key data (CKD), the home address (HA), and the record zero (R0) fields.

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

Interleaved: Interleaved arrays have two partitions striped across the array. The first partition is in interleaved mode and accounts for the major portion of the array capacity and is used to place 3390 or 3380 standard logical volumes (not custom volumes, not large 32760 cylinder volumes). LVs are placed within the partition in groups of four until there isn't enough storage space in the array to create another group of four still keeping 5,000 cylinders left over for the non-interleaved partition.

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

Interleaved applies to arrays with disk capacities of 9.1 GB, 18.2 GB and 36.4 GB.

Non-interleaved: A non-interleaved mode array 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 (3390-2 1.89 GB, 3390-3 2.83 GB, 3390-9 8.49 GB) or the size of the large capacity volume, 32760 cylinders. A custom volume can be from 1 to 10017 cylinders in size, or from 1 to 32760 cylinders in size for a system with Large Volume Support.

S/390 capacity tables can be found in "CKD logical volumes" on page 106 which show the various disk capacities and the number of logical volumes that can be defined.

2.2.4 ESS Copy Services

When using the ESS Copy Services functions offered by the ESS, the following considerations for capacity, pathing and ESS resources should be factored into the planning:

- FlashCopy provides a time zero (T0) image copy within the same LSS or LCU. Consideration needs to be made for the number of logical volumes that you may wish to use as target logical volumes since the disk capacity must be within the same LSS or LCU.
- PPRC will require additional disk storage reserved in the second (remote) ESS which will host the target volumes.
- XRC for S/390 will also require additional disk storage in the second (remote) ESS for the target volumes. You can use ESCON or FICON channels to connect the sites if the distance allows it, else you have to use channel extenders and telecommunication lines. One of the key elements in the sizing of an XRC configuration is the required bandwidth between the SDM (System Data Mover) and the primary ESS.
- PPRC requires that a primary and secondary ESS be connected by ESCON connections, for both S/390 and the open systems environments. 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. Remember to consider these points during planning, since they may effect the number and type of host adapters required.
- When using ESS Copy Services, one ESS cluster is the ESS Copy Services primary server which is the central place for collection of all information. A second ESS cluster can be nominated as the backup server. You will need to decide which ESS cluster to nominate for each role to ensure maximum availability. The IBM SSR will need this information during the ESS installation.

For detailed information on ESS Copy Services see *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757 and *Implementing ESS Copy Services on S/390,* SG24-5680.

INRANGE

Previous PPRC / INRANGE 9801 support is enhanced with greatly increased PPRC ESTABLISH or RESYNC performance at long distances. INRANGE supports new uses of ESS PPRC for data migration, data copy, or incremental resync at long distances.

- New ESS PPRC usages: open systems and/or S/390 data migration, data copy at long distances
- Utilizes new level of microcode in existing INRANGE 9801 Storage Network Subsystem hardware, combined with latest ESS microcode. See Figure 2-5.
- INRANGE supports OC3/T1/T3/E1/E3/ATM/ IP channel extension of PPRC
- ESTABLISH/RESYNC performance improvement: up to 40% reduction in elapsed time at short distances, up to 90% reduction in elapsed time at long distances

> **Note:** This enhancement does not change distance or performance for continuous synchronous PPRC.

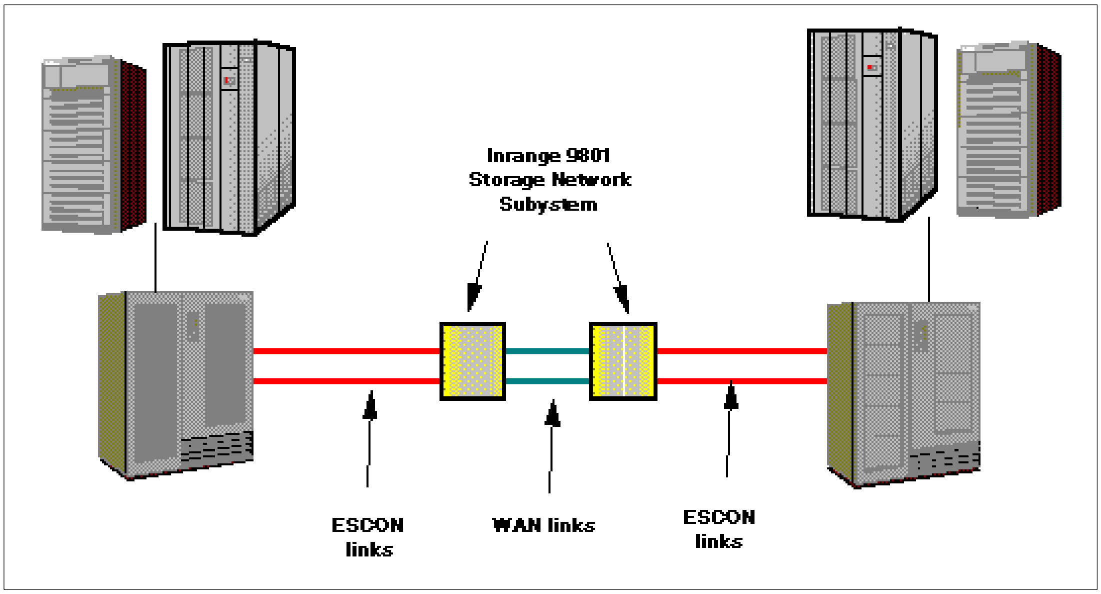

Figure 2-5 INRANGE 9801 storage

INRANGE can be used for long distance disaster recovery in two ways:

- Send log files long distance using PPRC as part of the log file archive process, with excellent performance.
- If the primary site can tolerate application quiesce as part of usage, then you can run PPRC pairs in a RESYNC/SUSPEND loop operation.

Do a RESYNC at appropriate intervals to create a backup copy.

This will require application-sponsored checkpoints or quiesces to insure production primary volume performance objectives are met.

2.2.5 Flexible configurations

The ESS has the ability to support mixed drive sizes (intermix). With drive intermix, a vast number of configuration options are available, so price and performance can be optimized to meet specific server and application requirements within a single ESS.

Intermix support is available on all ESS models. Any type and quantity of available disk 8-pack features can be intermixed as long as 8-packs of a given DDM capacity are purchased and added to a loop in pairs and supported on a given ESS model.

Intermix support requires the new Flexible Capacity Option features (#9500 and #9600).These new features replace the Standard Configuration features previously used when specifying ESS capacity and drive type. With the Flexible Capacity Option, capacity is now ordered by explicitly selecting the type and quantity of disk 8-pack features desired.

2.3 Logical configuration planning

The logical configuration of the ESS involves defining how the ESS is seen from the attached hosts. The UNIX, AIX, NT system hosts will see the ESS as SCSI and FCP generic devices, the AS/400 will see it as a 9337 or 2105 external disk, and the S/390 host will see it as a 2105 control unit with one or more logical control units (LCUs) attaching IBM 3390 devices.

The basic steps required are as follows:

- Draw up a logical map or 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 and 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 using the ESS Specialist during installation.
- Give the logical specification to 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 SSR input these requirements at installation time. For more information on the initial configuration process, please refer to Appendix A, “ESS configuration planning process” on page 293.

2.3.1 Logical subsystem mapping

Determine the number of SCSI and/or FCP 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 and FCP devices include the following:

- How will the hosts be connected to the ESS?
- Determine RAID or non-RAID format.
- Determine the LV or device size.
- Define the number of targets and the number of LUNs per target.
- Will LUNs be shared on two or more hosts?
- Which will be the host attachment: Fibre Channel, SCSI?
- Addressing requirements: Are specific SCSI IDs required for host and target? It is recommended to set the initiator ID to 7.
- Consider that iSeries and AS/400 6501 SCSI IOP can only access 1 to 8 logical disks.

A S/390 LSS relates directly to a S/390 logical control unit (LCU) with its associated devices or volumes. A four digit subsystem identifier SSID will need to be assigned to each logical control unit. Also when planning for S/390 CKD LSS mapping, PAV devices must be taken into account towards the 256 LSS device limit and also the 1024 devices limit per ESCON CHPID (16384 for FICON).

LSS/LCU to physical mapping

This is how the LSS physical mapping is implemented.

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 2-6, the CUADDR for cluster 1, adapter 1, loop B is 08.

SCSI and Fiber Channel: The mapping of the SCSI and Fibre Channel FB LSS to target/LUN is also dependent on the DA and loop association, and is automatically assigned during ESS Specialist setup. The SCSI and Fibre Channel FB LSS numbers are shown in Figure 2-6 below.

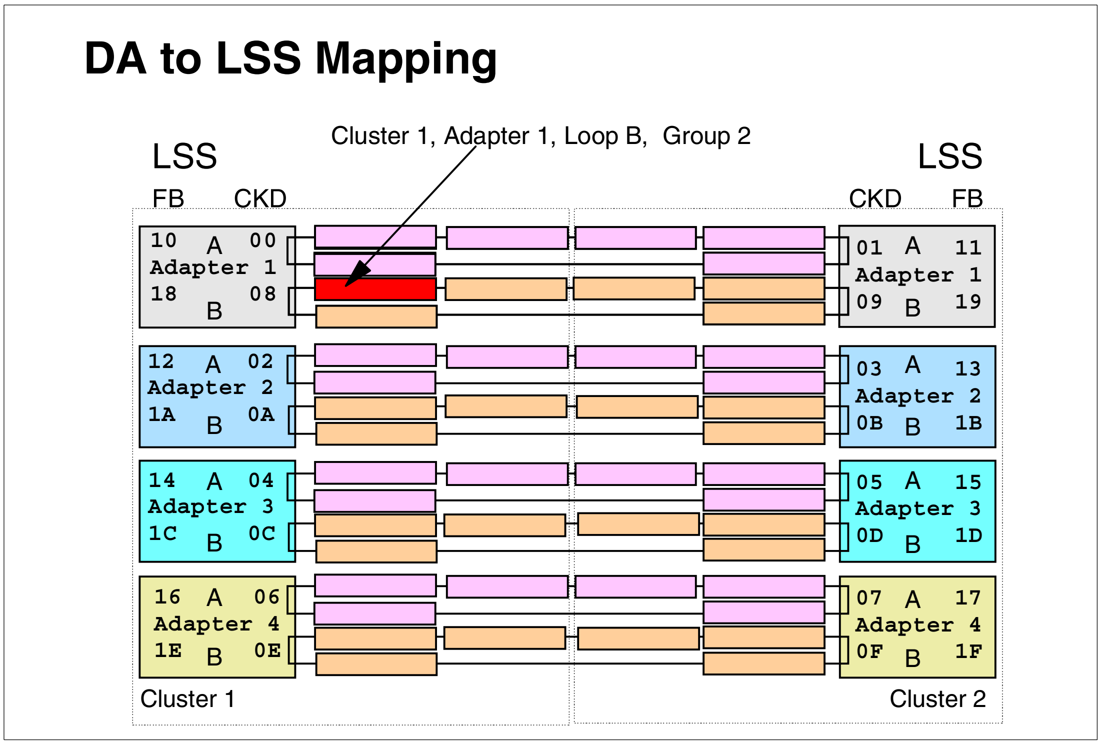

Figure 2-6 LSS logical to physical mapping

Disk group allocation

The disk 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 LSSs on each adapter. For example, LCU 00 or 08 will have access to disk groups 2, 4, and 8on the loop, with LCU 01 or 09 accessing disk groups 1, 3and 5. The ESS Specialist manages this for you.

Logical device mapping

This is how the logical device mapping is implemented.

Array allocation: The first step has to do with the definition of which arrays/ranks are to be assigned/mapped to which LSS/LCU.

For S/390 you must associate at least one disk group with the first CKD LSS for the device adapter. If there is at least one disk group associated with the first CKD LSS for the device adapter, other CKD disk groups can be associated with either the first or the second CKD LSS for that device adapter.

For open systems when a disk group is formatted as an FB array, if there are already 193 or more LUNs defined in the first LSS for a device adapter. the new disk group will be added to the second LSS for the device adapter.

RAID or JBOD: 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-5.

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

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

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

The ESS maximum number of SCSI LUNs (or LVs) per target is 64 and 16,000 for Fibre Channel, however, it is host operating system or I/O adapter dependent. For example, many servers operating system only support 8 LUNs per target for SCSI-attached storage and 32 LUNs per target for Fibre Channel attached storage. Similarly with Fibre Channel the non-Report LUNs host support 256 LUNs per target, while the Report LUNs hosts support 4096 LUNs per target.

S/390: Decide whether you will be using Interleaved or non-interleaved volumes. The CKD logical volumes defined in the interleaved partition, will match full 3390-2 (1.89GB) or 3390-3 (2.89GB) or 3390-9 (8.51GB). For the remaining non-interleaved partition, which will be the custom volumes you will be defining?.

A decision must be made on the logical device type to emulate, 3390 Model 2, Model 3, Model 9, or large capacity volumes, or 3380 track format. If using custom volumes, decide 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 option enabled via the ESS Specialist (*Note*: the PPRC feature has to be previously activated by the SSR during the ESS installation or later if this feature is added on an existing ESS). The volumes that require PAV aliases need to be identified, along with the number of aliases to be available to each volume.

Host connections

The number of connections between a server and the ESS must be decided based on the bandwidth requirements of the server. For zSeries servers, this will be the number of FICON or ESCON channels in channel path groups attached to the ESS. For open systems servers, it will be the number of SCSI busses or Fibre Channel I/O adapter-to-ESS port connections.

The ESS SCSI logical volumes are specifically assigned to an HA port via the ESS Specialist at ESS setup. If using the IBM SDD, the physical SCSI and/ or Fibre Channel connections should be spread across different HA bays.

For Fibre Channel open systems servers, LUNs are assigned to a sever host bus adapter by its world-wide port name (WWPN). A server host bus adapter can be restricted using the ESS Specialist to accessing its LUNs through selected ESS Fibre Channel host adapter ports.

When planning how much capacity to attach to a SCSI or Fibre Channel 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
- We recommend that no more than 400 GB are assigned per Fibre Channel port.

Note: The previous attachment guidelines are based on assumptions about the average sequential bandwidth requirements per GB of data stored.

For S/390 this is fairly straight forward since it is no different than standard device considerations. It is recommended that eight ESCON channels for each z/OS or S/390 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. For ESCON configurations be sure to review the logical path limitations. With FICON things become simpler because many ESCON limits no long exist. Less FICON channels will be used when compared to ESCON: a 4:1 consolidation can typically be achieved. Also with FICON, thanks to its greater bandwidth and increased logical pathing characteristics, more LCUs in the ESS can be logically daisy chained thus addressing more devices with less channels and consequently less host adapters needed.

PPRC and XRC also have an additional requirement for ESCON host adapter ports. PPRC primary or source volume ESS 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. For XRC (S/390 users) dedicated host adapter ports are recommended. Also in the sizing an XRC configuration the required bandwidth between the SDM (System Data Mover) and the primary ESS must be determined, having in mind which are the peak demands.

Planning 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? Also keep in mind that arrays are installed in pairs and must be installed in fixed order.
- 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-5 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. The topics in this section should be complemented with the corresponding information presented in the following chapters.

Disk drive capacity

There are configuration decisions to be made when ordering your ESS. One is the disk drive capacity. You have the option to choose 9.1 GB, 18.2 GB, 36.4 GB or 72.8 GB disks drives.

The greater capacity 72.8 GB drives provide superior seek and throughput characteristics, thereby allowing each drive to satisfy big demanding I/O loads. Moreover, access densities (I/O operations per second per GB disk space) continue to decline, so workloads migrated to 72.8 GB drives will often not notice any material difference in performance compared to smaller capacity disk drives.

Some customers have very high demand workloads that may be very cache-unfriendly or have a very high random write content. These workloads may still require the lower capacity drives (more drives are required for a given capacity). For these workloads, customers may want to consider purchasing lower capacity drives or, for environments with mixed workloads with different characteristics, an intermix of drive capacities.

Host adapters

When physically connecting a host system to the ESS host adapter ports (ESCON, FICON, Fibre Channel 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 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 at least 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 4-9 on page 76 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 (one host adapter per host adapter bay), the other associated with cluster 2, the second set of four device adapters (one host adapter per host adapter bay).

One eight-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 group should provide access to the LCUs defined on the other Cluster 2, CUADDR 1, 3, 5, 7, 9, B, D, F. See Figure 2-6 on page 24 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.

For FICON the approach to a good performance configuration is similar but the requirements differ because of FICON better performance characteristics and improved logical pathing limits. For installations where ESCON host adapters were in the 25-50% utilization, then one FICON host adapter is enough where four ESCON host adapters were used. Plan at least four FICON host adapters per ESS. A more typical configuration would have eight FICON HAs per ESS (eight FICON HAs may be necessary to exploit ESS full bandwidth for some workloads).

Spread FICON HAs across all four HA bays. This should result in minimally one FICON HA per bay, or in a typically configured ESS, two FICON HAs per bay.

For each control unit definition in HCD/IOCP, define a minimum of four FICON channels per ESS/LSS path group (or typically eight).

Subsystem Device Driver: The Subsystem Device Driver (SDD) for open systems can provide multiple SCSI or Fibre Channel bus paths to the same target/LUN group, thus spreading the I/O workload. This I/O load balancing improves the performance of the server I/O operations.

Cache

If you are only defining one type of LSS/LCU (CKD or FB) 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 SSR during installation. You must have in mind that if you later need to change the LSS definitions (if a requirement of mixing workloads appears in your installation) then you will need to re-IML both ESS clusters together.

Cache options for the ESS are 8 GB, 16 GB, 24 GB and 32 GB. Consideration for the application use of the ESS cache should be factored in when ordering the cache size. In general larger cache sizes can be very beneficial for S/390 workloads, but are much less likely to help open systems where the I/O caching is resolved in the host application buffers.

LSS/LCU

For a given host it could be beneficial to use as many LSS/LCUs as possible, from the configured LSS/LCUs in the ESS (Note: for S/390 ESCON path requirements it may be that 8 rather than 16 LSSs is the recommended configuration). This will enable to spread the logical volumes of the server across a greater number of LSS/LCUs, and therefore DAs and loops, thus getting more of the ESS back-end bandwidth.

Spreading the devices of a given server across more LSS/LCUs (using more DAs and loops) will also make the I/O to be directed across the two ESS clusters. This allows to take advantage of the bandwidth and processing capacity of both clusters.

If the host is dedicated to a specific application, this dispersion of the I/O could be beneficial and more easily be implemented and monitored. But when more than one application is running simultaneously then you may have overlapping I/So from the different applications/servers and then the benefits of spreading the logical volumes of a given server may not be so evident. Also you must consider that using more LSS/LCUs for the logical volumes of a given server will have to be considered when planning your application implementation of FlashCopy procedures (FlashCopy works with source and target within the same LSS/LCU).

Data isolation

For performance reasons, contrary to the spreading, you may wish to isolate volumes containing data with a particular activity pattern or attribute. You can use the ESS design to your advantage by isolating or grouping together volumes at either the Cluster, target LSS/LCU (DA), or array/rank level by selecting the disk groups where you allocate the logical volumes.

RAID-5

RAID-5 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 interleaved 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.

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. Custom volumes can also assist you in circumventing hardware serialization problems at the volume level, for example reserve/release.

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

Parallel Access Volumes (PAV): S/390

Parallel Access Volumes provide the ability to almost eliminate IOSQ time, a major contributor to response time. PAVs will improve the performance of volumes are experiencing large IOSQ values. In general 1 PAV for each 3390-3 or 3 PAVs for each 3390-9 is recommended.

Larger volumes should use more aliases to keep IOS queueing down, but also allow more aliases to be defined for the LCU since there will be less base defined devices. Please refer to Section 8.1.6, “Parallel Access Volumes” on page 110 for detailed considerations.

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 PAV and enabling I/O Priority Queueing.

3

Physical installation planning

This chapter discusses various physical considerations involved in planning the physical installation of a new ESS in your environment. Some of the topics covered are host attachments, electrical and cooling, floor loading, network, and communication requirements. You should consult the latest version of this guide for the current information as you plan for the physical installation of your ESS.

3.1 General considerations

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

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

Always look at the most recent information for physical planning in *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294.

3.2 Host attachments

If you are configuring open systems SCSI/Fibre Channel and S/390 ESCON/FICON logical disks in your ESS, you need to locate the ESS appropriately, within the various cabling schemes found in the information you have gathered from the references in the installation planning section. Please refer to Chapter 1 of the *IBM TotalStorage Enterprise Storage Server Installation and Planning,* GC26-7294.

3.2.1 Daisy chaining SCSI adapters

We do not recommend daisy chaining SCSI host adapters together. Although it is technically possible, it has proven to be problematic.

3.2.2 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 IBM TotalStorage Enterprise Storage Server Specialist (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 “Configure Host Adapter Ports” on page 247 to see the process involved to do this.

3.2.3 Fibre Channel host attachment

The ESS supports attachment to a Fibre Channel arbitrated loop and Direct native Fibre Channel connection or through a switch fabric.

Direct Fibre Channel connection or through 2109 Fabric switches

To complement the information in this section, please refer to *Implementing Fibre Channel Attachment on the ESS*, SG24-6113.

3.3 S/390 ESCON/FICON cables

This section discusses various considerations regarding S/390 ESCON/FICON cables.

ESCON cable considerations

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 with which they already have experience.

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.

FICON cable considerations

The ESS Models F10 and F20 support native FICON connection over short-wave or long-wave fibre links to provide FICON attachment to zSeries or S/390 servers. For the FICON cable options refer to the chapter 12 in the IBM Redbook, *FICON Native Implementation and Reference Guide,* GC24-6266.

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 detailed information on PPRC for the ESS, refer to Chapter 12, “ESS Copy Services for S/390” on page 197, and the corresponding chapter in *IBM TotalStorage Enterprise Storage Server,* SG24-5465.

TCP/IP Connection for PPRC: If the ESS 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 100BaseT 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

All ESS models, 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 *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, GC26-7294 for details on power consumption, voltage, and current requirements for the ESS.

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 3-1. A wall outlet typically can supply only 20 or 30 amps at 120 volts.

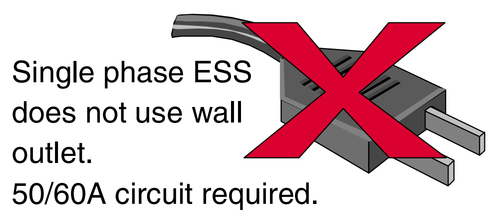

Figure 3-1 Single phase ESS

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

This section discusses considerations regarding power connections and remote control.

Power connectors

The cable connectors supplied with various line cords, and the required receptacles are given in Chapter 3, "Installation Planning" in the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, GC26-7294. The length of the line cords is 14 feet, except in Chicago, where the length is 6 feet.

Convenience outlet power

The ESS 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.

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.

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.

To check for specification changes and updates on the Web, go to:

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

Refer to Table 3-1 for dimensions of the Arbo crate.

Table 3-1 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 F10 and F20	Height: 75.25 (179.6) Width: 54.5 (138.3) Depth: 35.8 (91)	2160 (980kg)	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)	2590 (1175kg)	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 F10 or F20 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 *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide*, 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). 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 3-2.

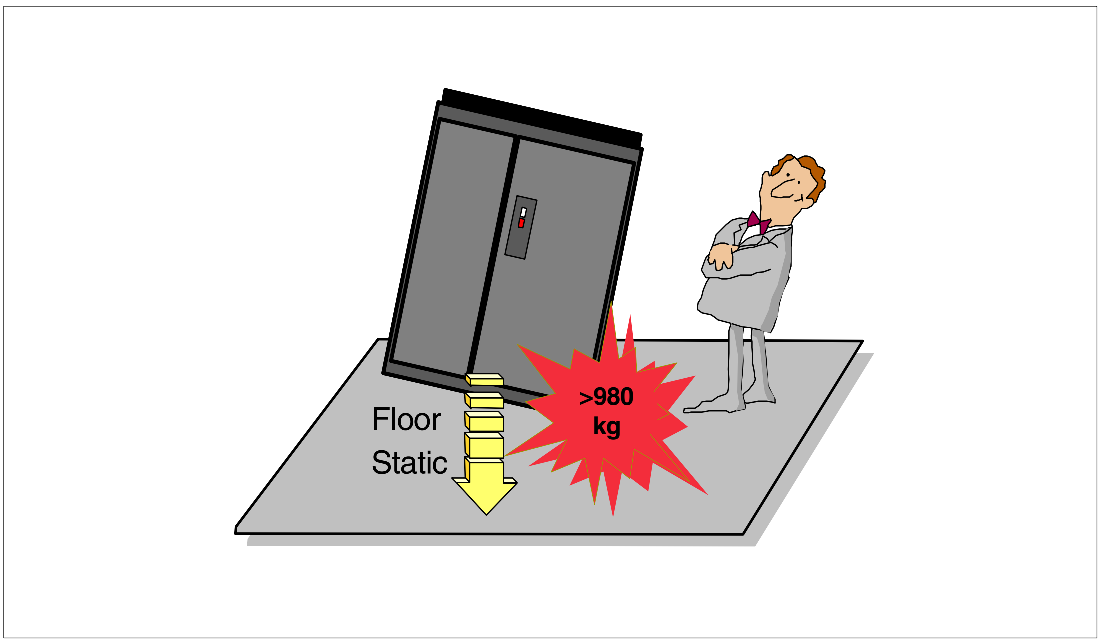

Figure 3-2 Caution: Floor loading

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

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 *IBM TotalStorage Introduction and Planning Guide, 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 ESS should be maintained within an operating temperature range of 20 to 25 degrees Celsius (68 to 77 degrees Fahrenheit). The optimum temperature is 22 degrees Celsius (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 IBM TotalStorage Enterprise Storage Server *Introduction and Planning Guide, G*C26-7294, for information on required temperature and humidity, and heat output of the 2105 models.

3.7 Call Home and Remote Support

Why is the “call home” feature needed?

- To allow the ESS to automatically notify IBM when service is required
- To allow remote access and service by IBM support specialists

A remote link capability enables call home support where the IBM TotalStorage Enterprise Storage Server can initiate a service call if it detects or anticipates a problem occurring. The remote link also enables service calls into the ESS for remote analysis and potential correction of problems.

Service alerts describe a problem and the action and/or parts needed to correct the problem. When a service alert is generated, an e-mail and/or SNMP can be issued to designated users, and a notification can be sent via the ESSs call home function to an IBM Service Call Screening Center. Upon notification, service personnel may be dispatched to provide repair service or remote service actions may be initiated. For the call home function to work, there must be an analog phone line available for the remote service support.

The Call Home and Remote Service Support functions have been enhanced with the introduction of the IBM TotalStorage Enterprise Storage Server Master Console.

3.7.1 IBM TotalStorage Enterprise Storage Server Master Console

The Remote Support Facility feature code (FC 2715), which included the ESSNet console, has been superseded by ESS Master Console (FC 2717) for new ESS orders.

The IBM TotalStorage Enterprise Storage Server Master Console (ESS Master Console) is provided with the first ESS installed in a site. This is a desktop machine preconfigured with Linux and a Netscape browser. It provides a Web client function and allows connection into the customer's LAN. The ESS Master Console is the means for:

- The IBM SSR to administer service functions. What previously was done using his Most (Mobile Solution Terminal) laptop, now can be done from the Master Console.
- The IBM SSR to set up the Call Home, Remote Support Facility and other ESS service related options.
- Launch the ESS Specialist. The ESS Master Console includes links to the ESS client user interface. When you click on one of these links, it initiates the Web interface to ESS Specialist and ESS Copy Services.

When ordering the first ESS, FC 2717 (ESS Master Console) must be included (Figure 3-3). This feature provides the hardware to support the ESS configuration, call home, and remote support capabilities by providing a dedicated console (processor, modem, monitor, keyboard) and networking components (hub and Ethernet cables and multiport serial adapter). Up to seven (7) ESS machines are supported per ESS Master Console. Six additional ESS machines can be connected into the ESS Master Console if FC 2716 (Remote Support Cables) is ordered with each additional ESS to be connected.

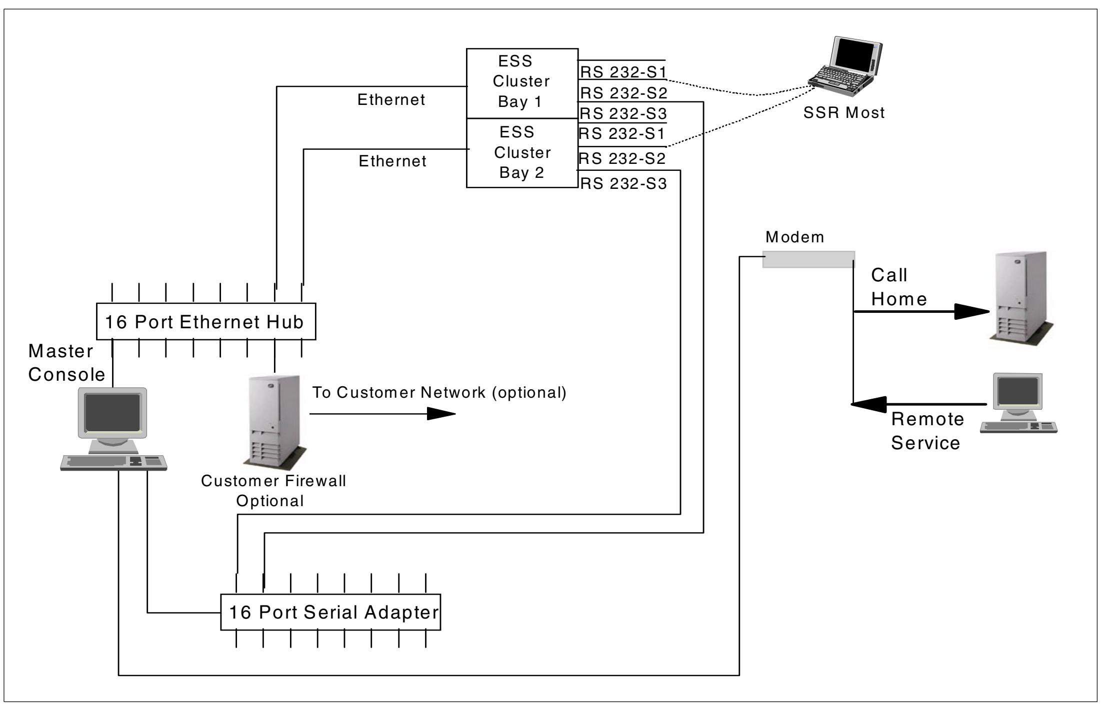

Figure 3-3 ESS Master Console: (FC 2717) — connections

Figure 3-3 shows the ESS Master Console connections and remote support functions. The same Master Console can be used to launch the ESS Specialist or the ESS Copy Services interfaces by clicking the **ESS Launcher** icon (see Figure 3-4). Optionally, you can access ESS Specialist from your workstation if you connect the ESS Master Console into your intranet by using one of the ports on the ESSNet hub as shown in Figure 3-3.

The desktop machine will display the panel as shown in Figure 3-4 when the ESS Master Console is installed or when the ESS Master Console icon is pressed from the desktop.

Note: The panels from the ESS Master console used in this section are not intended to be all inclusive. We provide these screens for you so that you recognize them as a component of the ESS Master Console.

Your IBM SSR will be trained in the use of all the console panels and is the intended user of the ESS Master Console.

Figure 3-4 ESS Master Console: Main screen

In order to allow remote server support, **Console Launcher** button must be clicked. Then the IBM SSR will proceed to click the **Remote Services Settings** button as shown in Figure 3-5.

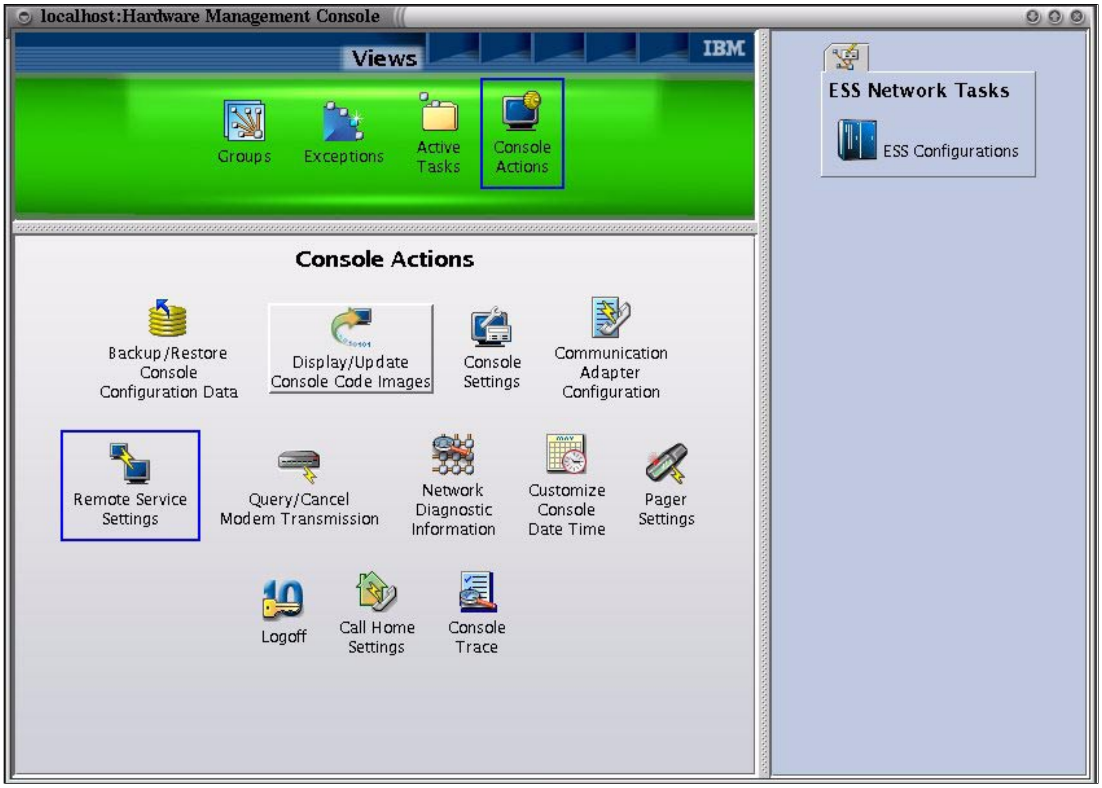

Figure 3-5 ESS Master Console: Selecting remote server services

You will then be given the panel shown in Figure 3-6 to enable remote services. Here you must select the use of the default password or supply a custom password.

Figure 3-6 ESS Master Console: Enabling remote services

3.7.2 ESS Master Console differences

The TotalStorage Enterprise Storage Server Master Console (ESS Master Console) improves on the previous ESSNet concept. While providing the same overall function, it also adds new capabilities. ESS Master Console does not use the ESS Ethernet connections and replaces them with serial port connection to each ESS cluster creating a serial network (see Figure 3-4 on page 40).

It permits the elimination of the modem expander that was used with ESSNet (see Figure 3-7). In its place is a Multiport Serial Adapter (MSA) which provides performance improvements giving each ESS the ability to transfer control information independent of other ESSs that might be connected to the ESS Master Console (see Figure 3-4 on page 40).

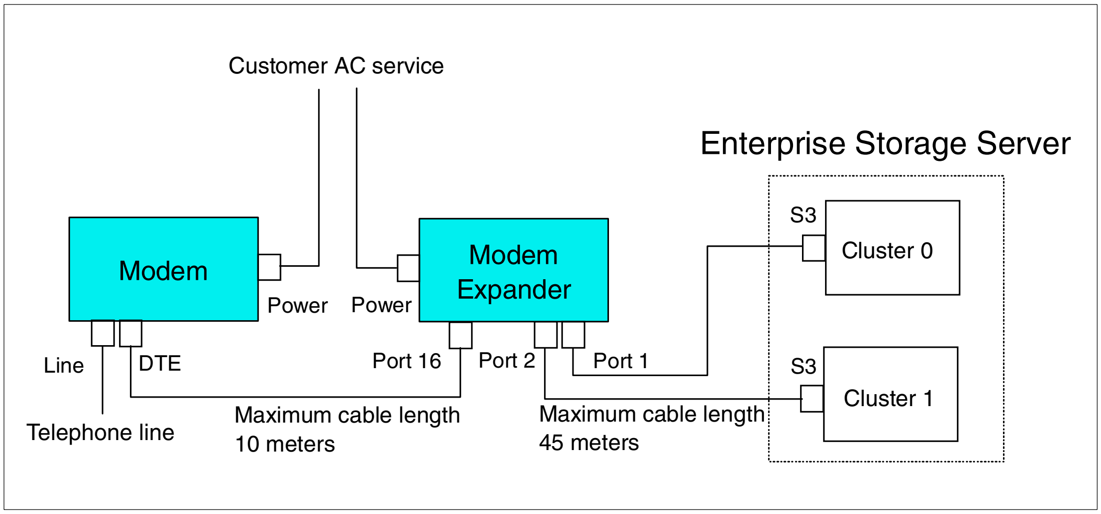

Figure 3-7 Call home configuration of the old ESSNet (FC 2715) with modem expander

Table 3-2 shows the differences between the ESS Master Console and the old ESSNet console.

Table 3-2 Differences between ESS Master Console and ESSNet console

Component	ESSNet console	ESS Master Console
Operating system	Microsoft Windows 4.0	Linux
Software	• Microsoft Internet Explorer (MSIE) browser • Netscape Browser (see note) • Virus protection software	• Netscape Browser (see note) • Terminal emulation • Virus protection software not applicable
Communication	• Modem • Modem expander	• Modem • 16-port serial adapter
Note: The default browser is Netscape 4.7.6. ESS Specialist does not support Netscape 6.0		

The ESS Master Console implementation offers, among others, the following enhancements to the remote support function over the previous ESSNet implementation:

- Activation of microcode engineering changes (ECs) from the ESS Master Console
- Reduction of long-distance telephone costs for call home service
- Improved data transmission rates and improved reliability for state saves and traces

3.8 Web client and Ethernet connection

The two clusters in the ESS communicate with each other through a 100baseT Ethernet connection. With the addition of a hub, this Ethernet is used also as a link to the Web client, and optionally can be connected to the user's LAN to provide connectivity from anywhere using TCP/IP.

3.8.1 ESS Specialist

The IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) is a software application that runs in the ESS clusters. It is the interface provided for the user to define and maintain the configuration of the ESS 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, "IBM TotalStorage Enterprise Storage Server Specialist" on page 79.

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 100BaseT 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.

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 4.7.6 running on Red Hat Linux 7.0 or NT. 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, Windows NT Version 4.0., or Red Hat Linux 7.1.
- 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 (***Note***: the security restrictions with the MSIE browser have been resolved with the later levels of MSIE and the G3+4 an later levels of ESS microcode).

For ESS Specialist Web browser interface prerequisites, refer also to 5.2, "ESS Specialist prerequisites" on page 81.

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 100BaseT cables, connected as shown in Figure 3-8.

The internal 10BaseT cable connecting the two clusters together is removed by the IBM service representative during ESS installation, and a 100BaseT 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 SSR 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.

When planning for the Web client connection remember that feature 2717 (IBM TotalStorage Enterprise Storage Server Master Console) consists of a dedicated console (processor, modem, monitor, keyboard) and networking components (hub, Ethernet cables and serial cables). Please refer to Section 3.7.1, "IBM TotalStorage Enterprise Storage Server Master Console" on page 39 for further description of these hardware components.

Multiple ESSs 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 ESSs in a site participate in PPRC, and this is to be managed using the ESS Copy Services function, they need to be connected to the same private LAN.

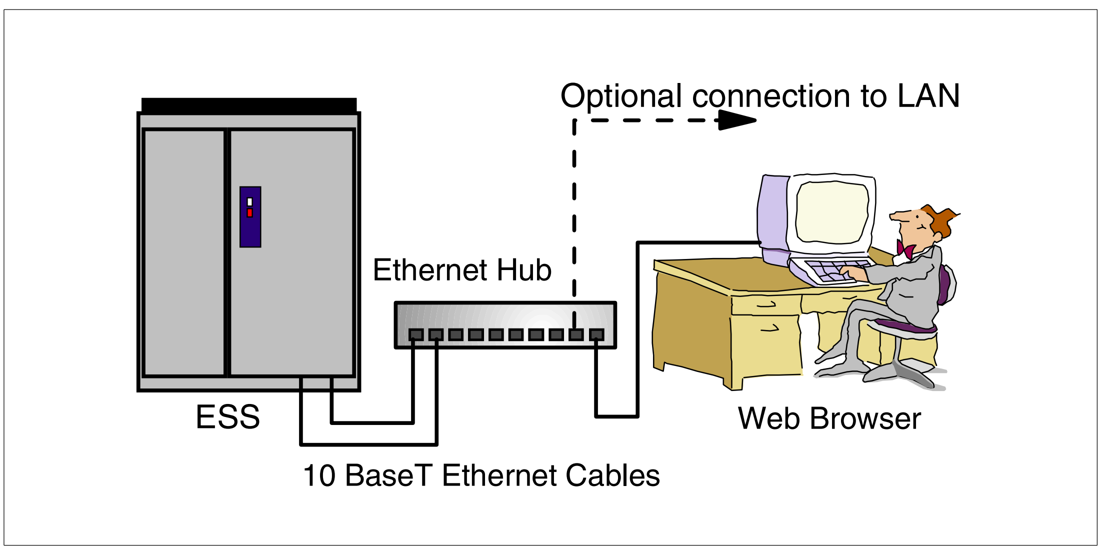

Figure 3-8 Ethernet private LAN configuration

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 3-8 for details. The user must supply any cables, bridges or hubs necessary to make the connection.

Why connect to the TCP/IP network?

These are some reasons why you may want to connect to the TCP/IP network:

- You may want to use the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist), or IBM StorWatch Enterprise Storage Server Expert (ESS 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 ESS Copy Services function of the ESS. This requires that the Web browser used for the ESS Specialist has TCP/IP connectivity to each local and remote ESS participating in PPRC.

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.
- Have the ESS Master Console and both ESS cluster TCP/IP addresses (hostnames) defined to the DNS server.

Use the Communication Resources Work Sheet at the back of the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294, to specify the TCP/IP address, network mask, default gateway address, hostname, and nameserver TCP/IP address for the ESS 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 four 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 ESS Copy Services behind the firewall

This section discusses ESS Copy Services behind the firewall.

Network Configurations and ESS Copy Services

To use ESS Copy Services in a network environment there are some questions you and your network administrator will have to answer. Each answer will depend on the configuration you are trying to implement.

For the purpose of narrowing the possible configurations, we have isolated the possible configurations to be within the following boundaries. As in any configuration, if you do not have the expertise available, IBM Global Services or your service provider could assist your administrators in setting up the right configuration for your environment.

Each ESS Server is located within a private Network along with the CLI Host or Client Browsers as shown in Figure 3-9.

The Primary and Alternate ESS Servers are separated by a Non-Secured or Public Network with the CLI Host and Browser located at either end.

The configuration of the sub-area network with Network Address Translation (NAT)/Firewalls in front of the ESS Server and placements of the CLI Host and Client Browsers.

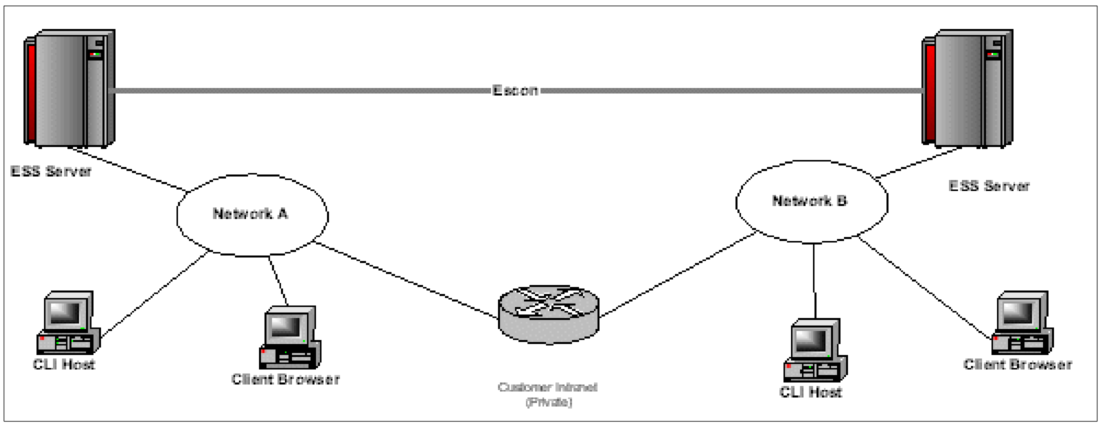

Figure 3-9 Network view 1

In Figure 3-10, the network connection attaching to the ESS Servers and CLI Hosts is connected within a customer intranet. This implies that each network is aware of each other and are routable to each other. The firewalls protecting the intranet are not between network A and network B. This is important when attaching the ESS server to your network. The required IP addresses have been obtained and are properly configured on the ESS server and CLI Host. There are no other required configuration changes to be performed to have the ESS Specialist console, CLI host and client browsers talk to the ESS servers. Note that the primary data transmissions and copies are done through the original connections.

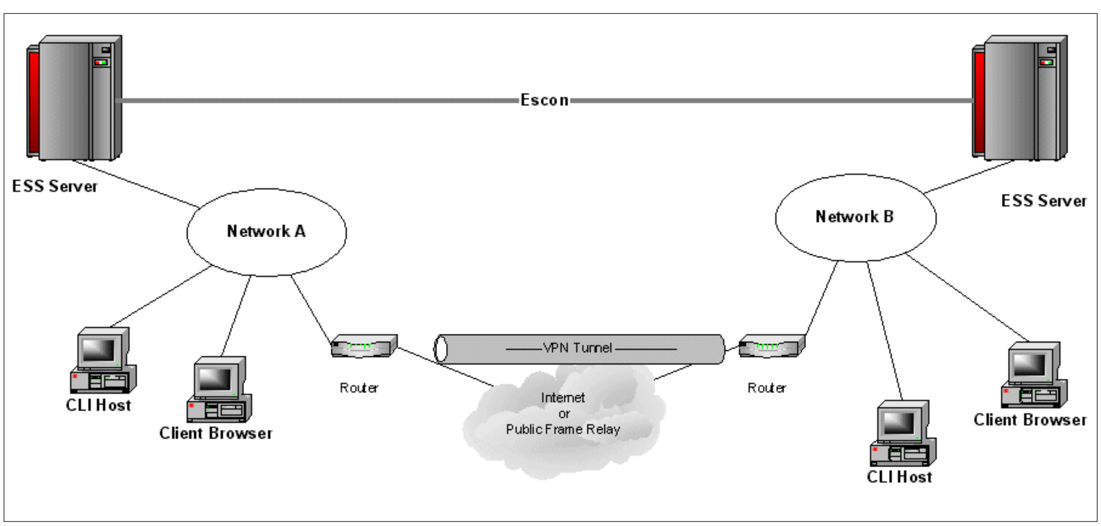

Figure 3-10 Network view 2

In Figure 3-11, a connection is defined to simulate the attachment of two offices located in locations remote to each other. The two networks are attached to the Internet and tunneling protocols have been implemented between the two sites to secure the management, configuration and processing data being passed between the CLI Host, Client Browser, and ESS Servers.

Again, because the two networks are unique to each other the same communication channels are used to pass the management, configuration and processing data as are used for communicating between the two remote sites. With the Virtual Private Network (VPN) Tunneling being established by the routers that are connecting the two sites, the data is being transmitted in a secured environment.

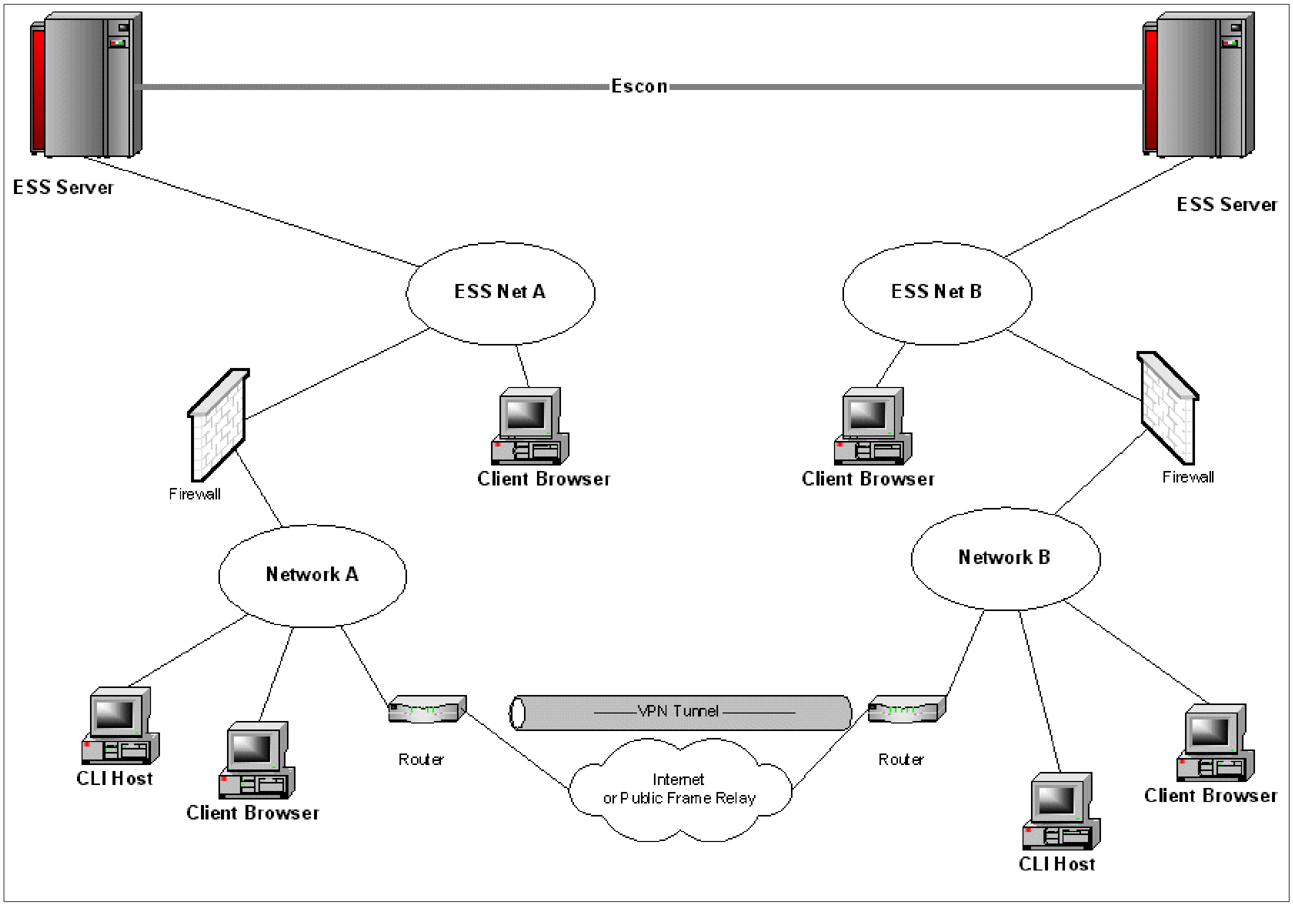

Figure 3-11 Network view 3

When connecting the CLI Host to an ESS Server using various security type procedures, such as Network Address Translation (NAT) and filtering routers, there are some procedures that should be followed to allow the passing of traffic from the CLI host to the ESS server. This may require verifying the configurations and port opens by your network security administrators.

Here is the list of filters which have been tested with filtering router firewalls that will have to be opened by the network owner of the router/firewall. Other techniques such as application firewalls and stateful firewalls could also be used. Follow the filters with: TCP any established, verifications, and/or stateful filter:

- TCP Port 80 - WWW (this means port 80 has to be opened to view your Web browser, over the World Wide Web)
- TCP Port 443 - HTTPS
- TCP Port 23 - Telnet
- UDP Port 53 - Domain (DNS)
- TCP Port 1703 - Nameserver - SOCKS
- TCP Port 1705 - appletserver - SOCKS
- TCP Port 1707 - applet status server - SOCKS
- TCP Port 1709 - client heartbeat - SOCKS
- TCP Port 1711 - cobra code - SOCKS/IIOP
- TCP Port > 32000 - Heartbeat opens between ESS servers to permit TCP any established

When using NAT you must be careful where you put your CLI host. If it is located as CLI host A, as shown in Figure 3-12, it will not be able to translate the Web updates correctly. This is because the translation is done outside the data frames being transmitted to the ESS Servers. The solution to having the CLI host located outside your firewall with NAT is to use VPN tunneling techniques.

This can be done in two ways. The first is to simply install the client VPN software supporting IPSEC Standard for your firewall/NAT device as in example A in Figure 3-13. The second is to install a device(s) which supports VPN Tunneling to connect CLI Host A to the ESS server as in example B in Figure 3-13. This would place the CLI Host in the same network. If you moved CLI Host A to ESS Net A, you would simply apply the filters required for the two ESS servers to allow them to talk to each other. CLI Host A now would be on the NAT side of the router/firewall, which would then translate correctly and find the other ESS servers.

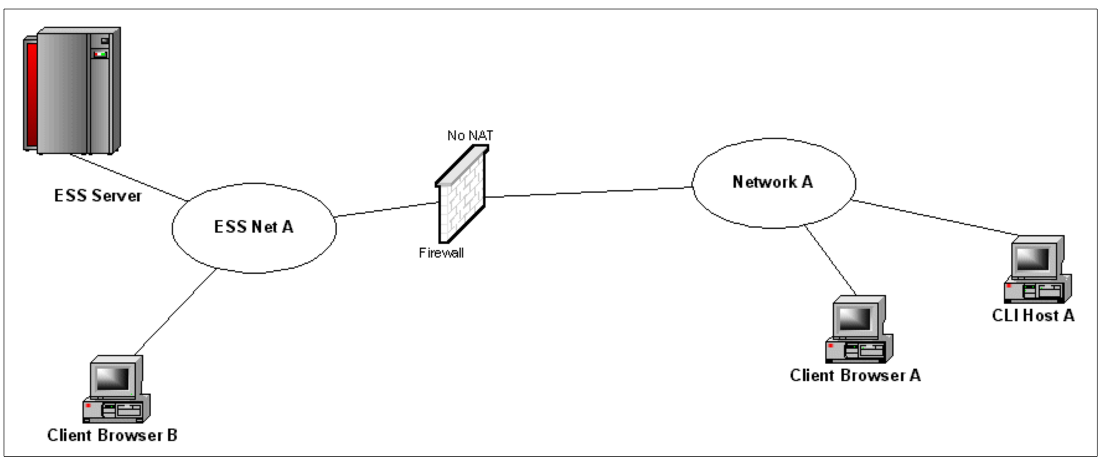

Figure 3-12 ESSNet A

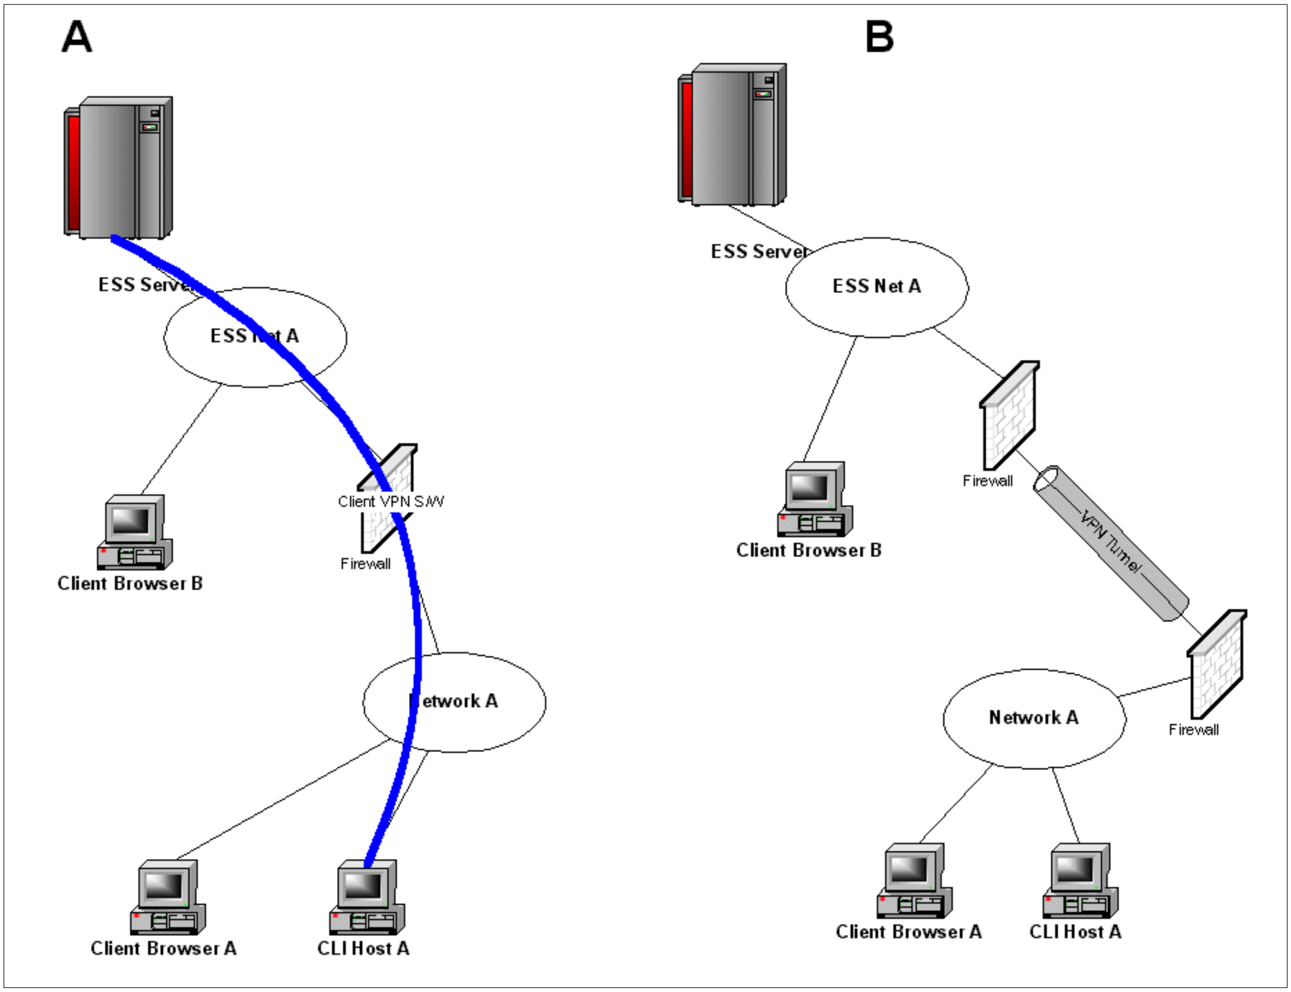

Figure 3-13 ESSNet A and B

Desk Top On Call (DTOC) or PCanywhere

Another solution, which is less complicated is to install DTOC or PCanywhere on an existing NT server behind the firewall. The server must have the CPU computing specifications and Memory allocation to handle ESS Copy Services remotely. By simply accessing this server remotely and using Netscape browser to connect to the Web interface of the ESS, you can perform the functions of the ESS Copy Services.

Customer and IBM personnel will use the information you enter in the work sheet during the installation process. For example, the IBM SSR 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.

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 3-8 on page 44. 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 many 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 section "TCP/IP Configuration" of the Communication Resources work sheet needs to be specified (see corresponding appendix in *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294).

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.

ESS Copy Services primary and backup server TCP/IP configuration

If you plan to use the ESS Copy Services function of the ESS to manage PPRC or FlashCopy, you need to specify the host name and TCP/IP address of the ESS 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 SSR 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 SSR, using the service terminal, to allow host name resolution. This needs to be done for each participating ESS subsystem.

The ESS Copy Services primary and backup server TCP/IP configuration is not included in the Communication Resources work sheet, it is included in the Communication Resources work sheets for ESS Copy Services (see corresponding appendixes in *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294).

ESS Specialist configuration

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

- View
- Operation
- Configuration
- Administration

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.

Call home and remote services

In this section of the Communication Resources work sheet 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.

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 Communication Resources work sheet.

E-mail configuration

Here you specify the following:

- Maximum error notification count per problem

  The ESS sends error notifications (0-9) to e-mail addresses and to the call home destination. Enter the total number of notifications (default is 1) that you want sent to each recipient for each problem. A number greater than 1 increases opportunities for delivery.

  Note: If you enter a value of 0, the ESS does not send an error notification.

- E-mail destinations

  The ESS sends error notifications and information to the destinations that you enter here if you have attached your LAN to the ESSNet external hub.

  – Enter the full e-mail address for each destination, for example:

    ```
    maria@host.com
    ```

  – For each recipient, check one box: **errors**, **information**, **all**, or **none**. This specifies which notification the recipient receives.

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).

4

Availability

In this chapter, we describe ESS high availability features. For example, in most cases of internal component failure, recovery by the ESS is automatic and requires no user intervention. 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 recovery situations, their effect on availability, and the impact of repair. In most cases of internal component failure, recovery by the ESS is automatic and requires no user intervention.

The main categories of possible ESS component failure and recovery are:

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

The ESS 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 failures.

4.1.1 Host adapters

A failure in an ESCON, FICON, FCP 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, FCP, ESCON, or FICON 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 with no disruption. Additionally, performance can be significantly improved by configuring multiple physical paths to groups of heavily used logical disks.

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.

FICON

Plan for a minimum of four FICON channels per ESS. Typical configurations will have eight FICON channels when the disk capacity reaches 3.4 TB.

As with ESCON adapters, spread the FICON host adapters across all the adapter bays.

Define a minimum of four FICON cards per path group, as reflected in the IOCP or IODEF.

SCSI and FCP

For SCSI and FCP, if the Subsystem Device Driver (SDD), 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 the IBM TotalStorage Enterprise Storage Server Specialist (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 Host Adapter Ports panel. Refer to 15.2.5, “Configure Host Adapter Ports” on page 247.

Add the desired logical volumes to the new port. See Section 15.2.7, “Modify Volume Assignments” on page 256.

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 I/O bay.

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 the 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 SSR to perform functions such as installation of microcode updates or to replace components when a failure does not cause automatic failure, for example the ethernet adapter, diskette drive, or serial ports.

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 cache installed in the cluster and write data is stored in the NVS of the opposite cluster. The CPC (processors and associated logic) in the cluster, together with the associated NVS in the other cluster comprise a subsystem. Refer to the diagram in Figure 4-1.

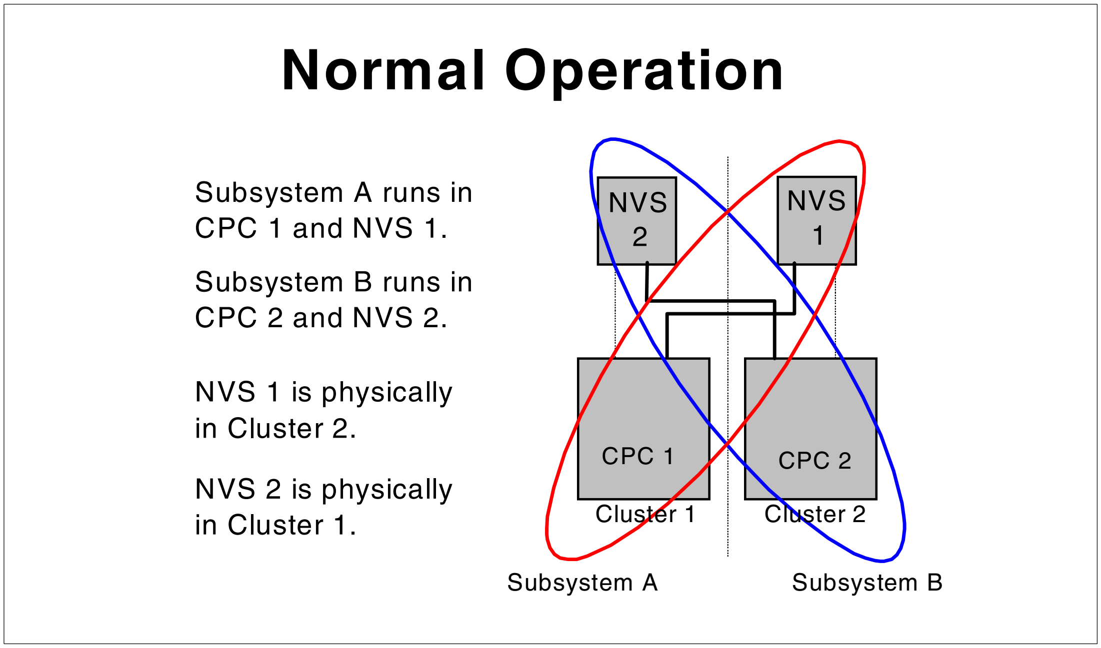

Figure 4-1 ESS cluster configuration for normal operation

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 exist in the NVS and cache of the other cluster. No data is lost.

The failover process is shown in Figure 4-2. 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 CPC2, and NVS 1, both physically located in cluster 2. Access to the disk arrays is through the DAs in CPC 2.

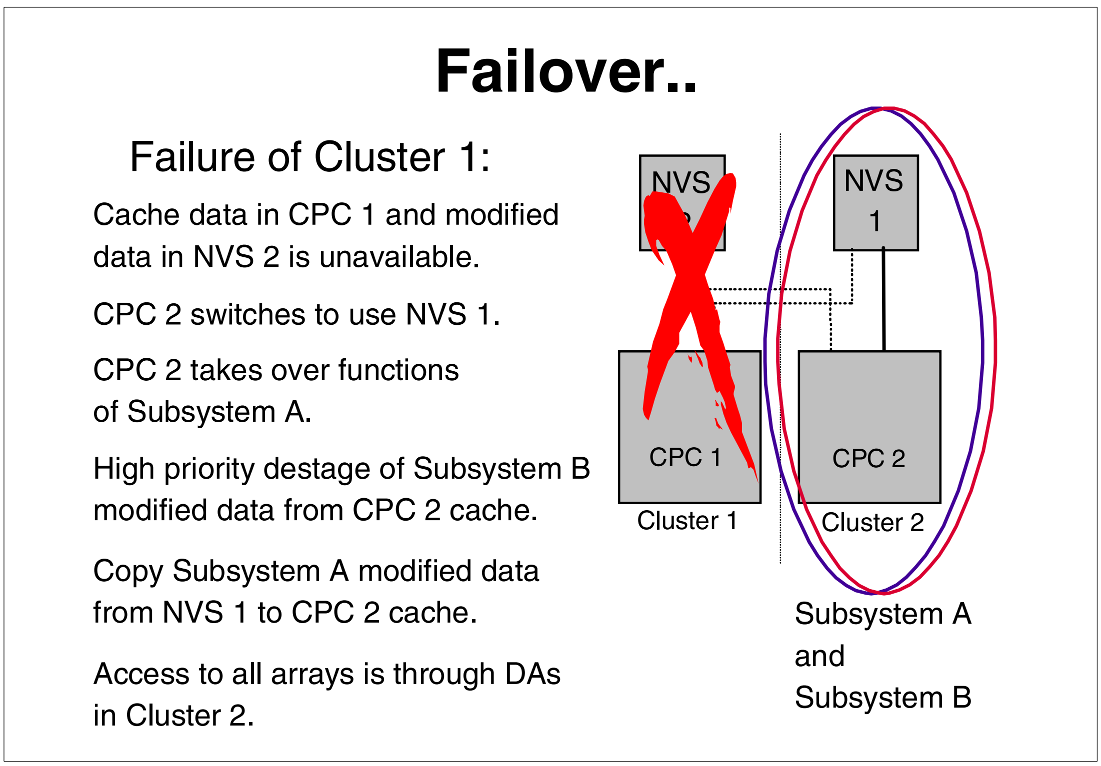

Figure 4-2 ESS cluster failover process

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.

Cluster failback

When servicing of the failed cluster is complete, the failback process is used to return to the normal configuration. Figure 4-3 shows the failback process.

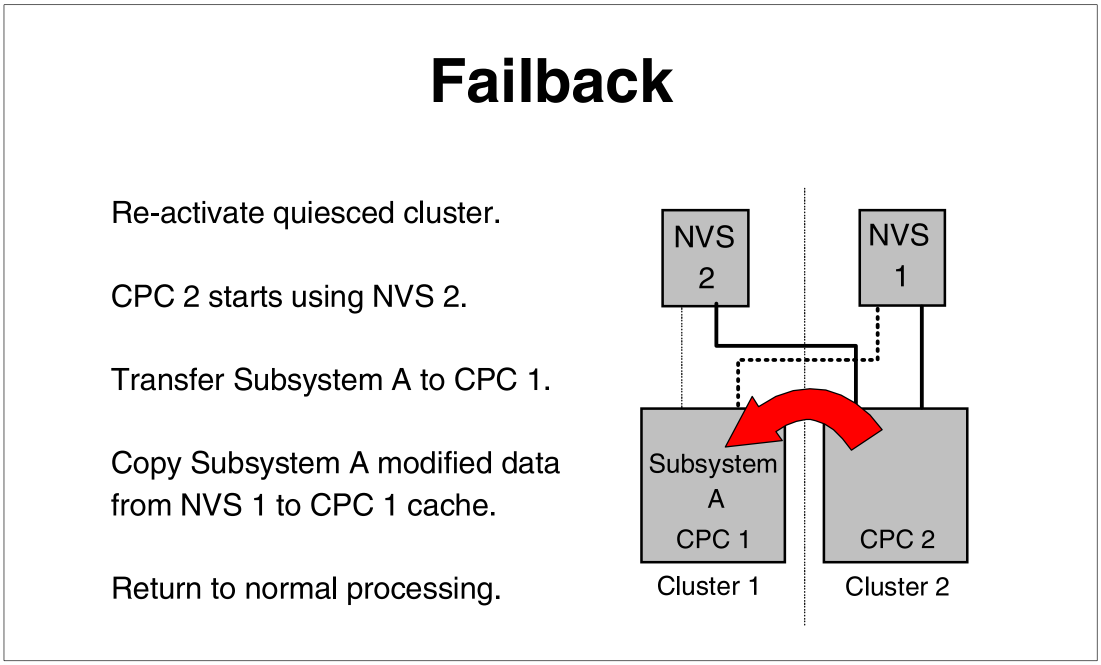

Figure 4-3 ESS cluster failback process

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

Cache

Each of the two storage clusters has a high performance cache consisting of 3GB, 4GB, 8GB, 12GB or 16GB of Error Correcting Code (ECC) Synchronous Dynamic RAM (SDRAM). The 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 among the following:

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.

Device adapters

At any given time, a disk group (RAID-5 array or JBOD) set 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 disk groups in 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 and failback in the normal manner for repair of cluster components.

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 device adapters each have one remaining path around the loop to all disks (apart from the failing disk) on the loop. The failure reduces the maximum bandwidth of the loop and may have a measurable effect on the performance of the logical devices located in disks on the loop if they are heavily utilized.

4.1.3 Disk drives

The disk drives of a disk group can be configured as a RAID-5 array or as a set of non-RAID disks (JBOD). For a disk drive failure, continued access to data depends on whether the failed drive was configured as part of a RAID-5 array.

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, 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 or 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 and 1/8 for a 7+P array.

For sequential read and write operations, 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 for sequential operations 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 failures 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.

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 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 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 is 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 disks drives.

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 and the data striping across multiple drives in the RAID 5 array.

4.1.4 Power and cooling

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

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.

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 an orderly shutdown. The batteries can be replaced concurrently by the IBM SSR.

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.

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
- Four storage cage DC supplies for one disk drive cage, or six 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.

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 4-4

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.

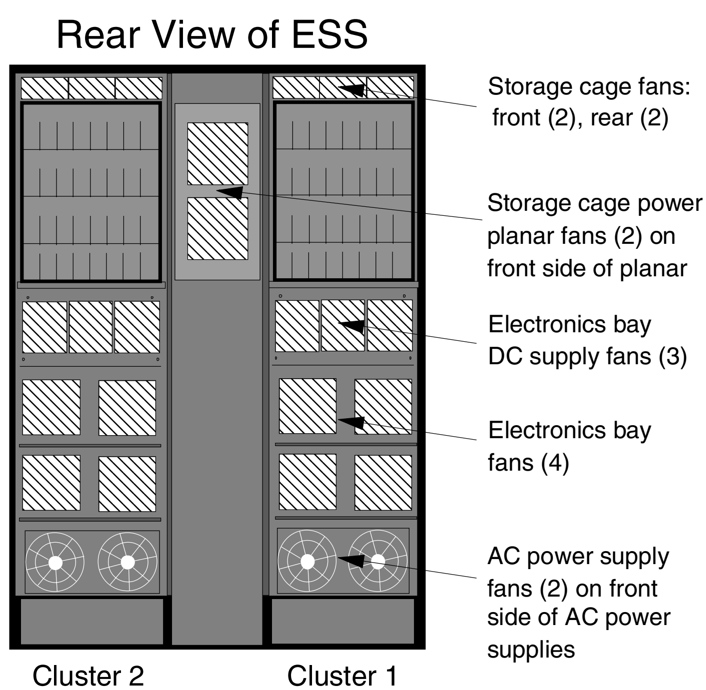

Figure 4-4 Cooling fan groups in the ESS base frame

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.

4.1.5 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 ESS Copy Services function of the ESS is used to manage PPRC, ethernet is used for communication between the ESS 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 nonoperational. 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 SSR 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 send 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 “ESS internal recovery” on page 54, 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 the following:

- **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 (See 3.7, “Call Home and Remote Support” on page 38). The IBM SSR is dispatched by IBM support to repair the problem.

The SSR uses the ESS Master Console (See Section 3.7.1, “IBM TotalStorage Enterprise Storage Server Master Console” on page 39 for a description) to access the maintenance procedures. This was previously done using the MoST (Mobile Service Terminal) service terminal connected to a cluster to access the maintenance procedures (See Figure 4-5). The problem is opened for service and the required resource is quiesced (for example, a cluster could be made available for maintenance by failover). The SSR 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.

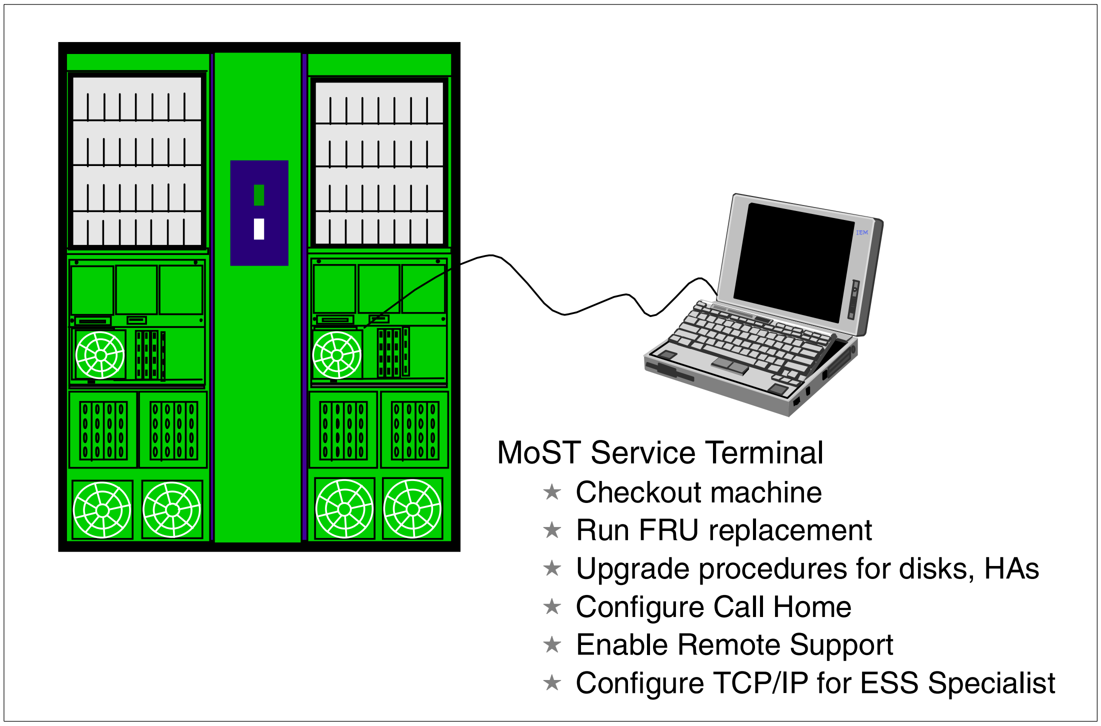

Figure 4-5 Service terminal functions

The MoST is an IBM Thinkpad running a terminal emulator that will be used by the IBM SSR if your ESS has not yet been upgraded to use the ESS Master Console feature (FC 2717). It is used by the IBM SSR 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.

Important: The functions of the MoST Service Terminal can now be performed using the ESS Master Console. See Section 3.7.1, “IBM TotalStorage Enterprise Storage Server Master Console” on page 39.

4.2.2 Upgrades

The 2105 Models E10, E20, F10 and F20 are designed to allow non-disruptive upgrade of cache capacity and NVS capacity, and addition of device adapters, ESCON, SCSI and Fibre Channel (SCSI-FCP for Open and FICON) host adapters, and disk drive bays.

The 2105 Models E10 and E20 comes 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. The 2105 Models F10 and F20 have options of 8 GB, 16 GB, 24 GB and 32GB of cache capacity.

An expansion frame can be non-disruptively attached to a base frame by an SSR 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 Non-disruptive 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.

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.

Host adapters

ESCON, SCSI, and FCP (FICON and Open Fibre Channel) 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 (four to eight ports) is quiesced. All the ports, ESCON, FICON, SCSI or FCP, 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, FICON, SCSI, FCP 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.

CPI cables

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

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, I/O 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 SSR 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.

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.

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 work without it. During the repair, each DA has access to each drive on the loop through only one path around the loop instead of the normal two paths.

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 SSR 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.

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 ESS 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 Section , “DC power supplies” on page 61. 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. The RPC cards, the electronics cage sense cards, and the storage cage fan and power sense cards can be replaced with the rack power on.

Cooling fans

These components are N+1, as described in Section , “Cooling fans” on page 62. 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 of 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
- Simple network management protocol (SNMP)
- E-mail
- Pager
- Service information message (SIM) - S/390 only
- Environmental recording and editing program (EREP) - S/390 only

Figure 4-6 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 SSR 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 IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) and IBM StorWatch Enterprise Storage Server Expert (ESS Expert).

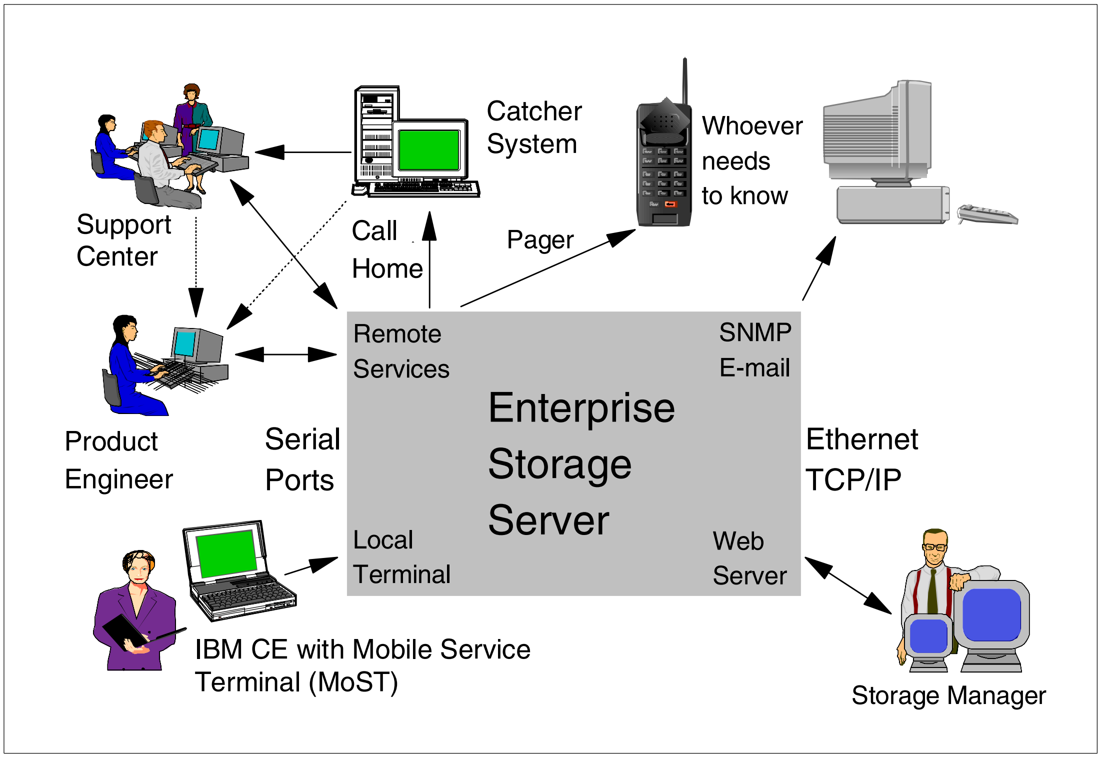

Figure 4-6 ESS communications

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 SSR 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 feature

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 IBM SSR during installation of the ESS.

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.

Remote support makes use of the same modem and serial ports in the storage clusters that are used for Call Home (note that 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, System Support Representative (SSR) 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 SSR 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, proceed as follows:

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

These are some examples of informational messages:

- A new level of Licensed Internal Code (LIC) has been installed.
- New hardware has been installed.
- A remote support specialist has dialled into the ESS.

- The SSR has run the customer-notification diagnostic test. This test verifies that e-mail messages are being received at the addresses set up by the SSR or by the customer, using ESS Specialist.

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. It is used by the IBM SSR to enter the maintenance procedures.
- SRN (Reference Number) — A detailed code used by the IBM SSR.
- Problem Status. See “Status” on page 71 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 SSR.
- 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

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 these 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 work sheet, refer to Appendix A on the Part , “Completing the Communication Resources worksheets” on page 295.

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.

Service information message (SIM)

In the S/390 environment, the ESS uses one of the 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. This message is sent to the system that presents the next active I/O. This system could be any system in the complex and not necessarily the production system.

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.

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.

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. See Section 11.5.1, “DEVSERV command” on page 188.

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 a host attachment 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 and FICON path failures include problems with host channels, problems in ESCON or FICON cables, patch panels, extenders or switches, and faults in ESS ESCON or FICON host adapters. SCSI and FCP path failures in open systems can be caused by problems in host SCSI initiators, other devices on a SCSI interface, or in physical cabling or fibre. 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.

Subsystem device driver (SDD)

For UNIX, Windows NT, or AS/400 environments, multiple SCSI and/or FCP paths to a device must be supported by software. IBM Subsystem Device Driver (SDD) is a software product that provides availability and I/O load balancing for the pSeries, Intel PC based servers, HP, Compaq and Sun environment. Other platforms may be supported by other non-IBM products. Please refer to Section 13.4, "Subsystem Device Driver (SDD)" on page 208 for support information.

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

SDD 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 4-7.

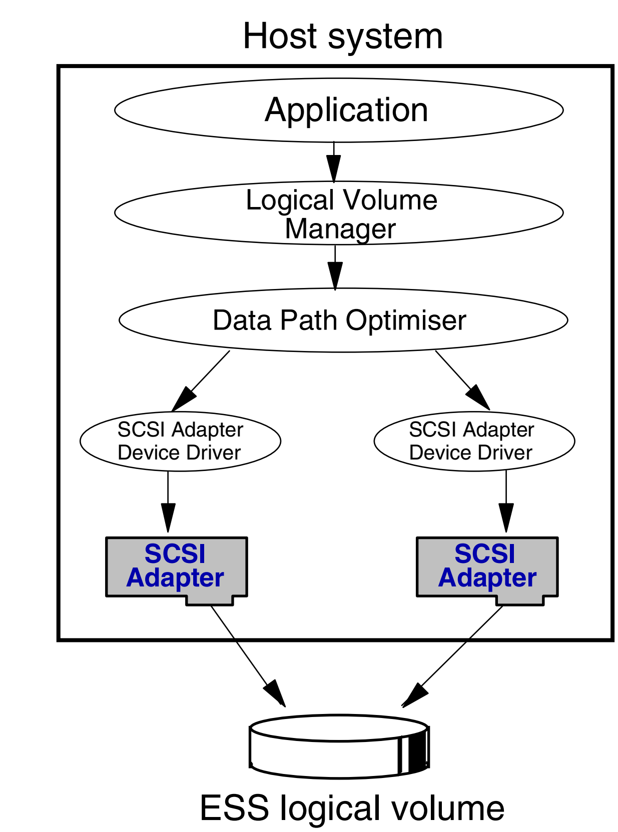

Figure 4-7 Function of subsystem device driver (SDD)

Shared LUNs are supported by SSD

SDD and HACMP can be run in an environment where more than one host is attached to the same LUN (a multi-host environment). This includes clustered hosts such as RS/6000 servers running HACMP and NT High Availability Clusters.

Supported systems

For a list of supported systems, please refer to Table 13-1 on page 207.

Windows NT

SSD for Windows NT is installed using the install shield. It is self configuring. The Subsystem Device Driver 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 Subsystem Device Driver.

Conversion scripts replace hdisk devices by vpath devices for volume groups. Affected file systems must be unmounted. The hdisk devices are still online, but to use the SDD 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. By simply typing in the command **datapath** you may view the options available to the command. Following are some examples:

```
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. SSD supports driver traces to assist in resolution of single path and intermittent problems that otherwise can be difficult to isolate.

Refer also to the *IBM Subsystem Device Driver Installation and Users Guide*. You can obtain a copy on the Web at:

http://ssddom02.storage.ibm.com/disk/ess/related.html

Once there, click the related PDF file.

We recommend that SDD, 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 iSeries 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.

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 4-8. 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.

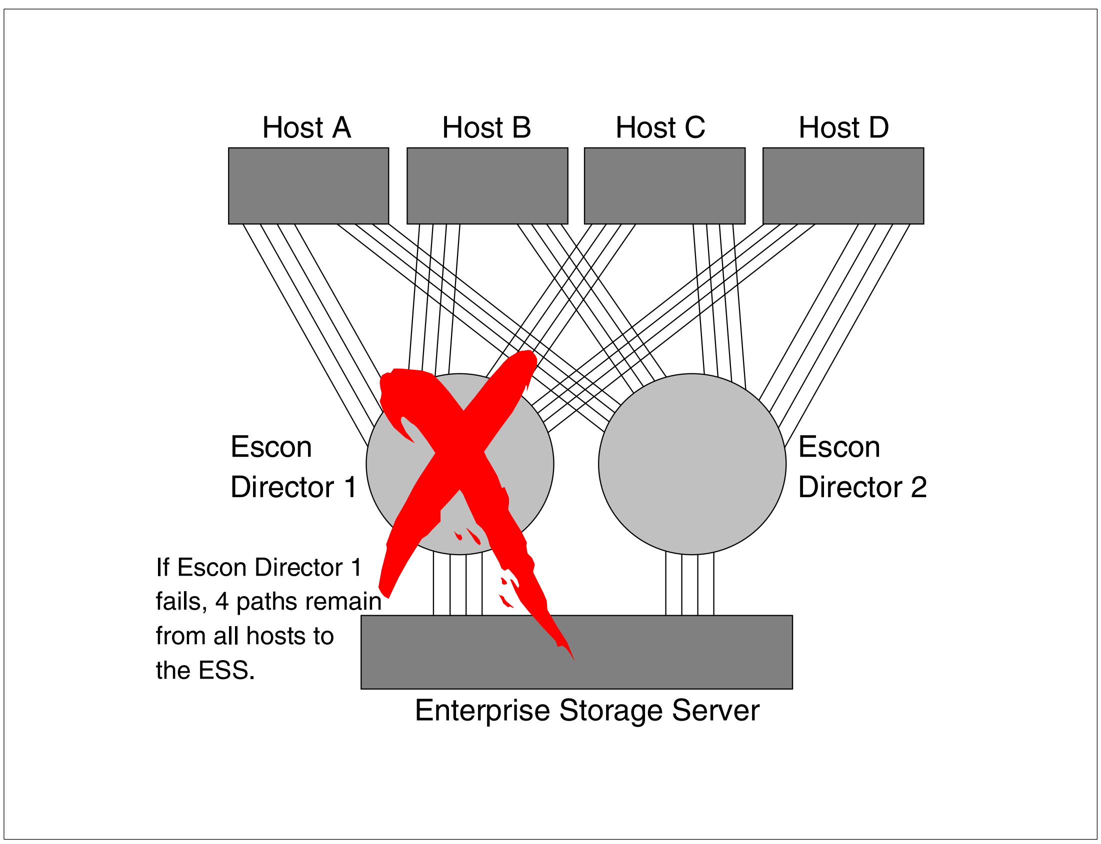

Figure 4-8 Multiple ESCON switch configuration

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 4-9. 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.

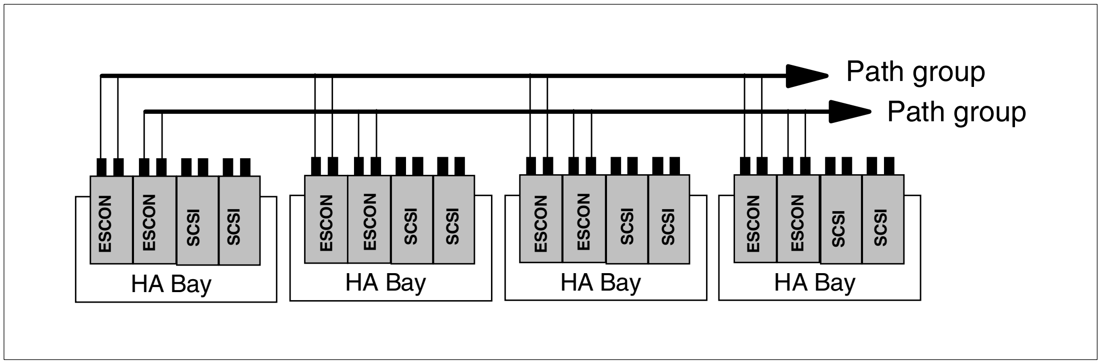

Figure 4-9 Physical configuration of ESCON paths

This configuration is recommended for best performance. Refer to "Host adapters" on page 27 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.

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 (using channel extenders). PPRC can be used for both S/390 volumes and FB volumes. PPRC for FB and CKD volumes can be managed using the ESS Copy Services function of the ESS. 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 storage subsystems. A feature code must be ordered to enable PPRC on the ESS.

Extended Remote Copy (XRC)

Extended Remote Copy (XRC) can be used in the z/OS environment to create and maintain asynchronous copies of ESS CKD volumes. The ESS can have 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.

The *IBM TotalStorage 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 Celsius (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.

5

IBM TotalStorage Enterprise Storage Server Specialist

In this chapter we describe the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist), and the Web interface for the ESS. This is the configuration and administration interface of the ESS. It is the most important interface between the hardware, the software, and the storage administrators.

5.1 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 IBM TotalStorage Enterprise Storage Server Master Console (ESS Master Console), or from a remote workstation using a Web browser.

Using a secure Internet connection (LAN with a Web browser), such as Netscape Navigator, or Microsoft Internet Explorer, your storage administrator can coordinate the consolidation effort and easily integrate storage capacity into the ESS.

The ESS Specialist provides you with the ability to do the following:

- 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 ESS status — Logical schematic of the ESS environment including the host attached ports, controller and cache storage, device adapters, devices and host icons may be checked.
- View and update the configuration — A color schemed view of the storage, including 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 may be viewed.
- 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 logically attached host connections to the logical volume.
  – 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 level.
- Select one of the following authorization levels for each user:

  – Viewer. A viewer can view the current configuration and status information.
  – Operator. An operator can perform view and operation functions, such as changing the remote service and PE password
  – 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 or revoke passwords and levels of authorization.
- Web support for ESS Copy Services (PPRC and FlashCopy).

The remainder of the chapter will give you general topics of the ESS Specialist. For ESS Specialist as it relates to S/390, refer to Chapter 8, “CKD storage configuration” on page 101. For ESS Specialist as it relates to fixed block storage, refer to Chapter 15, “ESS configuration for open systems fixed block storage” on page 235.

For detailed information of how to use the ESS Specialist, refer to *IBM TotalStorage Enterprise Storge Server Web Interface User’s Guide,* SC26-7346.

5.2 ESS Specialist prerequisites

The ESS Specialist interface (and the ESS Copy Services interface) is a set of Java applets, which are programs that are dynamically loaded by the browser, and which execute within your browser. When you request a change to the configuration, the Java applets communicate with the microcode running on the ESS clusters to retrieve the current configuration data, submit the requested configuration change, and display the outcome of the request.

The ability to run programs (applets) inside a Web browser is provided by Java technology. You must use a browser that contains the proper Java Virtual Machine (JVM) implementation to support these applets. The browser software provided by different companies, and even different versions of the same browser, vary widely with respect to their JVM support. Consequently, not all browsers are capable of supporting the ESS Specialist or ESS Copy Services.

The ESS Web interfaces support both the Netscape Navigator and the Microsoft Internet Explorer (MSIE) versions listed in Table 5-1.

Table 5-1 Web browsers supported by ESS Web interfaces

<table>
<tr><th>Netscape level (See Note 1)</th><th>MSIE level (See Notes 2, 3, 4)</th></tr>
<tr><td>Netscape 4.04 with JDK 1.1 fixpack</td><td>MSIE 4.x with Microsoft Java Virtual Machine (JVM) 4.0 or 5.0</td></tr>
<tr><td>Netscape 4.05 with JDK 1.1 fixpack</td><td>MSIE 5.x with Microsoft JVM 4.0 or 5.0</td></tr>
<tr><td>Netscape 4.06 (no fixpack required)</td><td></td></tr>
<tr><td>Netscape 4.5x (no fixpack required)</td><td></td></tr>
<tr><td>Netscape 4.7x (no fixpack required)</td><td></td></tr>
<tr><td colspan="2">Notes:
1. The ESS Web interfaces do not support Netscape above version 4.7.x
2. If your ESS is running with ESS LIC earlier than level 1.3.0 or SC01206, the performance of the ESS Web interfaces on MSIE 5.0 with JVM 5.0 is slower than with Netscape. It is recommended that you use Netscape as the browser or move to LIC level 1.3.0 or higher
3. MSIE 5.0 with JVM 4.0 is supported with all levels of ESS code. However, it is not recommend that you change JVM 5.x to JVM 4.0 on the ESSNet machine in order to improve performance. It is not trivial to change the JVM to a lower level.
4. The ESS Master Console running Linux does not support the MSIE browser.</td></tr>
</table>

The minimum recommended browser hardware is a 166 MHz processor and 32 MB memory, however, greatly improved performance can be achieved with larger configurations (233MHz and 128 MB RAM).

The ESS Specialist is not enabled by default. At installation time, the IBM SSR 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.

5.3 ESS Specialist components

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

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

The SFOI, Information server and 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 (see additional description in 5.2, “ESS Specialist prerequisites” on page 81).

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 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 Socket 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.

5.4 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. You can connect through either cluster but we recommend that once you start updating and modifying from one cluster that you continue to make your changes through that cluster, for the duration of the change. This is to give both clusters time to synchronize.

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 5-1. Then bookmark it for ease of use. You can also enter in the TCP/IP address assigned to the cluster (See Figure 5-2). When assigning the TCP/IP address, we recommend that you make the last field of the second cluster only one digit higher than the first cluster. For example if the left cluster or first cluster is 9.67.51.170, then make the right cluster of second cluster 9.67.51.171.

Figure 5-1 Entering the URL using the cluster name of your ESS

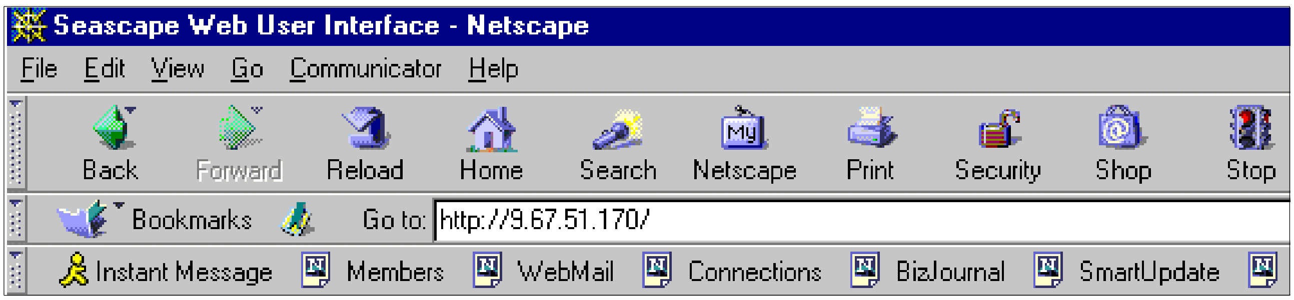

Figure 5-2 Entering the URL using the TCP/IP address

5.4.1 ESS Master Console

To aid you 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 (See Section 3.7.1, “IBM TotalStorage Enterprise Storage Server Master Console” on page 39). The ESS Specialist will always be available to the user or IBM support personnel 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.

5.4.2 ESS Specialist panels

The ESS Specialist is the major interface between your storage administrator and the ESS. It is used to configure the ESS logical storage environment and provide administrative and problem status information.

This section gives a very brief introduction to some of the ESS Specialist panels and features. For detailed information about the use and panels of the ESS Specialist, refer to *IBM TotalStorage Enterprise Storge Server Web Interface User’s Guide,* SC26-7346.

For ESS Specialist as it relates to S/390, refer to Chapter 8, “CKD storage configuration” on page 101. For ESS Specialist as it relates to fixed block storage, refer to Chapter 15, “ESS configuration for open systems fixed block storage” on page 235.

The ESS Specialist is frame enabled. When running inside a Web browser, it presents two main frames on the Admin panel — a navigation frame on the left with hyperlink buttons for accessing ESS Specialist functions and the working frame on the right. 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.

5.4.3 Welcome panel

From the screen similar to the one shown in Figure 5-3, you may select to enter the ESS Specialist, the ESS Copy Services, or Tools from the navigation frame.

Figure 5-3 ESS Specialist Welcome screen

Note: The screen shown in Figure 5-3 should not be closed, just minimize it. This is because closing it will affect the associated applets.

After selecting the ESS Specialist, options for managing the ESS can be selected from a panel similar to the one in Figure 5-4. This is the introduction and welcome screen. You can see the serial number, model, machine type and world wide node name of your ESS.

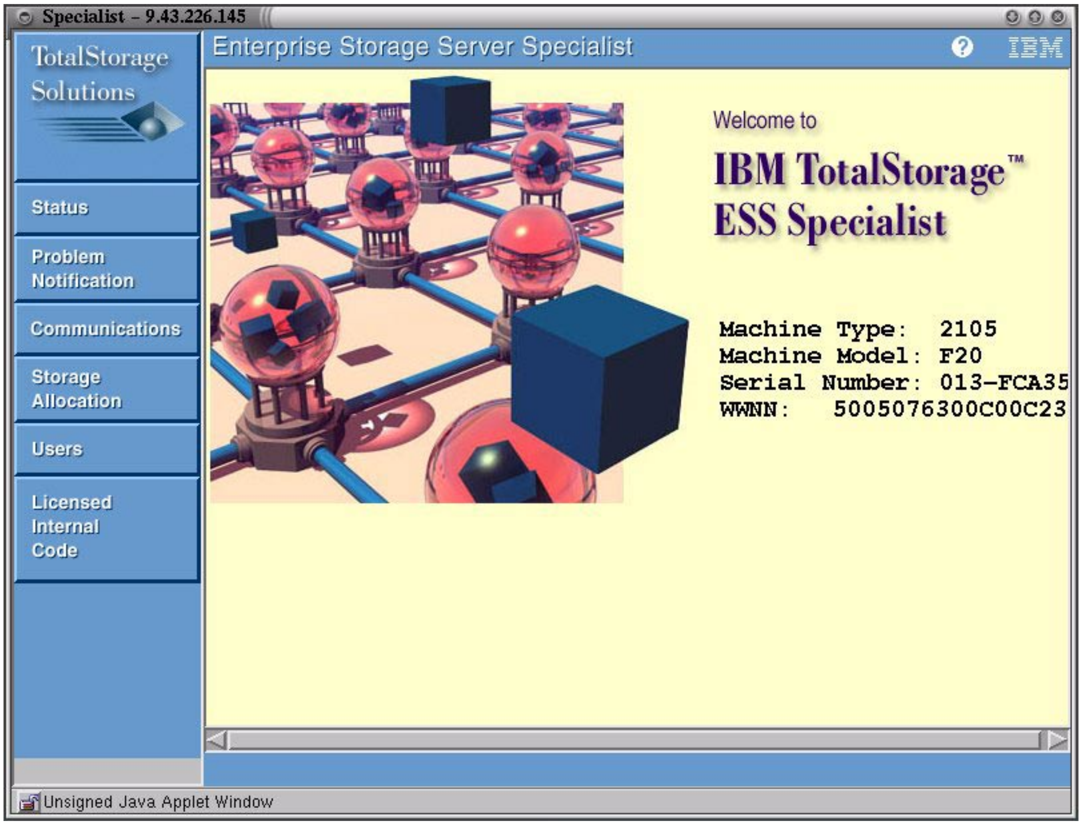

Figure 5-4 ESS Specialist function selection panel

5.4.4 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 as shown in Figure 5-5 is located on the top right-hand side of all panels.

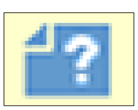

Figure 5-5 ESS Specialist help icon

5.4.5 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. By selecting the Users button from the ESS Specialist panel as shown in Figure 5-4, a user can be restricted to one of four different authorization levels — viewer, operator, configurator or administrator. See Figure 5-6.

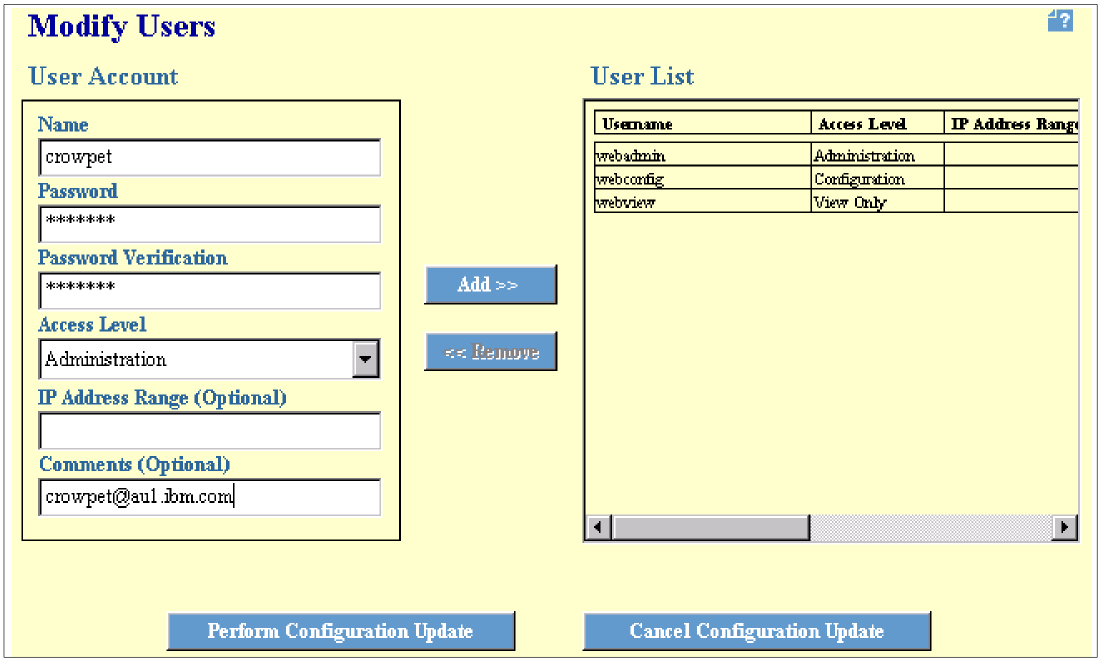

Figure 5-6 Modify users panel

A viewer can view the current configuration and status information, while a configurator can view the current configuration, 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.

6

IBM StorWatch Enterprise Storage Server Expert

In this chapter, we discuss the IBM StorWatch Enterprise Storage Server Expert (ESS Expert), an IBM offering which is a storage management tool for managing the ESS assets, capacity and performance. We explain what it is and how it relates to the ESS.

6.1 Using ESS Expert

The IBM StorWatch Enterprise Storage Server Expert (ESS Expert) provides storage resource management for the IBM TotalStorage Enterprise Storage Server and the IBM TotalStorage Enterprise Tape Library (ETL). This chapter will only cover the applications and functions of the Expert as it relates to the ESS (the ESS Expert).

The ESS Expert is a program product which can be purchased and used as a storage monitoring and management tool for the ESS. this tool is accessed using a Web browser interface.

The ESS Expert prepares reports and graphical charts from the data it collects. Data may be collected from any of the ESSs attached to the same network as the ESS Expert.

ESS Expert helps you with the following tasks:

- **Asset management:** Data is collected and presented in either summary or detail reports on the ESS. The type of information reported includes system names, LIC levels, device nodes, clusters and features installed.
- **Capacity management:** ESS Expert can be used for central tracking and reporting of the ESS capacity. Assigned storage, unassigned storage and freespace can be reported on each ESS.
- **Performance management:** Performance data is collected and prepared for presentation. The storage administrator may view the number of I/O requests based on cluster, device adapter, rank or volume. Performance reports on caching statistics and disk utilization can be observed to analyze the I/O subsystem performance.

The information collected by the ESS Expert along with host based tools will help the storage administrator to manage the capacity and performance of the ESS. Further details on using the ESS Expert can be found in *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102 which details the planning, implementation and usage of the ESS Expert.You can download a copy of this guide from the Web at the following location.

6.2 Accessing ESS Expert

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

To familiarize you with the ESS Expert, we are providing some screen shots. Enter the IP address of the ESS Expert from your Web browser and you'll be given the sign-on screen as shown in Figure 6-1.This is your doorway into the ESS Expert.

The ESS Expert interface consists of these major components:

- **The information frame:** This displays status and messages. You can press the help icon located in the information frame any time for additional help.
- **The navigation frame:** This presents folders to access the ESS Expert applications and tasks.
- **The work frame:** This is the main area in the center of the browser window which 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. Therefore, 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.

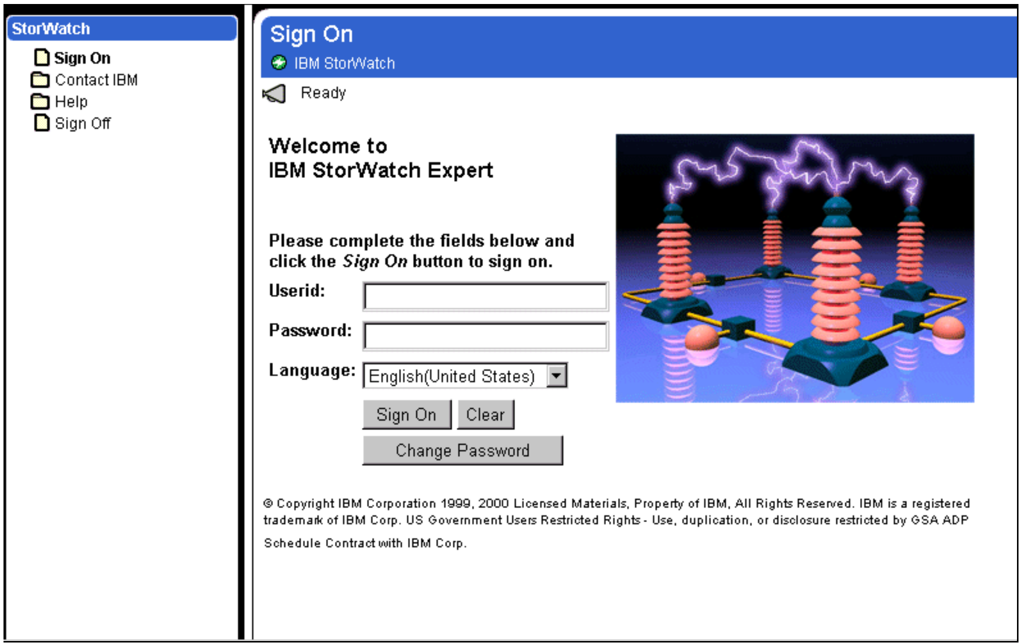

Figure 6-1 ESS Expert: Welcome panel

6.3 Navigating through ESS Expert

Once you have signed on to the ESS Expert, you'll be presented with the *Introduction* panel. We have expanded the Manage StorWatch folder in the navigation frame for you in Figure 6-2. From this folder you can perform tasks which are related to managing the ESS Expert system. These tasks include discovering and managing the nodes (clusters of the ESS), performing administrative tasks for users of the ESS Expert, administering the ESS Expert database, and monitoring the ESS Expert tasks. Further information on these tasks can be found by invoking the help function or by reviewing the IBM Redbook entitled, *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102.

Figure 6-2   ESS Expert: Managing the ESS Expert

Now we have expanded the Manage ESS folder in the navigation frame and presented it in Figure 6-3. It is from this application section that you access the tasks for managing the assets, managing the capacity and managing the performance of your ESS. For managing assets and capacity, you may want to set up a task to run periodically to gather the information on your ESS. You may then display pertinent information about the ESS, such as the ESS serial numbers and IP addresses, active, previous and next LIC levels. From a capacity perspective, you can view reports which show how the storage has been assigned, how much cache and NVS is available for each cluster and other capacity related information.

When using the managing performance for the ESS, you have the capability to set up tasks which can be run regularly or one time only to collect the performance statistics of the ESS being monitored. In order to view the reports of this collected data, you will need to run a data preparation task. Once the data is prepared in report format, you can view the performance of your ESS based on the cluster, device adapters or raid rank. Refer to *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102 for sample performance reports and how they can be used.

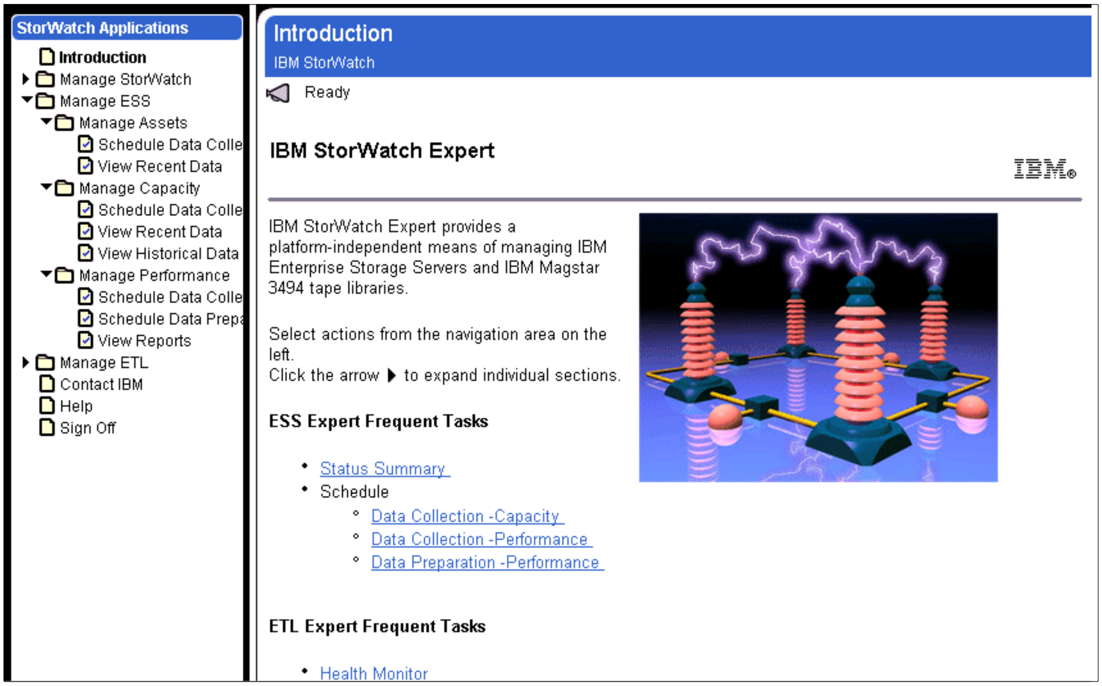

Figure 6-3 ESS Expert: Managing ESS

The reports presented by the ESS Expert supplement the host based tools which monitor system performance. The complementation depends on the host. For instance, the z/OS systems may use RMF from which they can obtain the caching statistics based on logical volume along with the rank statistics. With a UNIX or NT based system, **`iostats`** or the performance monitor will give you a further picture of the ESS by providing the caching statistics for you to monitor and use to tune your ESS.

6.3.1 Additional documentation on ESS Expert

We suggest that you go to *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102 for more detailed information on the ESS Expert. This publication shows how to plan for the ESS Expert and how to install, implement, and use the tool.

Part 2

Implementation in the zSeries environment

In this part we discuss the processes for implementing the IBM TotalStorage Enterprise Storage Server in the zSeries environment. We first cover software requirements for the zSeries operating systems, then discuss in length how to design and implement an ESS logical configuration that meets your needs. We also cover monitoring and reporting enhancements in support of the ESS, discuss migration scenarios, and the use of copy services with the ESS.

7

zSeries systems support

In this chapter we first cover the S/390 and zSeries operating system basic software support as of the initial ESS announcement. We then cover software support for the more recent ESS enhancements. We list the required minimum software release levels for the zSeries operating systems that support the ESS. We do not provide PTF level information in this redbook. For information regarding APARs and PTFs required in each case, please refer to the appropriate Preventive Service Planning (PSP) information, or contact your IBM representative.

7.1 Basic support

The ESS is supported by the z/OS, OS/390, z/VM, VM/ESA, VSE/ESA, and TPF operating systems.

7.1.1 OS/390 and z/OS support

For the currently supported releases of OS/390 and z/OS the ESS runs in what is called Exploitation mode. Exploitation mode means that the host recognizes the ESS as a 2105 control unit and can exploit the PAV capability. You use CNTLUNIT UNIT=2105 to define the ESS. The currently supported releases are DFSMS/MVS 1.4 and higher.

Exploitation support is provided with PTFs for OS/390 1.3-2.6 with DFSMS/MVS 1.4. Support is integrated in OS/390 2.7 with DFSMS/MVS 1.5 and above, and in z/OS 1.1 and above. ICKDSF (refresh) release 16, EREP version 3 release 5, and DFSORT for MVS/ ESA release 13 are also prerequisites.

The minimum software level for dynamic PAV support is OS/390 2.7 with DFSMS/MVS 1.5. Earlier releases only support static PAV. For dynamic PAV, the Workload Manager (WLM) must run in goal mode.

Note: For latest information on software levels and software maintenance requirements you should check the corresponding preventive service planning (PSP) bucket.

Previous releases of OS/390 and DFSMS/MVS now not supported, will have to run the ESS in either Transparency or Toleration mode. These modes of operating the ESS must be considered as transitional because the recommendation is to have all your systems at the currently supported levels of software and so capable to run the ESS in Exploitation mode.

Transparency support provides the base functions of an IBM 3990 Model 6 Storage Control. OS/390 sees the ESS as multiple 3990 control units. PAV support is not available in transparency mode. The host cannot share an I/O Definition File (IODF) with an exploiting system. You define the logical control units to the host system using IOCP specification CNTLUNIT UNIT=3990.

Toleration support permits a level of OS/390 which does not support PAV to share the ESS with releases of OS/390 and z/OS which do. The OS/390 host recognizes the ESS as a 2105 control unit and recognizes the base and alias addresses needed to support PAV. With toleration support, only non-PAV UCBs are built, so the system cannot utilize PAVs. However, the host can share an IODF with exploitation-capable systems. You define ESS logical control units using the new unit type, CNTLUNIT UNIT=2105.

7.1.2 VM/ESA support

In this section we discuss VM/ESA support.

CP native support

The following support is provided:

- 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 (no use of new CCWs, no PAV support for CP/CMS use).

Guest support

The following support is provided:

- **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

VM/ESA 2.4.0 support for z/OS and OS/390 guests

The following support is provided:

- **Parallel Access Volumes:** Supported for guest use only
- **FlashCopy:** Supported for guest use only

Utilization of ESS functions

PPRC is supported in VM/ESA and managed by ICKDSF or the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist). FlashCopy can be managed by the ESS Copy Services Web interface.

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.

7.1.3 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. PAV support is not available. New read/write track CCW is not supported.

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.

Utilization of ESS functions

PPRC is supported in VSE/ESA and managed by ICKDSF or the ESS Copy Services Web interface. FlashCopy can be managed by ESS Copy Services Web interface.

Multiple Allegiance and I/O Queueing are the ESS hardware functions. In a shared environment VSE/ESA could take advantage of it, but the priority byte, however, is not set by VSE/ESA and therefore I/O Queueing is not applicable.

7.1.4 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.

TPF support levels

ESS is supported with **TPF 4.1**.

Utilization of ESS functions

With applied corresponding PTFs TPF 4.1 is capable of using 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.

FICON attachment support

TPF supports native FICON attachment. This support requires TPF Version 4, Release 1. The ESS has to be an F10 or F20 at LIC level 1.5.0, or later.

PPRC and FlashCopy support

PPRC and FlashCopy are supported on TPF. This requires TPF Version 4, Release 1. Additional PTFs may be required. PPRC and FlashCopy are supported on the ESS models F10 and F20 with LIC level 1.5.0 or later.

7.1.5 Linux

On zSeries 900 and on S/390 servers you can run Linux natively as a stand-alone or as a logical partition (LPAR). In addition, the S/390 Virtual Image Facility for Linux, and z/VM V4 enable you to run more Linux images than can be deployed using LPARs and provide capabilities to help create and manage these images.

Current Linux for S/390 distributions supported with the ESS are SuSE Linux Enterprise Server for S/390 and Turbo Linux Server 6 for zSeries and S/390. For the most current information on the Linux for S/390 releases and distributions that are supported with the ESS, refer to:

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

7.2 FlashCopy support

The minimum levels for FlashCopy support are:

- DFSMS/MVS 1.3
- VM/ESA 2.4.0 provides FlashCopy support for guests
- VSE/ESA 2.5.
- TPF 4.1

See PSP bucket information for required PTFs.

7.3 FICON support

Native FICON attachment is supported on ESS models F10 and F20 with LIC level 1.5.0, or later, and using the Fibre Channel/FICON host adapters (FC 3021 or 3023). FICON support requires the following minimum levels of S/390 and zSeries operating systems:

- z/OS 1.1, or later, is required.
- OS/390 2.8 is the minimum level.
- System Automation for OS/390 1.3 is the minimum level.
- VM/ESA 2.2 is the minimum level.
- VSE/ESA 2.1.0 is the minimum level.
- TPF 4.1 is the minimum level.

PTFs are required for FICON support. See the PSP bucket for detailed information, or consult your IBM representative.

7.4 Large Volume Support

The ESS initially supported custom volumes of up to 10017 cylinders, the size of the largest standard volume, the 3390 model 9. This was the limit set by the operating system software. The ESS Large Volume Support (LVS) enhancement, announced in November 2001 has now increased the upper limit to 32760 cylinders, approximately 27.8 GB. The enhancement is provided as a combination of ESS licensed internal code (LIC) changes and system software changes, available for z/OS, OS/390, and z/VM.

Requirements

Large Volume Support is available on ESS models F10 and F20 with ESS LIC level 1.5.0, or later, and requires one of the following operating system levels:

- z/OS V1.1, or later + PTF
- OS/390 V2.10 + PTF
- ICKDSF R16 + PTF
- DFSORT R14 + PTF
- z/VM V3.1, or later + PTF

For DFSMS/MVS components this support is provided as a Small Programming Enhancement (SPE) on OS/390 Release 10 (HDZ11F0) and integrated into z/OS Release 1.3 (HDZ11G0).

z/VM supports these large volumes as native device as well as for guest support.

VSE does not support large volumes.

Install large volume support on your systems before defining large volumes on the ESS. Large volume support needs to be installed on all systems in a sysplex prior to sharing data sets on large volumes. Shared system/application data sets can not be placed on large volumes until all system images in a Sysplex have the large volume support installed.

Installation of PTFs on some components will require a system IPL to activate.

Please check PSP bucket information for required PTFs. Check with your OEM software product vendors for changes to their products which may be required in support of large volumes.

Coexistence support

For DFSMS/MVS 1.4 and 1.5, a coexistence PTF is provided that will allow these system levels to coexist in the same Sysplex with LVS systems. You must install this PTF in order to prevent unpredictable results that may arise from systems without large volume support accessing volumes that have more than 10017 cylinders. The coexistence PTF will:

- Prevent a device with more than10017 cylinders from being varied online to the system.
- Prevent a device from coming online during an IPL if it is configured with more than 10017 cylinder.

Coexistence PTFs will also be required by DFSMShsm on all releases prior to OS/390 R10 because of the updates being made to the record format in the DFSMShsm control data sets.

Coexistence support will not be available for DFSMS 2.10 and higher. Install the large volume support prior to using data on large volumes at these release levels. No coexistence support will be available for DFSMS/MVS 1.3 (unsupported release), or earlier.

7.5 CUIR support

In large configurations, quiescing channel paths in preparation for upgrades or service actions is a complex, time consuming, and potentially error prone process. Control Unit Initiated Reconfiguration (CUIR) automates the process of quiescing and re-enabling channel paths, reducing the time required for service actions and reducing the operations staff required efforts, and reduces the possibility for human error.

CUIR requires ESS licensed internal code (LIC) level 1.5.0, or later, and is supported on zSeries and S/390 platforms by the following operating systems:

- z/OS V1.1, or later + PTF
- OS/390 V2.10 + PTF
- z/VM V3.1, or later + PTF

Many of the support PTFs require a system level IPL to activate. CUIR support will also require host channel microcode updates (MCLs). Please check PSP bucket information for required PTFs.

8

CKD storage configuration

In this chapter we provide guidance and recommendations on how to configure an IBM TotalStorage Enterprise Storage Server (ESS) to meet your needs. We begin by introducing the logical configuration components. We then describe the procedures using the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) to configure the ESS for use with S/390 and zSeries hosts.

8.1 Basic concepts

For the S/390 and zSeries environments, the ESS appears to the host as a 2105 control unit with one or more logical control units (LCUs) attaching 3390 devices in either 3390 or 3380 track format. The number and type of storage controls and devices an ESS emulates, that is the *logical configuration* of the ESS, can be selected by the user. One of the key tasks of ESS implementation planning is to decide, what the logical configuration of an ESS should be. During installation, the logical configuration is defined on the ESS so that it will present to the host the desired image. Later, it can be modified if requirements change.

The target configuration you will implement depends on several factors, such as the installed capacity of the ESS, your host system configuration, host software level, and the type of logical volumes you need for your applications. You need to plan for the number and type of logical control units and logical volumes you will define. If your system supports Parallel Access Volumes (PAV), you need to plan how to allocate them. You will also need to decide on which configuration method to use. These items should be planned well in advance of the actual installation in order to guarantee a successful implementation of the ESS.

In this section we introduce the basic concepts you need to understand in order to configure an ESS for use on S/390 and zSeries systems.

8.1.1 ESS Specialist

The IBM TotalStorage Enterprise Storage Server Specialist (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 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, “ESS Specialist” on page 43 for ESS Specialist requirements, and also to Chapter 5, “IBM TotalStorage Enterprise Storage Server Specialist” on page 79 for an introduction to the ESS Specialist.

8.1.2 Logical Control Units

The ESS can emulate up to 16 ESCON or FICON attached Logical Control Units (LCUs) with emulated 3390 devices. The LCUs you want the ESS to present to the host must be configured on the ESS. The LCU and its logical volumes must also be defined in the host IODF to enable host communication with those volumes. Each LCU corresponds to one CKD format ESS Logical Subsystem (LSS). The terms LCU and LSS are therefore often used interchangeably in documentation for CKD storage. The term LCU only refers to a CKD LSS.

An LCU can be defined to emulate one of the following control units:

- 3990 Model 6
- 3990 Model 3
- 3990 Model 3 TPF
- 3990 in Siemens BS2000 compatibility mode.

Every installed ESCON and FICON host adapter port is capable of accessing all configured LCUs. ESCON and FICON attached hosts are identified to the ESS when you make the physical connection between the hosts and the storage server. These hosts are seen as a single ESCON or FICON net in the ESS Specialist for improved graphical presentation.

The LCUs are numbered from 0 to F. (The ESS Specialist typically uses a three digit notation, 0x000 to 0x00F.) These map directly to LSS numbers 00 to 0F. See Figure 8-1.

An LCU can have up to 256 logical volumes defined to it, one per unit address 00-FF. The ESS can therefore have up to 16 * 256 = 4096 CKD logical volumes allocated. If Parallel Access Volumes (PAVs) are enabled for the LCU, each PAV alias takes up one of the 256 addresses.

Each LCU is uniquely associated with one cluster, and one device adapter (DA). On each device adapter you can configure up to two CKD LCUs. The even numbered LCUs are managed by ESS Cluster 1, the odd numbered by Cluster 2. Under normal conditions all I/O to the LCU is done through that particular cluster and DA. Each LCU is also associated with half of the disk groups on the two loops attached to the device adapter. (See explanation of disk groups later in section 8.1.3, “Disk groups and ranks” on page 104.) A LCU can be assigned physical disk capacity from either or both of the loops on the DA. The associations between LCUs, DAs, and clusters are predefined and shown in Figure 8-1.

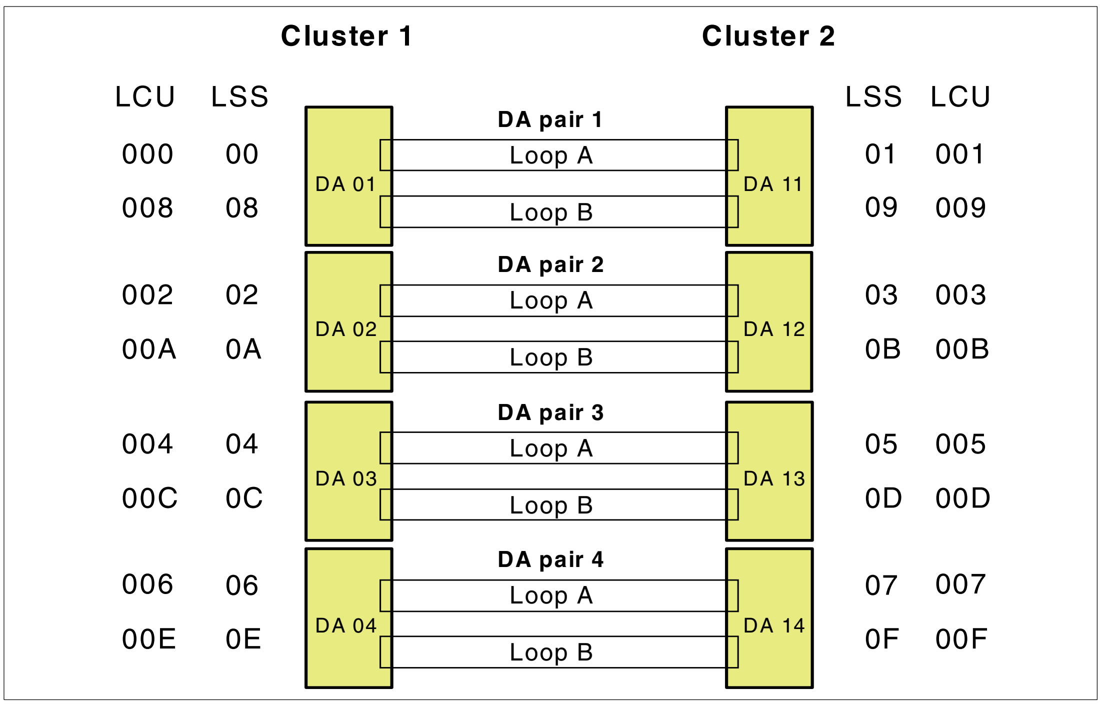

Figure 8-1 Relationship between S/390 LCUs and SSA loops

For example, LCUs 4 and C are managed by Cluster 1 and DA 03 (in adapter pair 3). The logical volumes belonging to LCUs 4 and C are physically located on one or more (maximum of six) disk groups on SSA loops A and B of DA pair 3. SSA loops always contain an even number of disk groups, evenly divided between the two clusters. Thus loops A and B of DA pair 3 may also have one or more disk groups that belong to LCUs 5 and D, and are managed by Cluster 2 and by DA 13 (also in DA pair 3).

In the event of a cluster failover, the other cluster can take over all LCUs and disk groups. Data would then be accessed through the DAs in the surviving cluster.

At ESS installation time the IBM System Support Representative (IBM SSR) can configure the ESS to support either 0, 8, or 16 CKD LSSs and 0, 8, or 16 FB LSSs. The default for both is 16. If you do not plan to use your ESS for open systems, have the maximum LSS value for FB set to 0. This will save some 2MB of cache memory for each LSS. However, we recommend that you keep the number of CKD LSSs in the default 16, as this will maximize the number of devices you can define on the ESS. This does not mean that you have to configure all 16 LCUs initially, it just means that 16 are available for you to configure should you later need them. Changing the Maximum LSS values after the initial configuration will require an IML of the ESS.

8.1.3 Disk groups and ranks

A *disk group* in ESS is a group of eight disk drive modules (DDMs) of the same size on an SSA loop. As physical capacity on the ESS is always installed in pairs of disk 8-packs of the same size on a loop, there can be 0, 2, 4, or 6 disk groups on any loop.

The grouping of DDMs into disk groups is done by the ESS when 8-packs are installed in the ESS. You cannot select which DDMs belong to each disk group. With the ESS Specialist you cannot display which DDMs belong to a certain disk group.

A disk group must be configured before its capacity can be allocated for host system use. You can configure the disks in a disk group to form a RAID-5 array, or leave them as eight individual non-RAID disks (JBOD). The RAID-5 array, or the non-RAID disks must further be configured in either 3390 track format or 3380 track format for use on zSeries hosts. Alternatively, they may be configured as FB storage for open systems use.

The formatted disks are referred to as *ranks*. Thus, a RAID array is called a RAID rank, and the individual disks are called non-RAID ranks (JBOD ranks). A rank is the unit on which logical volumes are later allocated. A CKD rank is assigned to a LCU when the rank is configured.

Each disk group is associated with and managed by one device adapter (DA) and one of the two ESS clusters. A disk group can only be assigned to one of the four LSSs (two CKD LSSs and two FB LSSs) that are associated with that particular DA. Figure 8-2 shows an example from a fully configured ESS. Any of the highlighted disk groups can be assigned to any of the four LSSs, regardless of which loop the group is in. Since there is always an even number of disk groups associated with a DA, you can divide them evenly between two LCUs if you plan to define both LCUs, or in smaller configurations assign them all to one LCU.

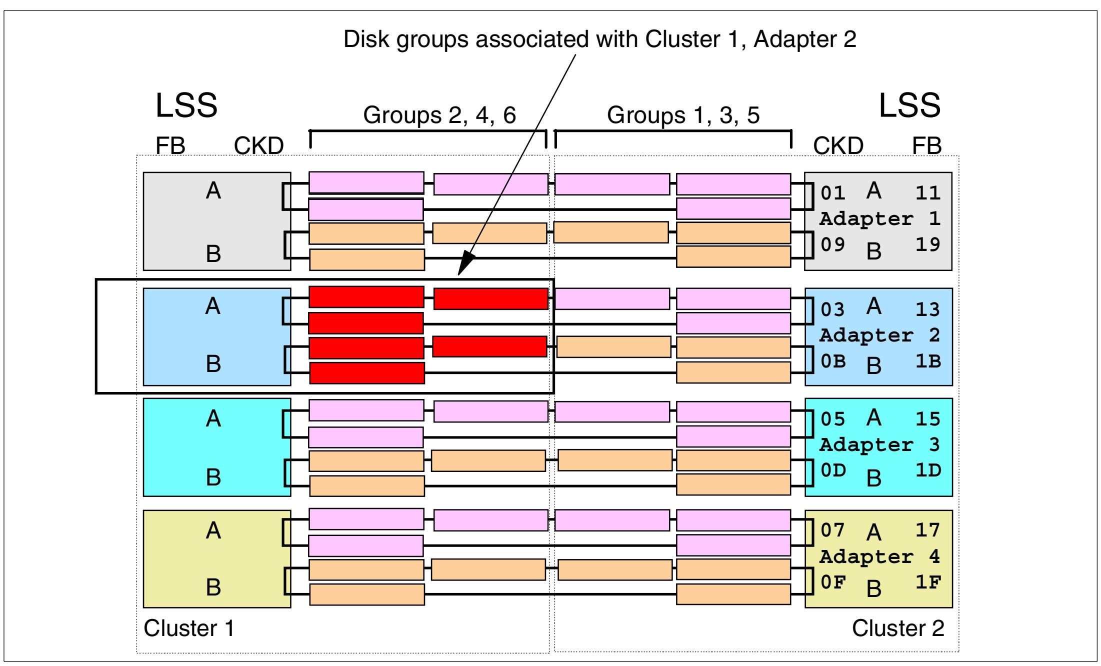

Figure 8-2 Disk group to LSS association

All I/O operations to a disk group are performed by the cluster and DA it is associated with. In the event of failure, the other cluster can take over all operations.

On the ESS Specialist *Storage Allocation* panel each rectangle in the middle represents a disk group (Figure 8-3). The panel only displays configured disk groups. There may be disk groups on the loop that have not yet been defined. Because there is always an even number of disk groups on each loop, you know for certain that there are unconfigured disk groups on the first and last loop in the figure. The panel displays both CKD and FB disk groups. The highlighted groups shown in the figure are CKD storage. They are highlighted because we clicked on the EsconNet host icon (not visible). The unhighlighted groups are configured as FB. The figure also shows which disk groups belong to which cluster.

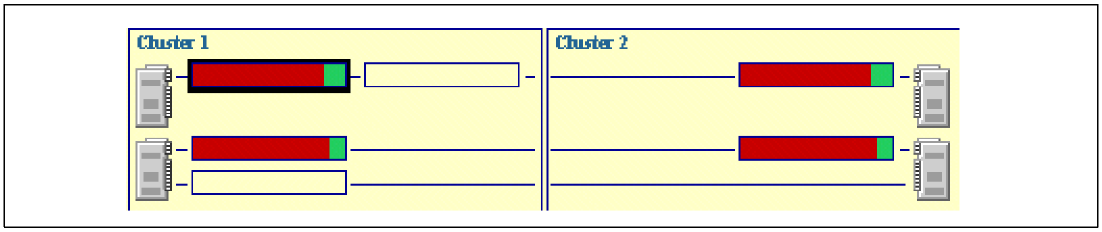

Figure 8-3 Disk groups in the Storage Allocation panel

The physical drives in a disk group typically reside in two 8-packs on a loop. When two 8-packs are installed on a loop, the physical disks drives form two disk groups, both containing four drives from each 8-pack. Over time, the DDMs belonging to a disk group may change due to disk sparing in RAID arrays.

See the *IBM TotalStorage Enterprise Storage Server,* SG24-5465 for more details on disk groups and ranks.

Here is a summary of disk group attributes:

- A disk group contains eight physical disk drives (DDMs).
- A disk group is contained on a single SSA loop. It cannot span both loops on a device adapter pair.
- All DDMs in a disk group are of the same size. A loop may contain different size DDMs.
- The number of disk groups on a loop is the same as the number of 8-packs installed on the loop. That is, it can be 0, 2, 4, or 6.
- The disk groups on a loop are numbered from 1 to 6. The even numbered disk groups are managed by Cluster 1. The odd numbered groups are managed by Cluster 2.
- A disk group can be configured in either of two ways: as a RAID-5 array or as non-RAID disks (JBODs).

8.1.4 RAID arrays

The RAID-5 arrays on ESS will have one of two array configurations:

- **6+P+S**. This setup leaves one spare disk (S) in an SSA loop. The capacity of six physical disks drives is used for storing data and one for storing the parity information (P). The first two disk groups of any given DDM size on a loop will automatically have a spare physical disk drive and the arrays will be 6+P+S.
- **7+P**. In this setup the capacity of seven disks is used for storing data, one for storing the parity. The third and subsequent arrays of a given DDM size in a loop will be 7+P.

Non-RAID ranks may also be used for CKD volumes, but we recommend the use of RAID arrays due to the availability they offer. For this reason our examples mostly cover RAID configurations.

When you configure a RAID array for CKD storage, you can optionally auto-allocate a group of standard volumes on the array. As a result, two partitions, the interleaved partition and the non-interleaved partition, are created on the array. The interleaved partition contains the auto-allocated standard volumes. The non-interleaved partition will initially be empty. Optionally, you may choose not to auto-allocate any volumes. In that case there will be no interleaved partition, the whole array capacity will be used for the non-interleaved partition.

The auto-allocate option is available at RAID array configuration time, either when the array is configured using the ESS Specialist, or when configured by the IBM Systems Support Representative (IBM SSR) using the Batch Configuration tool (see 8.3, “Configuration process” on page 114).

The interleaved structure has been designed to equalize the slight performance differences that logical volumes residing towards the edges of physical disk drives in the RAID array have due to longer average seek times. Interleaved logical volumes are divided into smaller stripes which are distributed over the array to provide the volumes with uniform performance characteristics. We recommend that you auto-allocate standard volumes, when possible.

The interleaved partition characteristics are described in the IBM Redbook *IBM TotalStorage Enterprise Storage Server,* SG24-5465.

You cannot auto-allocate standard volumes on a 72.8GB disk array. The capacity of a 72.8GB array is so large that it would not be feasible to allocate small standard volumes on it because that would quickly exhaust the 256 device address limit of the LCU. Especially in multi array LCUs you would not have enough room for alias devices to provide performance. (Refer to Table 8-4 on page 111 for examples.) Large arrays are more suitable for large volumes with PAV aliases. We recommend that you install the software support for Large Volumes and allocate large custom volumes on 72.8GB arrays.

Here is a summary of RAID array attributes:

- A RAID array is configured as either 6+P+S, or 7+P.
- The first two disk groups of a given DDM size on an SSA loop have a spare disk drive, and the RAID configuration is 6+P+S. This ensures that two spare disk drives of each DDM size are available on the loop.
- The third and subsequent arrays of a given DDM size on an SSA loop are 7+P.
- A spare disk drive can be used for any RAID array of the same DDM size on that loop, either CKD or FB.

8.1.5 CKD logical volumes

CKD logical volumes are mapped into a logical device map in a CKD LSS. The logical devices represent the device address of the logical volume. It ranges from 0x00 to 0xFF. Because each LSS is seen as a logical control unit, the zSeries systems will see it as a 3990-x with up to 256 devices. The logical devices need not be mapped to the ESCON or FICON host adapters, because the zSeries hosts have access to all the logical devices through any of the ESCON or FICON connections available. The set of logical devices accessed by any zSeries image is defined with the Hardware Configuration Definition (HCD) in the IODF file that the operating system uses to recognize its hardware topology.

Logical volumes are allocated on ranks, either RAID ranks, or JBOD ranks. Logical volumes cannot span multiple ranks. When a rank is configured, its track format is set as either 3390 or 3380. All logical volumes on the rank will be of the same track format. It is not possible to intermix different track format volumes on a rank. However, different ranks in an LCU may have different track formats.

An ESS *standard volume* is defined as a logical volume that is created by auto-allocation when an array is configured, either during the initial Batch Configuration process, or using the *Configure Disk Groups* panel on the ESS Specialist. Standard volumes reside in the interleaved partition of a RAID array, and are also referred to as interleaved volumes. You can select to allocate one of the following standard volume types:

- 3390-2, 3390-3, or 3390-9 in 3390 track format
- 3390-2 or 3390-3 in 3380 track format

The number of standard volumes in an array is predefined, and depends on the volume type, DDM size, and the array configuration. Refer to Table 8-1.

Custom volumes are defined using the *Add Volumes* panel on the ESS Specialist. Custom volumes reside in the non-interleaved partition of a RAID array, or on a JBOD rank. As the name suggests, custom volumes have customized sizes. You can specify any number of cylinders from 1 to 10017 when you define the logical volume. On systems that support large volumes (LVS), you can specify up to 32760 cylinders. The volumes you define with the *Add Volumes* panel are considered custom volumes, even when their size matches that of a standard volume. Custom volumes are also called non-interleaved volumes.

You can define as many custom volumes as you like, within the limits of available capacity, up to the maximum of 256 total devices per LCU.

Custom volumes are reported by the host system as:

- 3390-3 devices, when their size is between 1 to 3339 cylinders
- 3390-9 devices, when their size is more than 3339 cylinders

This implies that 3390-2 and 3390-2 (3380 track format) type devices can only be allocated as interleaved standard devices. 3380 track format custom volumes can have maximum of 3339 cylinders, and are reported as 3390-3 devices.

Table 8-1 summarizes the capacities of different RAID arrays, and the number of logical volumes that fit in the array. The number of volumes is indicated with an entry of format X+Y, where X indicates the number of standard volumes in the interleaved partition, and Y indicates the number of custom volumes of the same size that will fit in the non-interleaved partition (NI-partition). Where standard volumes are not supported, only the number of custom volumes is indicated. The number of standard volumes in each RAID array is fixed. If you do not configure any standard volumes, you have the whole array capacity available for custom volumes. Row LargeVol represents the largest possible custom volume, a 32760 cylinder 3390 volume.

Table 8-1 Array capacities

Logical device type	Logical device physical capacity (GB)	6+P+S array			7+P array		
		Physical capacity of array (GB)	Physical capacity of NI-partition (GB)	Number of logical volumes	Physical capacity of array (GB)	Physical capacity of NI-partition (GB)	Number of logical volumes
9.1 GB DDMs							
3390-2	1.96	53.81	6.77	24 + 3	62.79	7.91	28 + 4
3390-3	2.94	53.81	6.77	16 + 2	62.79	15.75	16 + 5
3390-9	8.82	53.81	18.53	4 + 2	62.79	27.51	4 + 3
3390-2 (3380)	1.82	53.81	10.13	24 + 3	62.79	4.55	32 + 4
3390-3 (3380)	2.73	53.81	10.13	16 + 3	62.79	8.19	20 + 5
LargeVol	28.86	53.81	53.81	1	62.79	62.79	2
18.2 GB DDMs							
3390-2	1.96	107.67	5.75	52 + 2	125.62	8.02	60 + 4
3390-3	2.94	107.67	13.59	32 + 4	125.62	8.02	40 + 2
3390-9	8.82	107.67	37.11	8 + 4	125.62	19.78	12 + 2
3390-2 (3380)	1.82	107.67	5.75	56 + 3	125.62	9.14	64 + 4
3390-3 (3380)	2.73	107.67	9.39	36 + 3	125.62	5.50	44 + 2
LargeVol	28.86	107.67	107.67	3	125.62	125.62	4
36.4 GB DDMs							
3390-2	1.96	215.38	11.54	104 + 5	251.28	8.24	124 + 4
3390-3	2.94	215.38	15.46	68 + 5	251.28	16.08	80 + 5
3390-9	8.82	215.38	38.98	20 + 4	251.28	39.60	24 + 4
3390-2 (3380)	1.82	215.38	4.26	116 + 2	251.28	11.04	132 + 6
3390-3 (3380)	2.73	215.38	7.90	76 + 2	251.28	11.04	88 + 4
LargeVol	28.86	215.38	215.38	7	251.28	251.28	8
72.8 GB DDMs							
3390-2	1.96	430.80	430.80	219	502.56	502.56	256
3390-3	2.94	430.80	430.80	146	502.56	502.56	170
3390-9	8.82	430.80	430.80	48	502.56	502.56	56
3390-2 (3380)	1.82	430.80	430.80	236	502.56	502.56	256 (1)
3390-3 (3380)	2.73	430.80	430.80	157	502.56	502.56	184
LargeVol	28.86	430.80	430.80	14	502.56	502.56	17
(1) 256 is maximum allowed. Full array capacity cannot be utilized using the volume type.							

Table 8-2 summarizes the characteristics of CKD logical volumes.

Table 8-2 CKD logical device capacities

Logical device type	**Cylinders**	**Bytes per cylinder**	**Logical device capacity (GB)**	**Physical capacity used (GB)**
		Standard volumes		
3390-2	2226	849960	1.892	1.962
3390-3	3339	849960	2.838	2.943
3390-9	10017	849960	8.514	8.828
3390-2 (3380)	2226	712140	1.585	1.821
3390-3 (3380)	3339	712140	2.377	2.731
		Custom volumes		
3390-3	1-3339	849960	0.00085-2.838	0.00176-2.943
3390-9	3340-32760	849960	2.839-27.844	2.944-28.869
3390-3 (3380)	1-3339	712140	0.00071-2.377	0.00163-2.731

The physical capacity used is the amount of space required on the rank to configure the logical volume. The difference between logical and physical capacity is mainly due to the fact that data on the physical disks is stored in 524 byte sectors of which 512 bytes is available for the logical volume.

You need the physical capacity information to calculate how many logical volumes you can configure on a given ESS subsystem. Use the following formulae to *approximate* the physical capacity of custom volumes. The formula do not completely reflect the algorithms of the ESS as it allocates space for a logical volume.

The amount of physical capacity for 3390 devices can be approximated by:

Physical capacity = (((# Cylinders + 1) * Bytes per Cylinder * 524) / 512) * 1.013 bytes

The amount of physical capacity for 3380 devices can be approximated by:

Physical capacity = (((# Cylinders + 1) * Bytes per Cylinder * 524) / 512) * 1.122 bytes

The equations compensate for any overhead in the logical device such that the result is always greater than or equal to the physical capacity required to configure the logical device.

Table 8-3 summarizes the number of volumes that will fit in the different RAID arrays. The upper section of the table shows the figures for ESS standard volumes. This information is essentially an extract from Table 8-1 on page 108. In the lower section we have included selected custom volumes for your reference. Use the above formula to calculate the figures for other custom volume sizes.

Table 8-3 Number of volumes in a RAID array

Volume type	Vol size (cyl)	9 GB array		18 GB array		36 GB array		73 GB array	
		6+P	7+P	6+P	7+P	6+P	7+P	6+P	7+P
Standard volumes									
3390-2	2226	27	32	54	64	109	128	219	256
3390-3	3339	18	21	36	42	73	85	146	170
3390-9	10017	6	7	12	14	24	28	48	56
3390-2 (3380)	2226	27	36	59	68	118	138	236	(1)
3390-3 (3380)	3339	19	25	39	46	78	92	157	184
Custom volumes									
	20000	3	3	6	7	12	14	24	28
	30000	2	2	4	4	8	9	16	19
	32760	1	2	3	4	7	8	14	17
(1) More than 256 volumes would be required to utilize full array capacity.									

Logical volume characteristics can be summarized as follows:

- All logical volumes on a rank, both standard and custom volumes, have the same track format, either 3390 or 3380. No intermixing is possible.
- It is not possible to intermix CKD and open systems FB volumes on a rank.
- A 3380 track format custom volume can have maximum 3339 cylinders.
- JBOD ranks only contains custom volumes.

8.1.6 Parallel Access Volumes

The Parallel Access Volume (PAV) feature is performance feature that the IBM TotalStorage Enterprise Storage Server brings specifically for the z/OS and OS/390 operating systems. It allows your applications to access logical volumes in parallel. That is, your system can access a single volume from a single host with multiple concurrent requests. This capability represents a significant performance improvement by the ESS over traditional I/O processing for the z/OS and OS/390 environments.

PAV is an optional priced feature of the ESS. In order to use PAVs, the feature must be installed on the ESS. You then have to configure both your ESS and the MVS operating system to use PAVs.

- Use the ESS Specialist to define PAV alias devices for base devices on the ESS. This creates a unit address relationship between the bases and aliases in the ESS hardware, allowing concurrent I/O operations to a single volume by means of the multiple device exposures.
- Use HCD to define the corresponding device types on the host. You can specify in HCD whether PAV bases and PAV aliases should be associated statically or dynamically, resulting in static or dynamic PAVs.

The ESS does not distinguish between static and dynamic PAV - there are only PAVs. Dynamic alias management is a function of the host software. When created, an alias device must be assigned to a base device in the same LCU. WLM can then reassign the alias from one base to another within the same LCU.

See the *IBM TotalStorage Enterprise Storage Server,* SG24-5465 for more details on PAV.

The ability to do multiple I/O requests to the same volume nearly eliminates IOSQ time, one of the major components in response time. Traditionally, access to highly active volumes has involved manual tuning, splitting data across multiple volumes, and more. With PAV and the Workload Manager, your manual performance tuning efforts are minimized.

WLM Dynamic Alias Tuning will allow you to plan the I/O configuration at a coarser level of granularity. Instead of needing to plan the exact relationship of PAV-alias volumes to PAV-base volumes, and the number of PAV-alias volumes for each device, you just need to plan the total size of the PAV-alias pool. This allows for a higher PAV base to PAV alias ratio, allowing you to address more data behind a single ESS subsystem (some installations work well with one dynamic alias for each three 3390-3s, and one alias per each two 3390-9s). Data set placement and hot spot analysis will not be required as the operating system will determine the PAV-alias to PAV-base relationships that best achieve the work load goals defined by you.

To achieve the benefits of PAVs, you need to design your ESS configuration so that there are enough PAV alias devices available. The general rule of thumb is that you should have one alias for a 3390-3 volume (2.8 GB) and three aliases for a 3390-9 volume (8.5 GB). This translates to approximately 3 GB per alias. This applies to static PAVs. With dynamic PAV, you will typically manage with less aliases, but you can keep 3GB per alias as a goal. The lower the ratio is, the better. The larger the ratio gets, the more IOS queuing you can expect.

Larger volumes require more aliases to keep IOS queuing down, but also allow more aliases to be defined for the LCU as there are fewer base devices. Consequently, they give better GB per alias ratios, and typically provide better performance, particularly with dynamic PAV. Table 8-4 shows the maximum number of PAV base and alias devices, and the GB per alias ratio for different volume sizes in three different LCU configurations - an LCU with one 6+P+S array (column 6P), a LCU with one 6+P+S and one 7+P array (column 6P+7P), and finally an LCU with one 6+P+S and two 7+P arrays (column 6P+7P+7P). The GB per alias ratio is the total capacity of logical volumes in the LCU divided by the number of aliases. The LargeVol row represents arrays that are configured with 32760 cylinder volumes.

Table 8-4 PAV alias ratios

Ranks:	**6P**			**6P+7P**			**6P+7P+7P**		
	Base devices	**Alias devices**	**GB per alias**	**Base devices**	**Alias devices**	**GB per alias**	**Base devices**	**Alias devices**	**GB per alias**
36.4 GB DDMs									
3390-3	73	183	1.1	158	98	4.6	243	13	53.0
3390-9	24	232	0.9	52	204	2.2	80	176	3.9
LargeVol	8	248	0.9	17	239	2.0	26	230	3.1
72.8 GB DDMs									
3390-3	146	110	3.8	256	0	(1)	256	0	(1)
3390-9	48	208	2.0	104	152	5.8	160	96	14.2
LargeVol	14	242	1.6	31	225	3.8	48	208	6.4
(1) More than the maximum 256 volumes would be required to utilize full LCU capacity									

Notice how larger volumes provide better GB per alias ratios. Notice also how in large capacity configurations the 256 device limit per LCU becomes a constraint with smaller volume sizes - not enough addresses are available for aliases. Larger volumes will relieve the constraint, allow more PAV aliases, and consequently may provide better performance, particularly with dynamic PAV. For 73GB arrays especially we recommend large volumes.

8.2 Designing the LCU address layout

In this section we give guidance on how to design your Logical Control Unit (LCU) address layout, that is how to divide the LCU address range between base and alias devices. You must define the configuration on the host system IODF and in the ESS and the two must comply with each other, so it is important that you plan before starting the actual implementation. All LCUs on an ESS will not necessarily be identical, the layout depends on the capacity of each LCU and the type and number of volumes in each LCU.

The maximum number of devices in an LCU is 256. Each device has a unit address in the range of 0x00-FF. This applies to both base devices and PAV alias devices. The devices are assigned an address when they are configured on the ESS. The UNITADD parameter of the CNTLUNIT macro can be defined with a 128 or 256 range value (Note: for planning purposes consider that if you later have to increase the unitaddress range of the control unit, this will be disruptive for the HCD).

When a logical volume is defined on the ESS, it is assigned the lowest available address in the LCU address range. Thus, the base addresses typically occupy the low end of the LCU address range (from 00 upwards). When alias devices are defined, they are assigned addresses downwards from what is called the *Starting PAV* address (see 8.4.6, “Configuring logical control units” on page 124). You can select the *Starting PAV* address when you define the LCU with the ESS Specialist. The options are 0, 63 (0x3F), 127 (0x7F) and 255 (0xFF). If all addresses below the *Starting* PAV address are already reserved, then the alias is assigned the lowest available address above the starting address. (This is always the case when *Starting* PAV address is 0.) Thus, the alias devices typically occupy a contiguous address range above the base device address range.

The S/390 Storage Requirements table in the *IBM TotalStorage Enterprise Storage Server Configuration Planner,* SC26-7353 is useful for documenting your configuration. You may want to add in the table rows to tie the ESS configuration with your IODF. See Table 8-10 on page 140 for an example.

Layout 1: All 256 addresses are used

When PAVs are configured, you need to decide how many PAV aliases you can and want to have, and where in the LCU address range they should be located. The base devices (logical volumes) must be configured first and will occupy the low end of the address range as in the first example. A straightforward approach is to use for aliases all the addresses left over from the bases. In terms of performance this is a good solution as it will minimize IOS queuing times. Assume the same configuration as in the previous example, but add the maximum number of PAVs. We will have 52 bases and 204 aliases. See Table 8-5.

Table 8-5 Device address layout example 2 - 256 devices

Host I/O configuration	ESS unit addresses
CNTLUNIT UNIT=2105,UNITADD=((00,256)),... IODEVICE UNIT=3390B,ADDRESS=(9000,52),UNITADD=00,... IODEVICE UNIT=3390A,ADDRESS=(9034,204),UNITADD=34,...	 0x00-33 Base (52) 0x34-FF Alias (204)

When defining this configuration using the ESS Specialist, set Starting PAV Address for the LCU to 255 (0xFF). Then define four aliases for 48 bases and three aliases for the remaining four bases. Starting PAV Address of 0 would give you essentially the same result. The addresses would be assigned in different order, but would occupy the same range.

PAVs are supported when the ESS has a PAV feature installed, and the host runs in Exploitation mode.

Layout 2: Plan for future growth

So why would you want to define less than the maximum number of aliases? Each alias device is assigned a UCB which requires some memory, but since the alias UCBs are located above 16MB in virtual memory, that most often should not be an issue. One possible reason could be, that you want to design your address layout in anticipation of future growth. For example, assume you expect to upgrade the example subsystem to full 11TB capacity in the future, and want to define your IODF already at this point so that you don't have to change it any more when the upgrade takes place. In a full configuration each LCU will hold 80 full size 3390-9 volumes which leaves 176 addresses for the aliases. See Table 8-6.

Table 8-6 Device address layout example 3 - non-contiguous address ranges

Host I/O configuration	ESS unit addresses

`CNTLUNIT UNIT=2105,UNITADD=((00,256)),...`
`IODEVICE UNIT=3390B,ADDRESS=(9000,80),UNITADD=00,...`
`IODEVICE UNIT=3390A,ADDRESS=(9050,176),UNITADD=50,...`
	`0x00-33  Base (52)`
	`0x34-4F  Future base (28)`
	`0x50-63  Future alias (20)`
	`0x64-FF  Alias (156)`

On the ESS Specialist, after defining the 52 base devices, select Starting PAV Address 255 (0xFF) for the LCU, then define three aliases for the 52 bases to come up with this configuration. This leaves address range 0x34-63 undefined on the ESS. That's where we would later add 28 new logical volumes and 20 more aliases.

On the host we define the full address range in anticipation of the coming upgrade. Until the upgrade takes place, addresses 0x34-4F appear to the host as non-existing devices, while addresses 0x50-63 appear as unbound alias devices. (Refer to Example 8-1 on page 143 for a sample D M=CHP command output.) The ESS Specialist would let us define alias devices in range 0x34-4F, but we don't want to do that because it would create an address mismatch between the host IODF and the ESS. Addresses 0x34-4F are unusable at this point.

Layout 3: More connectivity

Another reason for defining less aliases is to reduce the LCU size. For example, an ESCON port can address maximum 1024 devices, counting both base and alias devices. If you define your LCUs with 256 addresses, one ESCON port can only access four LCUs. By limiting the LCUs to 128 addresses, each ESCON port could access eight LCUs. By reducing the LCU size you gain more connectivity. In our example of 52 base devices, an LCU size of 128 would leave us with 76 alias addresses. See Table 8-7.

Table 8-7 Device address layout example 4 - 128 devices

Host I/O configuration	ESS unit addresses

`CNTLUNIT UNIT=2105,UNITADD=((00,128)),...`
`IODEVICE UNIT=3390B,ADDRESS=(9000,52),UNITADD=00,...`
`IODEVICE UNIT=3390A,ADDRESS=(9034,76),UNITADD=34,...`
	
	`0x00-33  Base (52)`
	`0x34-7F  Alias (76)`

On the ESS Specialist, select Starting PAV address 0 or 127 (0x7F) for the LCU and define two aliases for 24 bases and one alias for the remaining 28 bases to come up with this configuration.

The 1024 address limitation is removed with FICON. A FICON channel can address up to 16384 devices.

Layout 4: Add aliases to improve performance

For a final example, assume you determine that, for performance reasons, you need two aliases for all the base devices instead of what we had in the previous example. This will make 52 bases and 104 aliases, for a total of 156 addresses. This will reduce connectivity, each ESCON channel now being able to access six LCUs of that size. (Again, this is not the case for FICON.) See Table 8-8.

Table 8-8 Device address layout example 5 - more aliases

Host I/O configuration	ESS unit addresses

`CNTLUNIT UNIT=2105,UNITADD=((00,156)),...`
`IODEVICE UNIT=3390B,ADDRESS=(9000,52),UNITADD=00,...`
`IODEVICE UNIT=3390A,ADDRESS=(9034,104),UNITADD=34,...`
	
	`0x00-33 Base (52)`
	`0x34-9B Alias (104)`

Using the ESS Specialist with Starting PAV Address set to 0 or 127 (0x7F), add one new alias for the 28 bases which in the previous example only had one. The ESS will automatically assign the aliases from 0x80 upwards since all the lower addresses have been used.

8.3 Configuration process

There are two ways of configuring the ESS. The two methods are referred to as the *custom configuration process* and the *standard configuration process*.

Custom configuration is the process of configuring the ESS using the ESS Specialist Web interface. This method allows you to define fully customized configurations. In this redbook, we describe the custom configuration process in detail.

The standard configuration process, also called the batch configuration process, is performed by the IBM System Support Representative (SSR) on the ESS service console using a tool called the ESS Batch Configuration Tool. It is a simplified configuration method, which provides support for a set of common configurations and standard volume types. The SSR may in some cases use the ESS Specialist to complete the standard configuration process.

For the IBM SSR to perform the standard configuration process, you need to provide him with the necessary configuration information in the form of worksheets. See 8.3.2, “Configuration worksheets” on page 116.

The Batch Configuration Tool is designed to perform the initial configuration of the ESS. Subsequent changes to the configuration should be done with the ESS Specialist although the ESS Batch Configuration Tool can be used for previously unallocated (unassigned) 8-packs on an SSA adapter card, for example when capacity is added to the ESS.

The ESS Batch Configuration Tool is only available to the IBM SSR. The ESS Specialist is available for anyone to use.

8.3.1 Selecting the configuration method

The standard configuration process can be used if the target logical configuration consists of the following elements, which are supported by the Batch Configuration Tool:

- Disk groups defined as RAID-5 arrays
- 3990-6 type control unit images
- Standard 3390-3 or 3390-9 volumes
- 3390 track format volumes only

The Batch Configuration Tool creates LCUs, assigns RAID arrays to the LCUs, and auto-allocates standard volumes in the RAID arrays. All the volumes in a LCU will be of the same type, while different LCUs may contain different volume types.

As part of the standard configuration process, the IBM SSR may use the ESS Specialist to complete the configuration with items that the Batch Configuration tool does not support:

- Add in each array full-size custom volumes of the same type as the standard volumes
- Add a selected number of PAV alias addresses to each configured base device

Table 8-9 summarizes the type and number of logical volumes that are supported by the standard configuration process. In a table entry of format X+Y, X is the number of standard volumes allocated on the array, and Y the number custom volumes of the same size that may optionally be allocated.

Table 8-9 Number of volumes supported by the standard configuration process

	9 GB disk array		18 GB disk array		36 GB disk array	
	6+P+S	**7+P**	**6+P+S**	**7+P**	**6+P+S**	**7+P**
3390-3	16+2	16+5	32+4	40+2	68+5	80+5
3390-9	4+2	4+3	8+4	12+2	20+4	24+4

If your target configuration is supported by the Batch Configuration Tool, we recommend that the ESS be configured using the standard configuration process. It saves you time and creates a balanced configuration. If the standard process meets your requirements except for the last two items, the custom volumes and PAVs, (for example, you want different size custom volumes), you can still use the standard process - simply leave the custom volumes and PAVs out of the batch configuration process, and later configure them yourself using the ESS Specialist.

In this redbook we do not describe the Batch Configuration Tool in detail, because this tool is only available to the IBM SSR. All you need to do is to complete configuration worksheets which describe your planned configuration. The IBM SSR will then use the worksheets to configure the ESS.

On the worksheets you essentially specify:

- The number of LCUs to configure.
- The volume type (3390-3 or 3390-9) on each LCU
- Whether to add custom volumes or not
- The number of PAV aliases per base

If the standard configuration process does not meet your configuration requirements, use the custom configuration method instead. Custom configuration is required, for example, if you want to configure:

- Other than 3990-6 type control unit images
- 3390-2 volumes
- 3380 track format volumes
- Different types of volumes on one LCU
- Custom volumes of non-standard size
- Large volumes with more than 10017 cylinders.

You can configure the ESS yourself using the ESS Specialist or have an IBM representative do it for you. In this redbook we look in more detail on how to use the ESS Specialist to create custom configurations.

8.3.2 Configuration worksheets

ESS configuration worksheets are intended to help you plan and document your ESS logical configuration. The worksheets, together with instructions, are documented in the *IBM TotalStorage Enterprise Storage Server Configuration Planner,* SC26-7353 which you can download from the Internet at the ESS Web site:

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

Select `Reference information` at the bottom of the page (see 1.5.3, "ESS documentation Web site" on page 10 for information on how to access the ESS documentation). Then refer to the manual for information on how to fill in the worksheets.

The CKD storage worksheets are:

- Common worksheets:
  - S/390 Storage requirements worksheet
  - S/390 Add custom volumes worksheet

  Use these worksheets to identify your requirements for ESS storage. You should always fill these in prior to the installation. Once you have completed the applicable common worksheets, give them to either your IBM representative or to the help-line personnel. They use the information recorded in the custom worksheets, and a copy of the ESS order form to decide which configuration method to use.
- Worksheets for standard configuration process:
  - S/390 Batch Configuration worksheets for different capacities

  Based on the information on the common worksheets, the IBM representative or IBM help-line will generate Batch Configuration worksheets that are used by the IBM SSR during the installation and configuration of the ESS. These worksheets are only required for the standard configuration process.
- Worksheets for custom configuration — these worksheets help you plan and perform an ESS custom configuration:
  - Define LCUs and PAVs: Use this worksheet to document the LCUs.
  - Define Disk Groups: Use this worksheet to document RAID or non-RAID disk groups and standard volumes.
  - Add Custom Volumes: Optionally use this worksheet to document custom volumes.
  - Configure Fibre Channel Ports: Optionally use this worksheet to document Fibre Channel ports for FICON attachment.

You can also fill your configuration information online. Go to the above ESS documentation Web site and select `ESS Configuration Worksheets.` On the Web page you can fill in essentially the same information as on the common worksheets. Once complete, the information is submitted to IBM Technical Support Help Line for verification.

Optionally, you can request the Help Line to contact you for assistance. The help line will process the worksheet you submitted and return to you via e-mail a completed copy of S/390 Storage requirements table and a completed Batch Configuration worksheet that corresponds to your ESS. Provide the completed worksheets to the IBM SSR for use during the installation of the ESS.

8.3.3 Communication Resources worksheet

The Communication Resources worksheet allows you to define how various communication functions within the ESS should be configured. IBM service personnel use the information from the worksheet to establish these communication links during ESS initial installation. The worksheet includes such configuration items as TCP/IP addresses for the ESS components, call home and remote support options, modem settings, and telephone numbers.

The worksheet can be found in the *IBM TotalStorage Enterprise Storage Server Introduction and Planning Guide,* GC26-7294.

There are two items on the Communication Resources worksheet that are specific to zSeries hosts, namely the level of Service Information Message (SIM) reporting, and the use of Control Unit Initiated Reconfiguration **(**CUIR).

- **Service information messages (SIMs) for S/390:**

  You select the number of times that the ESS sends a SIM to S/390 or zSeries hosts for entry in the Error Recording Data Set (ERDS). You also select the severity of SIMs that are to display a message on the host console. The ESS sends SIMs to host consoles for the following three types of SIMs (default reporting level shown in parenthesis):

  DASD SIM This SIM tracks disk drive module (DDM) failures and problems (default NONE).

  Media SIM This SIM tracks data check problems on the media (default ACUTE).

  Storage server SIM This SIM tracks control-unit, power-control, and other hardware problems (default ACUTE).

  You select for each SIM type the severity level of the SIMs the ESS sends to your console. SIMs that are below the severity level that you select do not display a console message. You also select the number of times that the ESS sends a SIM to the host. The ESS can send a SIM 1-5 times at eight-hour intervals or not at all (0). The severity levels are:

Selection:	**Severity reported:**
acute	A major subsystem resource is disabled. Performance may be severely degraded. System and application outages may have occurred.
serious	An unrecoverable error or a data check with loss of access to data.
moderate	Performance degradation is possible in a heavily loaded environment. A primary subsystem resource is disabled. Significant performance degradation is possible.
service	No system or application performance degradation is expected in any environment.
none	None

  See 11.6.3, “IEA480E SIM message” on page 194 for an example of the SIM message.

- **Allow CUIR to automatically vary paths off and on:**

  Check **Enable** to allow the ESS to initiate the reconfiguration for service. The ESS will request the attached hosts to automatically vary paths offline for service, and back online after the service is complete.

  Check **Disable** (default) to set the paths to the ESS cluster offline for service. The system operator has to manually vary the affected paths offline, and back online.

With CUIR enabled, the ESS can request that an operating system verify that one or more subsystem resources can be taken offline for service. The ESS uses this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation.

See 11.2.2, "Control Unit Initiated Reconfiguration (CUIR)" on page 178.

8.4 ESS custom configuration for CKD storage

In this section we describe ESS Specialist tasks and panels used for performing a custom configuration of the ESS for CKD storage. For detailed information on how to use the ESS Specialist, refer to *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7346.

8.4.1 Tasks

Configuration of the ESS for use with zSeries systems involves the following main tasks which are performed using the ESS Specialist:

1. **Configure logical control units (LCUs)**

   Define the LCUs for your configuration. Select the control unit emulation mode, specify the Subsystem Identifier (SSID), and enable or disable Parallel Access Volumes (PAVs) for each LCU.

2. **Configure disk groups for each LCU**

   Assign one or more disk groups to an LCU. Select the storage type (RAID-5 or non-RAID) and track format (3390 or 3380) of the ranks. Optionally, auto-allocate a group standard volumes on a RAID array.

3. **Add logical volumes for each LCU**

   Allocate logical volumes on the disk ranks you defined previously. The volumes you add are custom volumes, that is, you can specify the size of each volume individually.

4. **Configure PAV alias devices for each LCU**

   If your system supports PAVs, configure PAV aliases for the logical volumes in the LCU. You can configure multiple aliases for each base device.

5. **Configure FICON ports**

   Configure FICON ports for use on zSeries systems. This task is not required for ESCON ports.

During initial installation, tasks 1 to 4 are performed for each of the up to 16 LCUs, generally in this order, although in large configurations you may need to repeat steps 2 and 3 multiple times for an LCU. Task 5 needs to be done only once for the ESS. Later, you can modify the configuration.

8.4.2 ESS Specialist panels for CKD storage allocation

The following ESS Specialist panels are used for configuring CKD storage on the ESS:

Storage Allocation This panel gives you a graphical view of disk rank allocation status. From this panel you enter either the S/390 Storage or the Open System Storage configuration functions. Enter this panel from the Navigation frame.

S/390 Storage Use this panel to display your CKD configuration and to access the following panels and their associated configuration tasks. This is the main CKD storage configuration panel. From this, you enter the specific task panels, which are used to:

Configure LCUs	Configure the LCUs.
Configure Disk Groups	Allocate physical disk capacity for an LCU.
Add Volumes	Define custom volumes for an LCU.
Modify PAV Assignments	Define and delete PAV aliases.
Configure Host Adapter Ports	Configure FICON host adapter ports.

Figure 8-4 is a graphical presentation of the CKD storage configuration panels and the functions you can select on them.

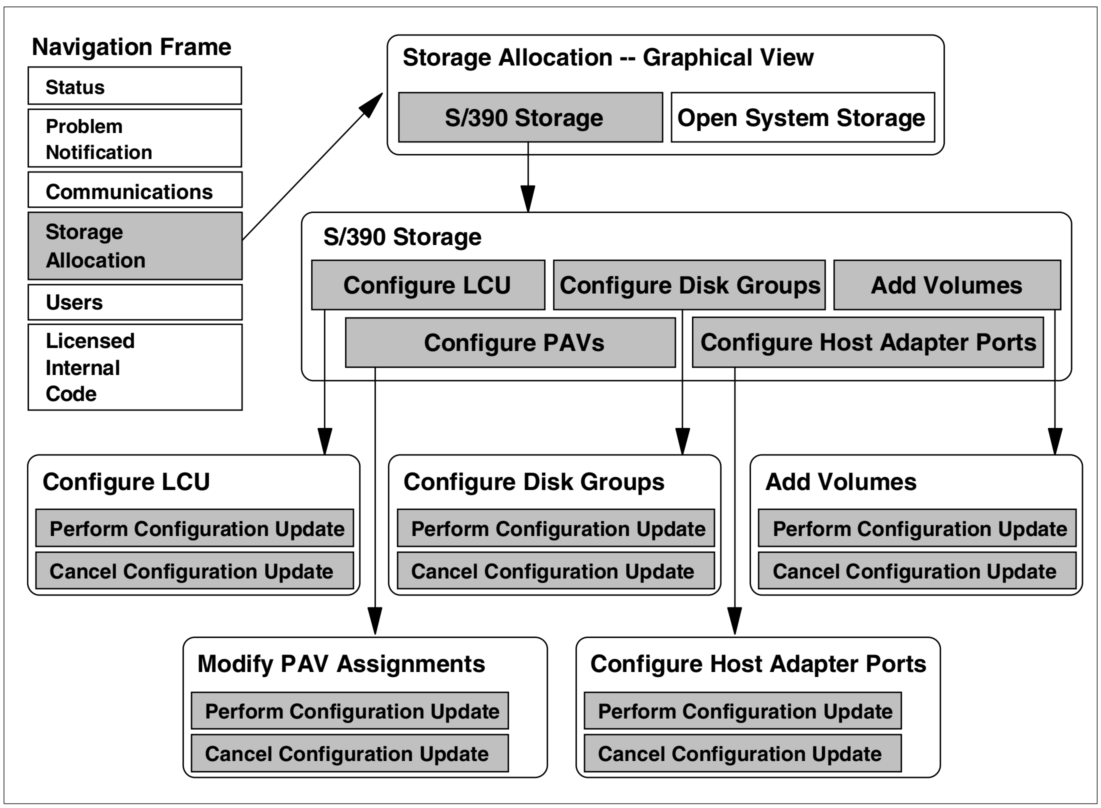

Figure 8-4 ESS Specialist: S/390 panels

The **Navigation frame** is visible on each panel. On most screen images in this chapter we do not show the Navigation frame. It is shown in Figure 8-7 on page 122.

On each of the task panels, you enter information, set attributes, and make selections required to define your configuration. When you have made the necessary definitions, perform the actual configuration action by clicking the **Perform Configuration Update** button, or cancel it by clicking **Cancel Configuration Update**. In either case, once the action is complete, you will return to the *S/390 Storage* panel, where you can select the next task.

To get back to the *Storage Allocation - Graphical View* panel, click the **Storage Allocation** button on the Navigation frame.

8.4.3 ESS Specialist messages

During configuration procedures, the ESS Specialist will present messages in the Message frame at the bottom of the screen. The two most commonly displayed messages are:

- **2108:** Retrieving the machine configuration data ... please wait.
- **1533:** Downlevel 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 will receive this message:

- **2122:** Retrieving the current configuration data ... please wait.

If this is again followed by message 1533, wait a few minutes while the ESS Specialist is synchronizing configuration data in the two clusters, then re-try.

ESS Specialist also displays messages and status information in separate windows. When you perform a configuration action that is expected to take longer than 30 seconds to complete, the ESS Specialist displays a window showing the estimated remaining time for the action (Figure 8-5):

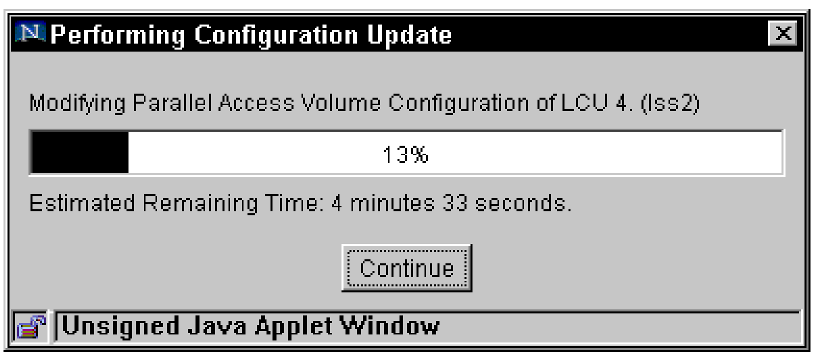

Figure 8-5 ESS Specialist: Estimated time

If you click the **Continue** button, the window closes, and you can continue working with the ESS Specialist. The actual configuration action continues in the background in the ESS. However, you cannot initiate another configuration action until the current one is complete.

Note: The ESS does not allow configuration changes to be performed using the ESS Specialist while a service representative is logged on the service console. If you get error message 1109 (see Figure 8-6) indicating that ESS components are in service mode, this is probably the cause.

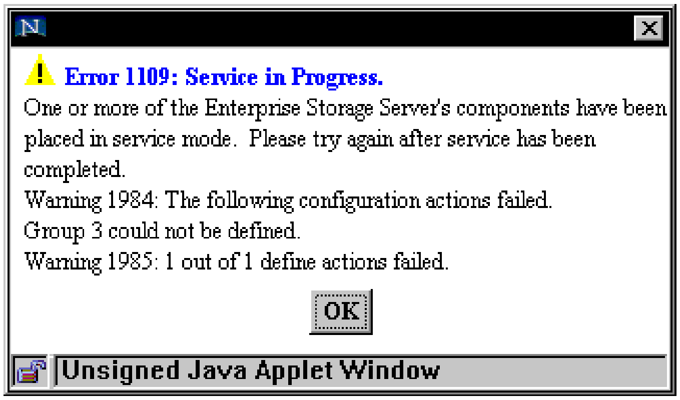

Figure 8-6 ESS Specialist: Error 1109

8.4.4 Storage Allocation panel

Begin ESS configuration by connecting to the ESS Specialist. When connecting, you may be presented with a New Site Certificate window. View the series of certificates, then select ESS Specialist on the Navigation frame. The Username and Password Required panel is displayed. Proceed to enter a user ID with configuration authority and the password. Refer to Chapter 5, “IBM TotalStorage Enterprise Storage Server Specialist” on page 79 for instructions on how to start an ESS Specialist session.

Once connected, the *ESS Specialist Welcome* screen appears on a separate browser window. Select **Storage Allocation** from the Navigation frame to start the storage configuration function.

You should now see the *Storage Allocation - Graphical View* panel (Figure 8-7). It gives you an overall view of the ESS configuration, showing defined host systems, installed host adapters, disk ranks, and a storage allocation summary.

You do not deed to define zSeries hosts on the ESS Specialist. When ESCON or FICON host adapters are installed on the ESS, a pre-defined EsconNet or FiconNet icon automatically appears on the panel. The icons are pseudo-hosts representing all zSeries hosts attached through ESCON or FICON adapters, respectively. The icons appear even if there is no physical connection between the hosts and the ESS.

Unlike with FB storage, you do not need to assign CKD logical volumes to a particular port or host. Every ESCON and FICON port in the ESS is potentially able to access any CKD logical volume (up to 256 per LCU) in any of the defined LCUs (up to 16). Which LCUs and logical volumes a particular host channel will actually access depends on the host I/O configuration.

The rectangles in the middle represent configured disk groups. If you click the EsconNet or FiconNet icon, CKD disk groups display colored and with thick frames, and the ESCON or FICON adapter ports highlight in yellow. This indicates that all installed ESCON/FICON ports are automatically associated with the EsconNet/FiconNet host. The empty rectangles in this case represent FB storage.

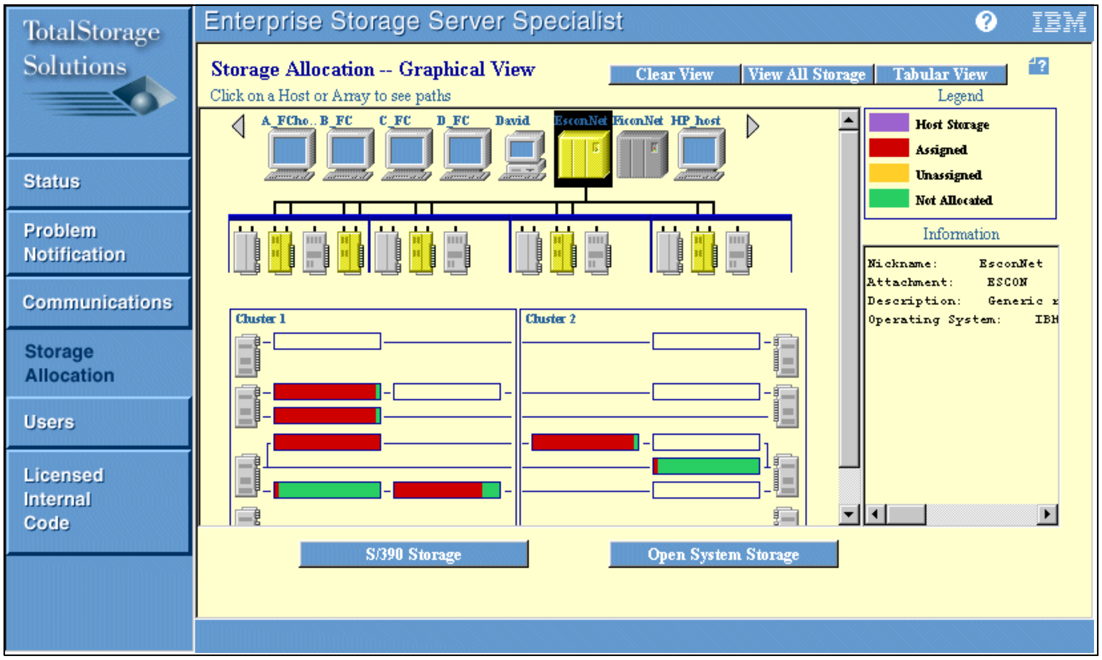

Figure 8-7 ESS Specialist: Storage Allocation - Graphical View panel

On the *Storage Allocation* panel, you can select Tabular View. This displays a table showing all the configured logical volumes.

On the *Storage Allocation* panel, click the **S/390 Storage** button to enter the storage configuration functions for ESCON and FICON attached hosts. The **Open System Storage** button will take you to SCSI and FCP attached Fixed Block storage configuration functions.

8.4.5 S/390 Storage panel

The *S/390 Storage* panel (Figure 8-8) is the main panel for allocating CKD storage. From it, you enter the specific configuration tasks, and once the task is complete, return back to it.

To display the *S/390 Storage* panel, click **S/390 Storage** on the *Storage Allocation - Graphical View* panel.

The *Logical Control Units* table on the panel shows the LCUs for the ESS. Initially, LCUs 000 to 007 are visible, all undefined. The upper LCUs 008 to 00F are not shown until you have assigned a disk group to the corresponding lower LCU on a device adapter.

Figure 8-8 shows a partially configured system, with some LCUs already configured, some still undefined.

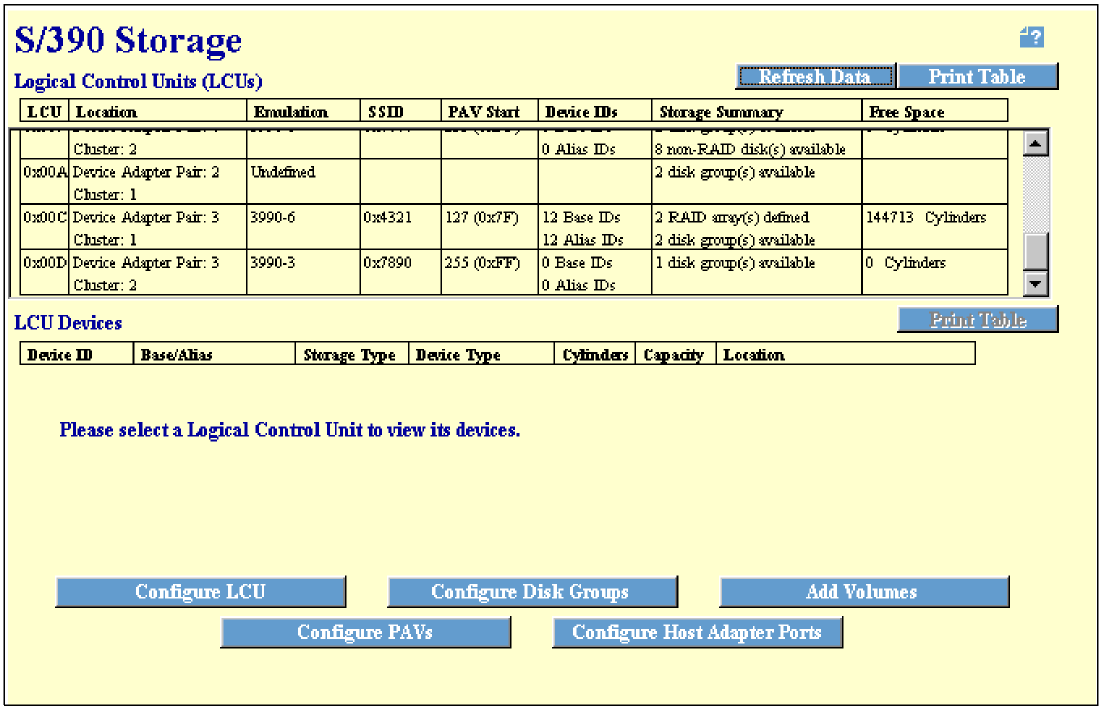

Figure 8-8 ESS Specialist: S/390 Storage panel

When you click a row on the *LCUs* table to select an LCU, the underneath *LCU Devices* table will show the currently defined logical devices for the LCU (Figure 8-9). If no devices have yet been configured for the LCU, the *LCU Devices* table is empty.

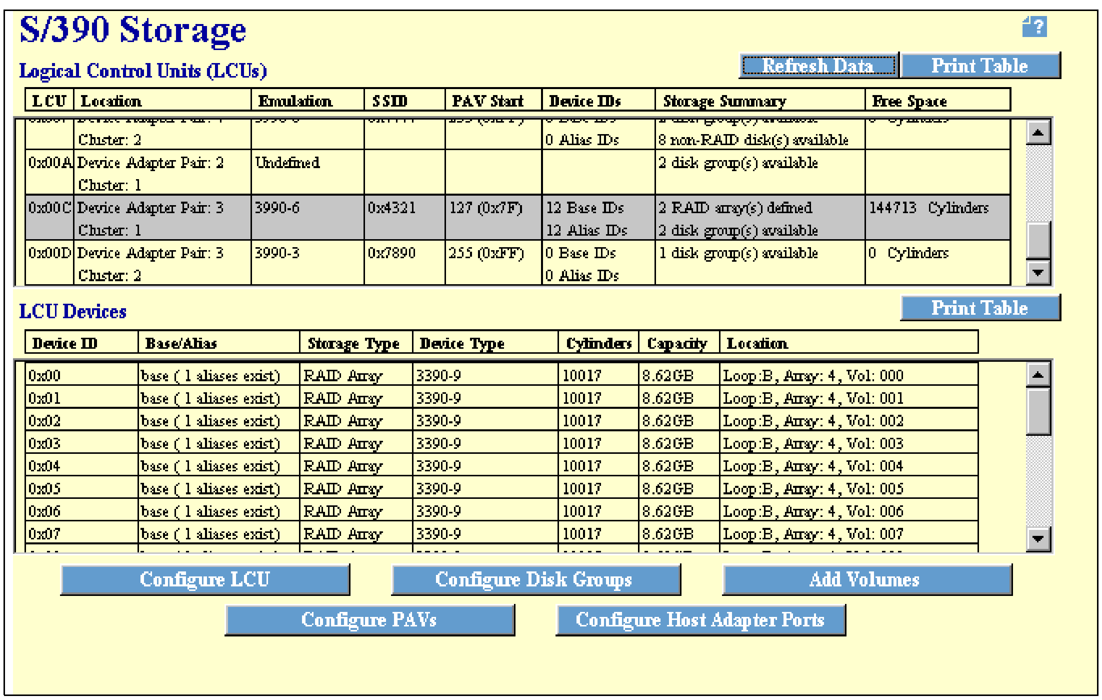

Figure 8-9 ESS Specialist: S/390 Storage panel with devices

In the *Logical Control Units (LCUs)* table:

- The **Location** column shows the cluster and device adapter where the LCU resides. Refer to Figure 8-1 on page 103 for the correspondence between LCUs, device adapters, and clusters.
- The **Storage Summary** column shows how many disk groups have been assigned to the LCU, and how many unassigned groups are still available for the LCU.
- The **Free Space** column shows the number of cylinders that are available in the LCU for allocating new logical volumes.

In the *LCU Devices* table:

- The **Device ID** column shows the unit address of the logical device. Once you have initialized the logical volumes using ICKDSF, it will also display the volume serial number (VOLSER).
- The **Capacity** column shows the size of the logical volume (GB = 1,000,000,000 bytes). The figure is slightly higher than the number of cylinders multiplied by the effective 3390 or 3380 cylinder size that is normally used in capacity calculations. The figure given here includes fields of the CKD format track that do not account for the usable capacity.
- The **Location** column shows the SSA loop and disk rank where the volume is allocated, along with a sequential number within the rank. For alias devices, it shows where the corresponding base device is located.

To start a configuration action, select in the *Logical Control Units (LCUs)* table the LCU that you want to work with. The row will highlight in grey. Then click one of the function buttons in the lower part of the panel. You need to select the LCU first for all other functions except for the Configure Host Adapter Ports function.

If you want to return to the *Storage Allocation* panel from here, click the **Storage Allocation** button in the Navigation frame (not shown in the figure).

8.4.6 Configuring logical control units

Use the *Configure LCU* panel to define and modify Logical Control Units. On the *S/390 Storage* panel, select the LCU you want to work with. It will highlight in grey. Then click **Configure LCU** to display the *Configure LCU* panel for the selected LCU (Figure 8-10).

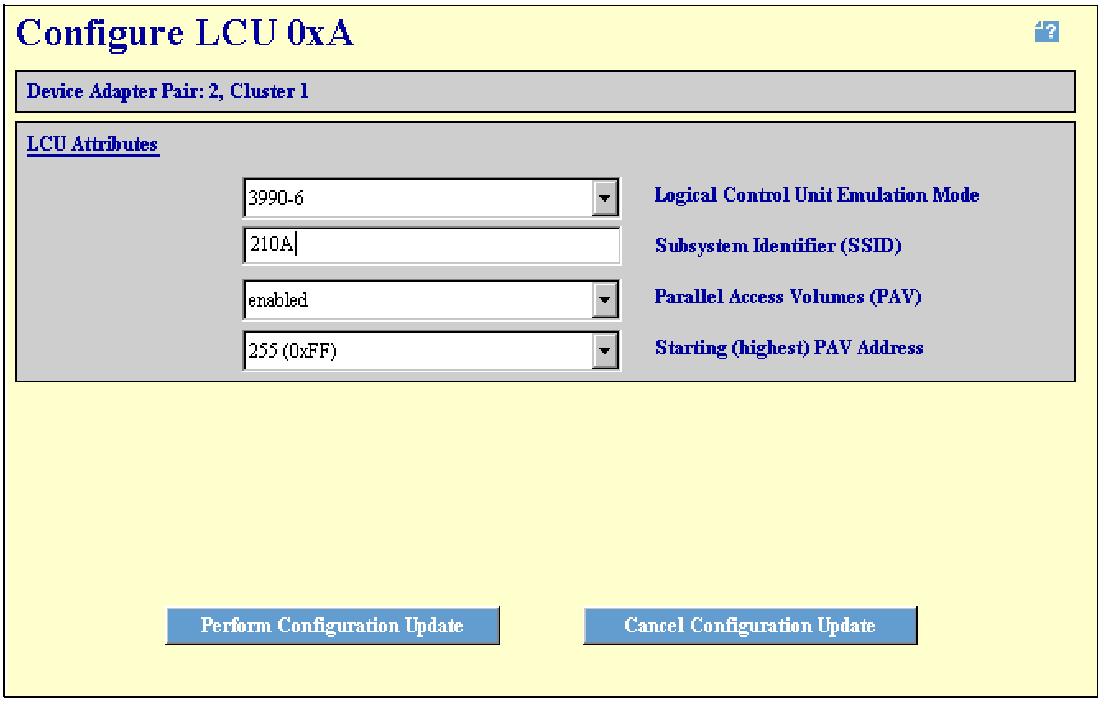

Figure 8-10 ESS Specialist: Configure LCU panel

Defining an LCU

These are the steps to be performed in defining an LCU:

1. Select the **Logical Control Unit Emulation Mode**. The options are:
   - 3990-6
   - 3990-3
   - 3990-3 TPF
   - 3990 (Siemens BS2000 Compatible).

   Select 3990-6 unless you have a specific reason, such as a software requirement, to select one of the other options. This is the only entry field displayed for an undefined LCU. Once you select the mode, other fields appear on the panel.

   Note: the IBM TotalStorage Enterprise Storage Server is qualified for the Operating System BS2000/OSD V4, allowing the integration of the ESS in solutions with Fujitsu-Siemens servers. ESCON is the only attachment supported currently. An RPQ must be submitted for getting IBM approval before attachment.

2. Set the **Subsystem Identifier (SSID)** for the LCU. The SSID is a four digit hexadecimal number in the range of 0x0001 to 0xFEFF (it must not have FF as the first two digits). The SSID must be unique in the location, no other storage subsystem should have the same SSID.

   A practical convention is to use the first two digits to identify the ESS subsystem, and the two last to indicate the logical subsystem (00-0F). This is useful for PPRC, for example, as the commands to establish PPRC paths and pairs take the SSID and the two digit LSS number as parameters.

3. Set the **Parallel Access Volumes (PAV)** attribute to enabled or disabled. Select enabled if your system supports Parallel Access Volumes. If the PAV feature is not installed, you can only select disabled. If you add the PAV feature to the ESS later, you can change the PAV attribute from disabled to enabled.

   You must enable PAV here in order to be able to configure PAV aliases for the LCU later.

4. If you enabled PAV for the LCU, set the **Starting (highest) PAV Address**. The options are:

   – 0 (0x00)
   – 63 (0x3F)
   – 127 (0x7F)
   – 255 (0xFF).

   The Starting PAV address controls where in the LCU address range the PAV alias devices will be assigned when you configure them using the Configure PAVs function (discussed later in 8.4.9, “Modify PAV assignments” on page 137).

   Select starting address 0 to allocate the alias devices contiguously right after the base devices. Select one of the other options only if you want to create an address layout where the base and alias device address ranges are not contiguous.

   PAVs are assigned the highest unused address available, beginning at the PAV starting address, and proceeding in decreasing address order. Once all the addresses below the Starting PAV address have been used, then addresses are assigned in ascending order starting from the lowest unused address upwards, up to the limit of 256 addresses that are available in an LCU. With the PAV starting address set to less than the total number of bases and PAV aliases you will be defining, this results in a contiguous range of device addresses for the LCU. Starting PAV address 0 always gives you a contiguous range.

   You may refer to section 8.2, “Designing the LCU address layout” on page 112 for a detailed discussion on PAV addresses.

5. When all entries are complete, click **Perform Configuration Update**, to apply the configuration step for the LCU.

You do not have to define all the LCUs that are available. In smaller configurations you may only need four or eight LCUs. If the storage capacity on some DAs is intended only for open systems use, you do not need to configure the LCUs that correspond to those DAs.

Figure 8-11 shows an example of how you would configure an LCU.

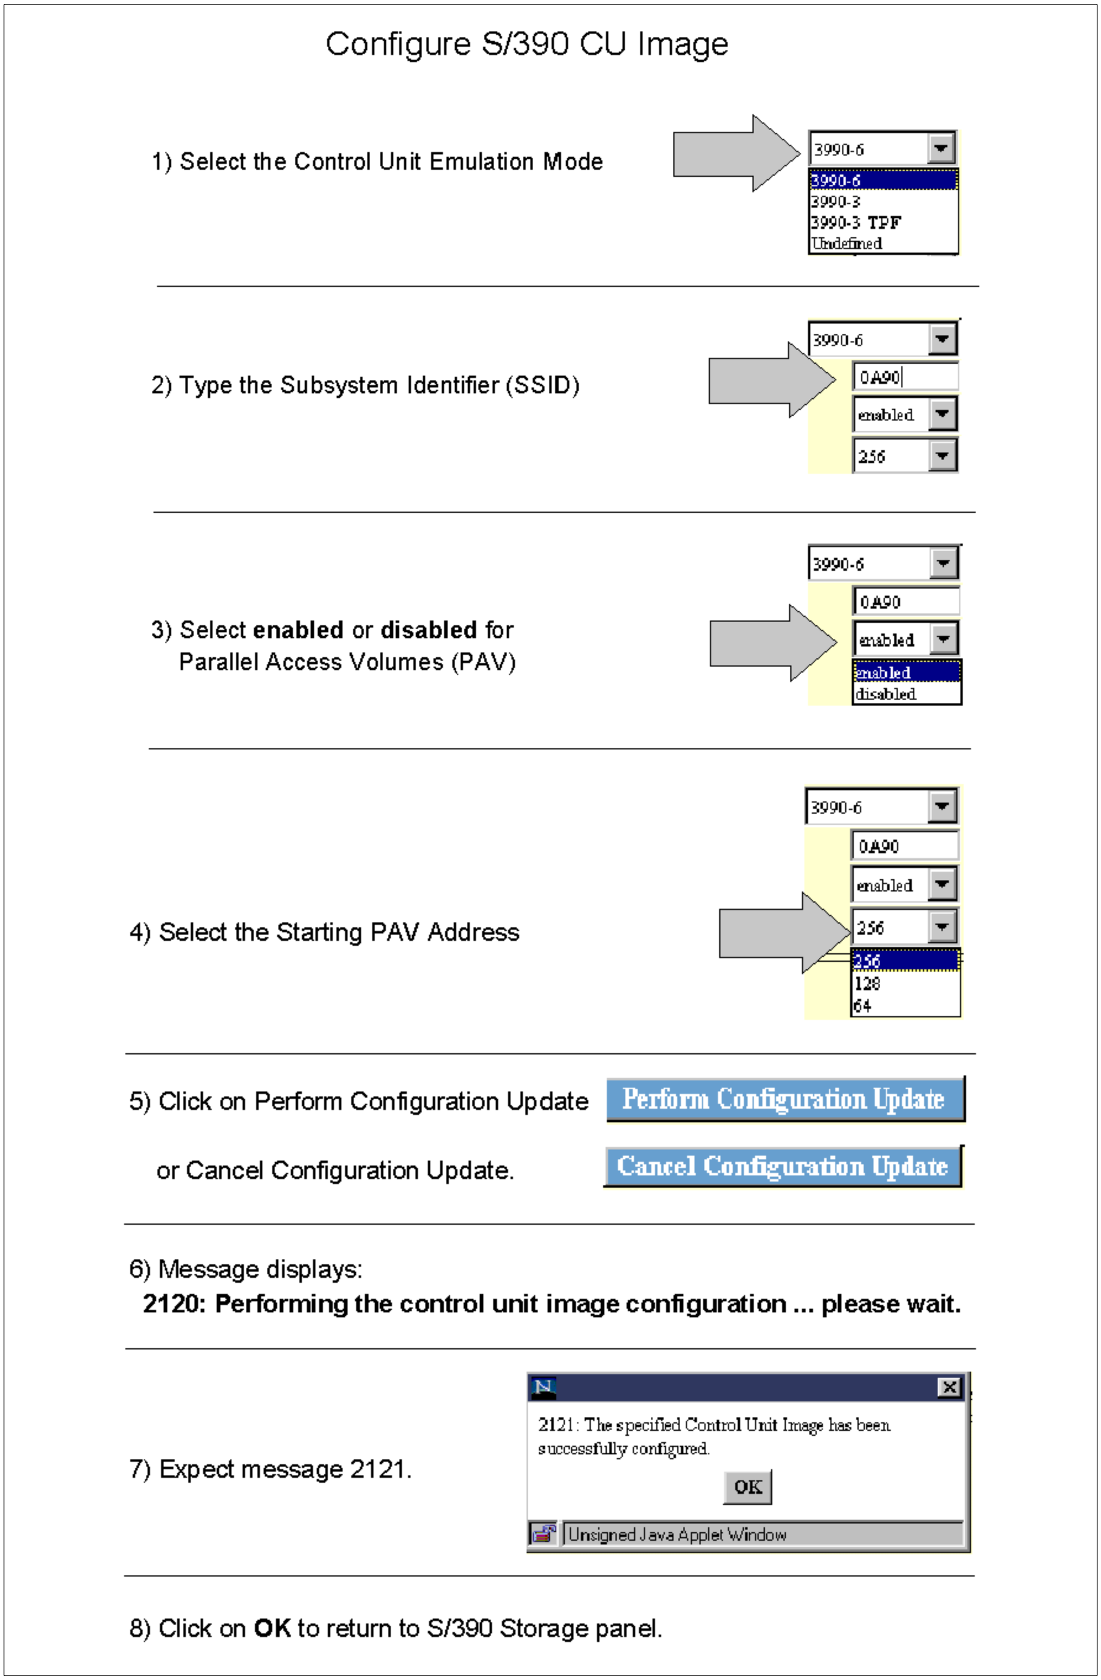

Figure 8-11 ESS Specialist: Procedure to configure an LCU

Note: The screen images on this figure may slightly differ from those that you see when using a later level ESS Specialist.

After successful configuration of an LCU, you will return to the *S/390 Storage* panel. The *Logical Control Units (LCUs)* table now shows the LCU as Defined, with the attributes you selected. The Free Space is zero, until you configure disk groups for the LCU.

Modifying an LCU

You can change the following LCU attributes:

- You can change the Starting PAV Address for an LCU, at any time.
- You can change the Parallel Access Volumes attribute from disabled to enabled for a configured LCU, but you cannot change it back to disabled. To disable PAV, you must first undefine, then redefine the LCU.
- You can change the Subsystem Identifier. However, you must first terminate host functions that depend on the SSID. Also all drives must be varied offline before changing the SSID so software can correctly rebuild their control blocks. For example, PPRC path and pair definitions use the SSID to identify the source and target LCUs. In order to change an SSID, you must first delete all PPRC pairs and paths defined for the LCU. Figure 8-12 shows the warning message you will get when trying to change the SSID.

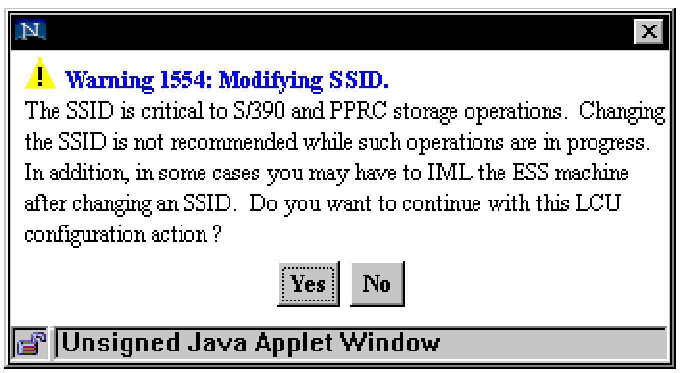

Figure 8-12 ESS Specialist: Warning 1554

You cannot change the Control Unit emulation mode once the LCU has been configured. To change the emulation mode you would have to first undefine the LCU.

Undefining an LCU

You can undefine an existing configured LCU by setting its Logical Control Unit Emulation Mode to **Undefined** on the *Configure LCU* panel. When you click **Perform Configuration Update**, the warning message shown in Figure 8-13 is displayed.

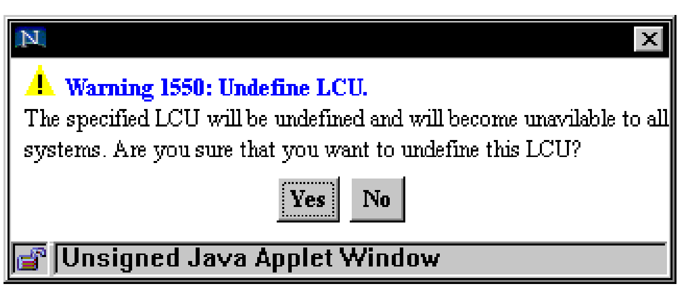

Figure 8-13 ESS Specialist: Warning 1550

Attention: If you click **Yes**, the LCU definition is deleted, with its associated ranks and logical volumes. All data on the logical volumes in that LCU is discarded. Make sure that you have backed up all data on the volumes before you delete an LCU.

You would need to undefine an LCU in order to change its emulation mode.

8.4.7 Configure disk groups for an LCU

Use the *Configure Disk Groups* panel to assign disk groups (that is, physical storage capacity) to logical control units, and to set the disk group attributes. On the LCU table on the *S/390 Storage* panel, select the LCU you want to add disk groups to. The Storage Summary column of the table must indicate that disk groups are available. The LCU row will highlight in grey. Click **Configure Disk Groups**. Panel shown in Figure 8-14 will appear.

You cannot add disk groups to an undefined LCU. You must first configure the LCU.

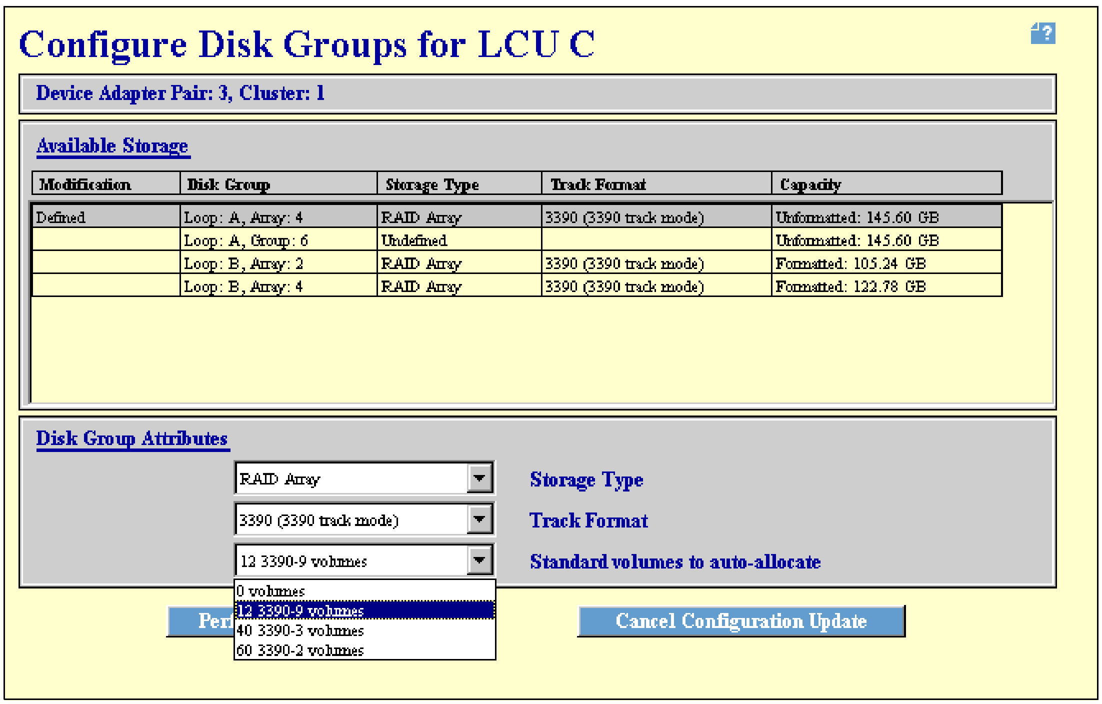

Figure 8-14 ESS Specialist: Configure Disk Groups panel

Defining RAID Storage

1. Select an undefined disk group to configure by clicking on the corresponding row in the *Available Storage* table. The row highlights in grey. You can select a disk group either in loop A, or loop B.

   The *even* numbered LCUs (which are located in Cluster 1) are associated with the *even* numbered disk groups (2,4,6) on an SSA loop. The *odd* numbered LCUs (in Cluster 2) manage the *odd* numbered disk groups (1,3,5).

2. Select the **Storage Type** in the *Disk Group Attributes* section. The options are:

   – RAID Array
   – non-RAID
   – Undefined (to undefine an previously configured group).

   Once you change an attribute, the selected row in the *Available Storage* table will be updated. The Modification column now shows Defined indicating that you have changed a disk group attribute.

3. Select the **Track Format** for the array. The options are:

   – 3390 (3380 track mode)
   – 3390 (3390 track mode)

   All logical volumes on the rank will be of the track format you select here.

4. In the **Standard volumes to auto-allocate** field, select the type of standard volumes you want to auto-allocate on the array. You cannot select the number of volumes, just the type. The number is predefined, and depends on the array capacity, track format, and volume type. The example pull-down in Figure 8-14 shows the options for an 18GB 7+P array. See Table 8-1 on page 108 for a summary of the number of standard volumes in each array.

   When you auto-allocate volumes, two partitions are created in the array. The interleaved partition contains the auto-allocated standard volumes and takes up most of the capacity in the array. The second, non-interleaved partition will initially be empty. You can later add custom volumes in it using the *Add Volumes* panel.

   If you select **0 volumes**, no standard volumes will be auto-allocated, just the track format of the array will be set. An interleaved partition is not created, and the non-interleaved partition will cover the whole array. This option allows you to fill the array with custom volumes.

5. If you have more than one disk group to configure, you may repeat steps 1 to 4.
6. Click **Perform Configuration Update** to apply the configuration changes you defined.

Figure 8-15 shows an example of how you would configure a disk group.

Configure Disk Groups for CU Image - RAID 5

1) Select RAID Array.

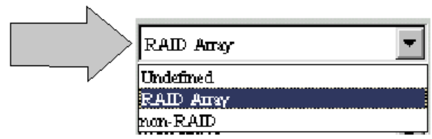

2) Select 3380 or 3390 track mode.

3) Select Volume Type:
3390 Model 2, 3, or 9
to auto-allocate interleaved volumes.
OR
Select **0 volumes** to create
a non-interleaved RAID partition.

4) Click on Perform Configuration Update

or Cancel Configuration Update.

5) Message displays: **2201: Performing disk group configuration ... please wait.**

6) Expect message 2202.

7) Click on **OK** to return to S/390 Storage panel.

Figure 8-15 ESS Specialist: Procedure to configure disk groups

Note: The screen images on this figure may slightly differ from those that you see when using a later level ESS Specialist.

The volumes auto-allocated in each array will not appear on the *LCU Devices* table until formatting of the disk group is complete. This may take more than an hour depending on the array size. Several disk groups can be formatting simultaneously.

After successful configuration of a disk group, the *S/390 Storage* panel appears as in Figure 8-9 on page 123. The auto-allocated volumes now appear in the LCU Devices table. The Storage Summary and Free Space columns of the *Logical Control Units (LCUs)* table have been updated to reflect the new configuration.

Defining non-RAID storage

1. Select a disk group in the *Available Storage* table (Figure 8-14).
2. Select **non-RAID** as the **Storage Type**. When you select non-RAID, the row in the *Available Storage* table expands to eight rows, one for each of the drives in the disk group. The Storage Type field in the table is non-RAID. The Capacity field has the unformatted capacity of a single disk drive.
3. The rows for the non-RAID disks can be individually selected in the *Available Storage* table. Select a disk to configure.
4. Select **3390** or **3380** for the **Track Format**.

   The Standard volumes to auto-allocate attribute is set to zero. You cannot change it. No volumes can be auto-allocated on a non-RAID rank. You need to allocate the volumes later using the Add Volumes panel.
5. Repeat steps 3 and 4 for the other non-RAID disks.
6. Click **Perform Configuration Update**.

A non-RAID disk group may contain both CKD and FB format disks in any combination. If a disk group is defined as non-RAID using this panel, then any disks in the group that are not formatted as CKD (3390 or 3380 track format) are available to be formatted in FB track format using the Open Systems Storage function. Similarly, if a disk group is defined as non-RAID using the Open Systems Storage function, any disks in the group not formatted as FB appear in the *Configure Disk Groups* panel for the LCU and are available to be formatted as CKD.

Undefining a RAID array

You can undefine an previously configured array. This will delete its contents and make the disk group eligible for configuration with a different logical volume setup, or as non-RAID storage.

1. Select the array in the *Available Storage* table.
2. Set the Storage Type attribute to **Undefined**. If there are volumes allocated on the group, warning message in Figure 8-16 is displayed.

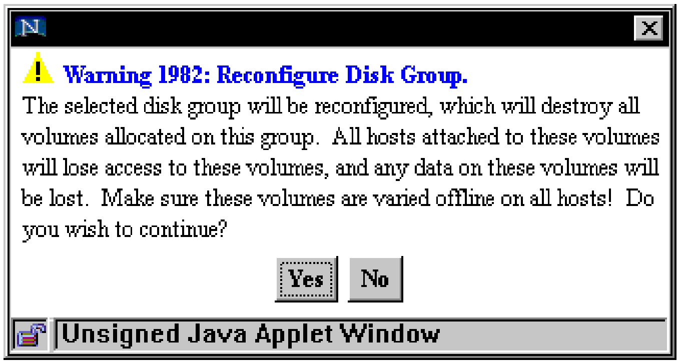

Figure 8-16 ESS Specialist: Warning 1982

When you click **Yes**, the **Storage Type** changes to **Undefined**, the **Track Format** changes to **None** (unused disk), and the Modification column in the *Available Storage* table will show **Undefined**. At this point you can still cancel the change. The actual reconfiguration is not applied until you perform the configuration update.

3. If you want to undefine more arrays in the LCU, repeat steps 1 and 2.
4. Click **Perform Configuration Update**.

Attention: When you click **Perform Configuration Update**, the disk group will be reconfigured. It is deleted from the LCU, with all its associated logical volumes. Any data on those volumes is discarded. Make sure you have backed up all data on the volumes before you delete an disk group.

The reconfiguration is a time consuming action. The physical disks in the array are reformatted and this may take 30 to 60 minutes, even longer to complete. The disk group will not appear on the Configure Disk Groups panel until the background formatting process has completed. The progress of the reformatting process is listed with a % complete indication on the Information window on the ESS Specialist Status panel.

Undefining non-RAID ranks

You undefine an individual non-RAID rank by changing its Track Format to **None (unused disk)** on the Configure Disk Groups panel. Once complete, the disk is again available to be formatted as either CKD or FB rank.

Attention: When you change the track format of a non-RAID disk to None (unused disk), all data on the disk will be lost. Make sure you have appropriate backups.

On this panel you can only undefine CKD format non-RAID ranks. Use the Open System Storage Configure Disk Groups panel to undefine FB format ranks.

In order to undefine a non-RAID group, you first have to undefine all individual disks, both CKD and FB format disks, in the group by changing their track format to None (unused disk). Once you have undefined the individual disks in the group, you can proceed to undefine the non-RAID group itself. The procedure is the same as for undefining RAID arrays. You can undefine a non-RAID group using either this panel or the corresponding Open Systems Storage Configure Disk Groups panel.

8.4.8 Add volumes to an LCU

To allocate custom volumes, select an LCU on the *S/390 Storage panel*, then click the **Add Volumes** button. You should then see the *Add Volumes* panel (Figure 8-17).

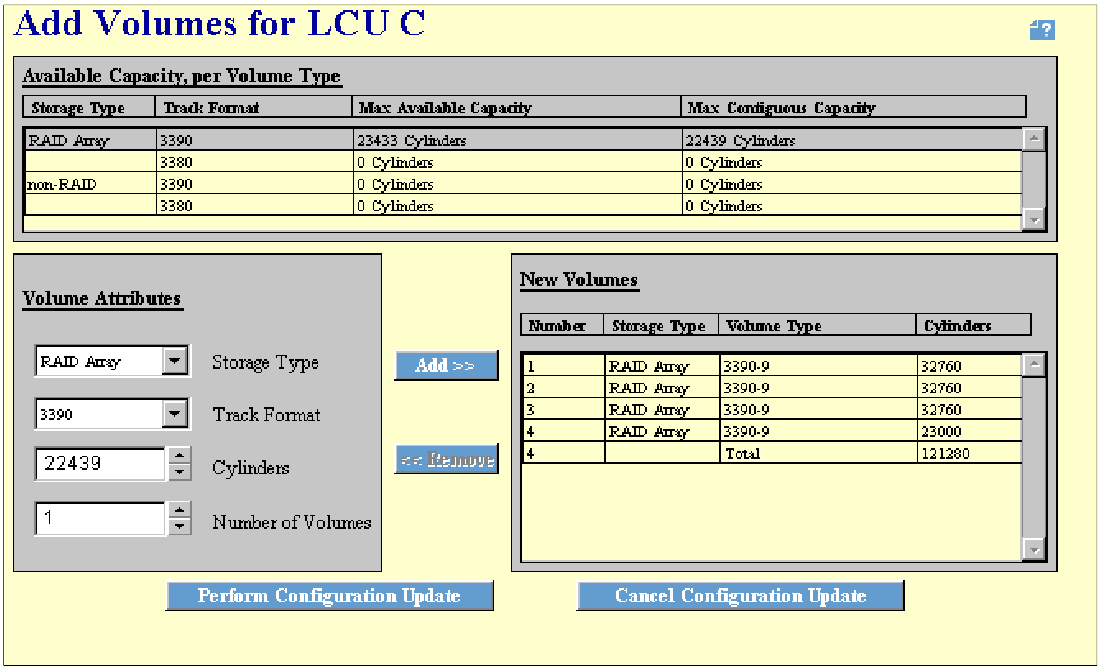

Figure 8-17 ESS Specialist: Add Volumes panel

1. In the *Volume Attributes* section, select the **Storage Type** and **Track Format** of the volumes you want to add. This will automatically select the corresponding row in the Available Capacity, per Volume Type table. Alternatively, select a row in the Av*ailable Capacity, per Volume Type* table that corresponds to the type of volume you want to add. This will automatically set the **Storage Type** and **Track Format** fields in the *Volume Attributes* section.

   The *Available Capacity, per Volume Type* table shows the capacity that is available for adding different types of volumes in the LCU, classified into four categories:

   – RAID Array in 3390 track format
   – RAID Array in 3380 track format
   – Non-RAID in 3390 track format
   – Non-RAID in 3380 track format

   The total capacity of the volumes you add cannot exceed the available capacity. If Max Available Capacity is 0 cylinders, you cannot add any volumes of that type.

   If you auto-allocated 3390 track format volumes when you configured the disk group using the *Configure Disk Groups* 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 CKD RAID array is formatted entirely in a single track format, either 3390, or 3380.

2. Set the size of the volumes you want to add by entering the number of cylinders in the **Cylinders** field. The maximum number you can set is 32760 if the system supports Large Volumes, otherwise 10017, or the number of cylinders in the Max Contiguous Capacity field in the *Available Capacity* table, whichever is smaller. If you have selected a category where no capacity is available, you cannot add new volumes.

The volumes you add here are custom volumes, that is, you can specify any number of cylinders from 1 to 10017 (or 1 to 32760 if Large Volume Support is installed). If you want to emulate the standard capacities for real 3390 volumes, specify:

1113	for 3390 Model 1
2226	for 3390 Model 2
3339	for 3390 Model 3
10017	for 3390 Model 9

For real 3380 track format volumes, the number of cylinders per volume are:

885	for 3380-STD, 3380-D, or 3380-J
1770	for 3380-E
2655	for 3380-K

3. Set the number of volumes you want to add in the **Number of Volumes** field.

   The available space for the category may be distributed over multiple ranks. Because a logical volume cannot span ranks, the number of volumes of a given size you can add may be less than the Max Available Capacity divided by the size of the volume. The number is dictated by the size of the contiguous segments on the ranks. The remaining space can be used for allocating smaller volumes. You can add volumes of different sizes in one category of available capacity, until no usable space remains.

4. When you have defined the attributes of the volumes, click the **ADD>>** button to add the volumes to the *New Volumes* list. You can repeat this several times, with different volume sizes and numbers.

   If the total capacity of volumes you want to add is too large for the available space, the ESS Specialist displays error message 1901 (Figure 8-18).

Figure 8-18 ESS Specialist: Error 1901

   The volumes that fit in the available space are added to the *New Volumes* list.

   The Max Available Capacity and Max Contiguous Capacity fields are updated. The number of cylinders subtracted from the available capacity may slightly exceed the total number of cylinders in the volumes added. This is because of overheads associated with each volume, for example, metadata. The Cylinders attribute field is updated to show the largest single volume that can be added in the remaining space.

5. Prior to performing the configuration update, you can remove a volume that has previously been added to the *New Volumes* list by selecting its row in the table and clicking the **<<Remove** button.

6. When you have defined the desired volumes, click **Perform Configuration Update** to apply the change. Volumes on the *New Volumes* list will be configured.

Attention: After creation, the logical volumes cannot be individually deleted. The only way to delete a logical volume, or to change its type is to undefine the whole disk group where the volume resides. When a disk group is undefined using the *Configure Disk Groups* panel, all the logical volumes on the group are deleted, and any data on them is discarded.

Figure 8-19 shows an example of how you would add volumes to a disk group.

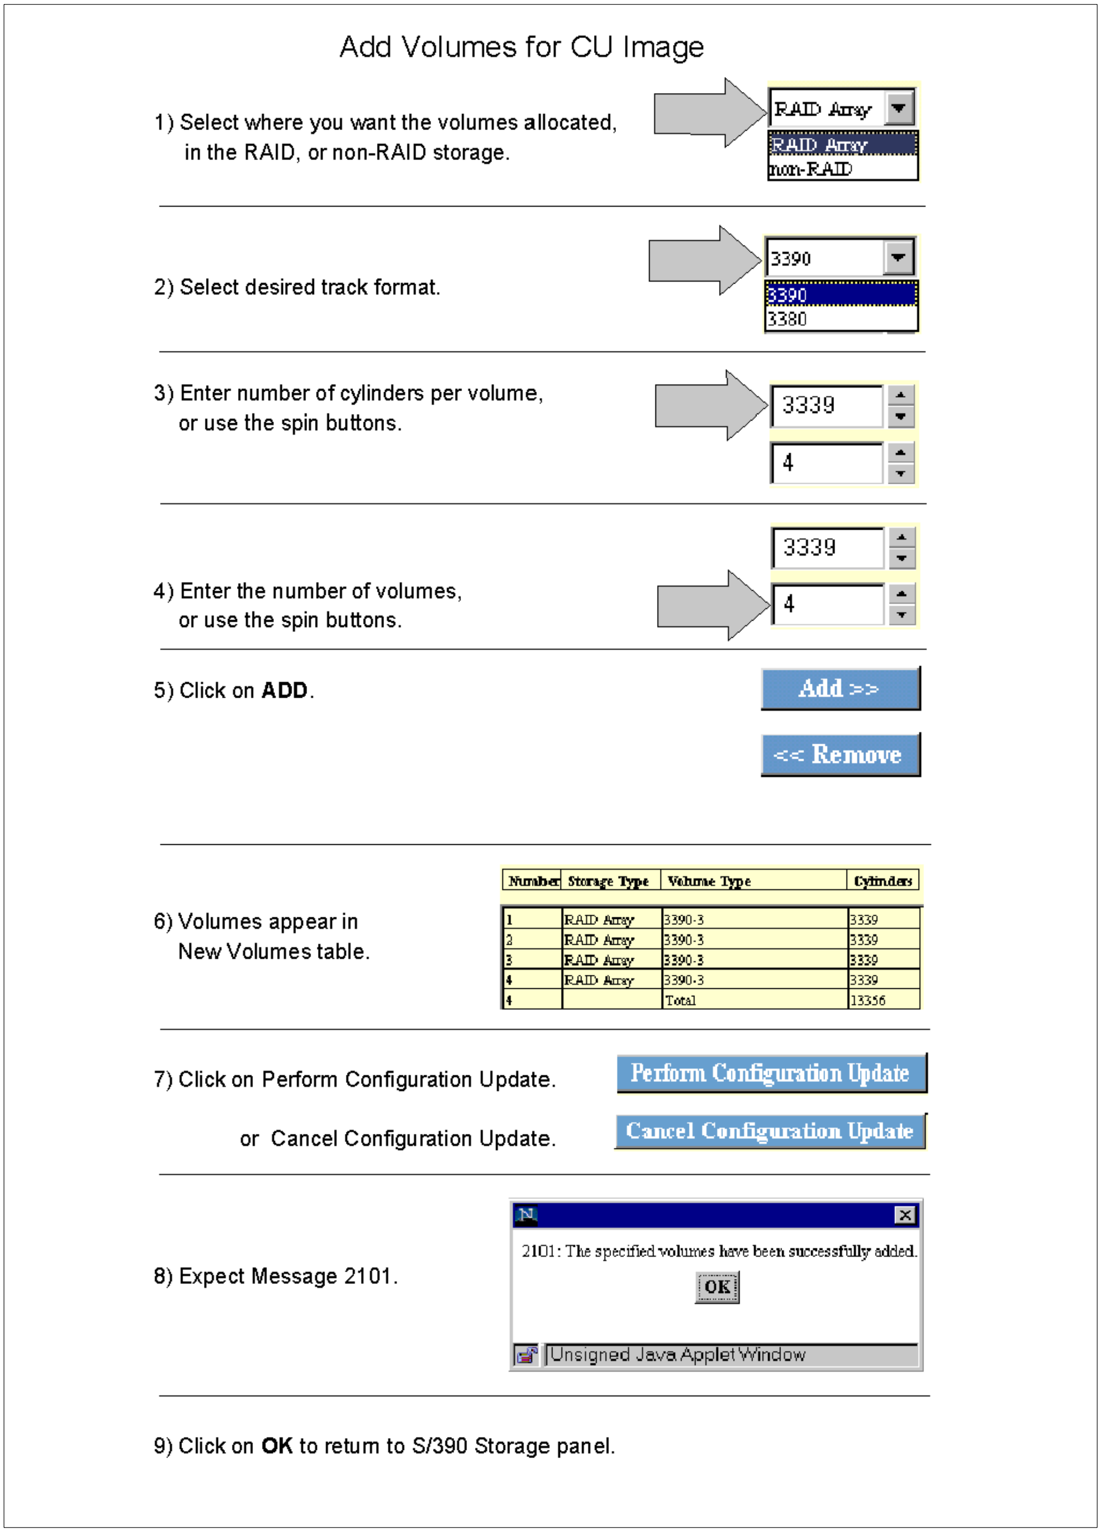

Figure 8-19 ESS Specialist: Procedure to add volumes

Note: The screen images on this figure may slightly differ from those that you see when using a later level ESS Specialist

Once the configuration action completes, the *S/390 Storage* panel will be displayed. Select your LCU; the *LCU Device* table now shows the custom volumes you added.

The new volumes are assigned available device IDs (unit addresses) in the order the volumes are created. The lowest available device ID in the LCU is used first. The device IDs map to the device numbers as seen by the zSeries operating system.

8.4.9 Modify PAV assignments

After you have defined an LCU, configured one or more disk groups for it, and created base volumes, you can assign PAV aliases to the volumes in the LCU. *This task is not necessary for an LCU for which PAV has been disabled (on the* Configure LCU panel).

To assign PAVs, select the LCU on the *S/390 Storage* panel, then click **Configure PAVs** to display the *Modify PAV Assignments* panel (Figure 8-20).

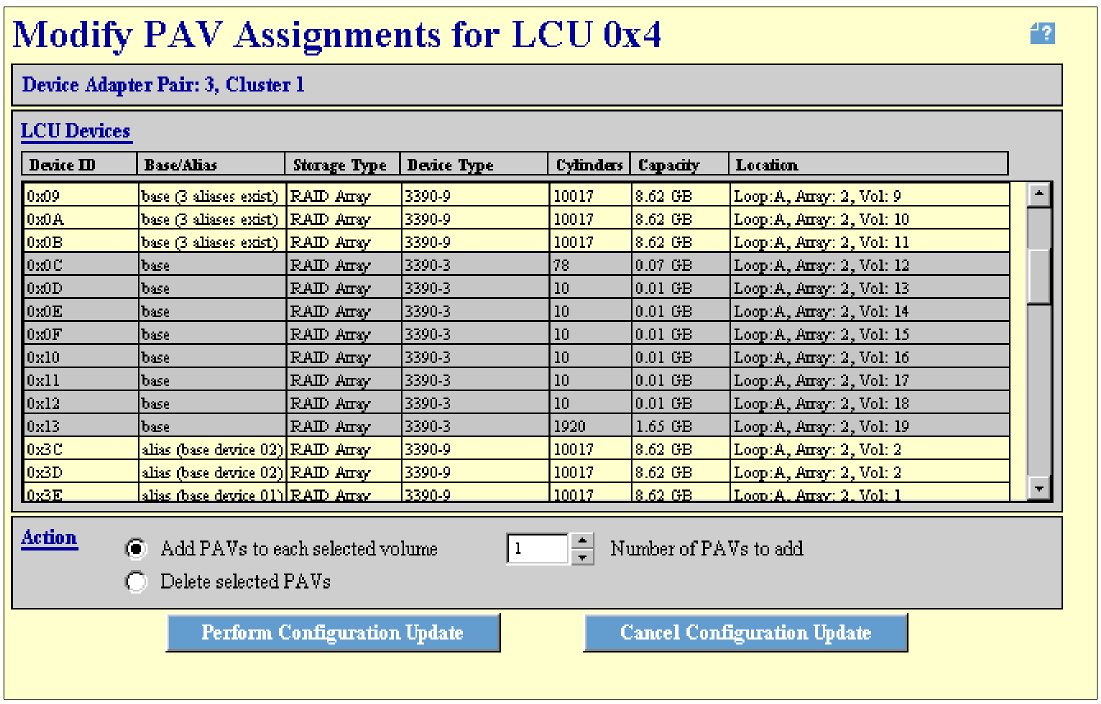

Figure 8-20 ESS Specialist: Configuring PAV aliases

1. On the LCU Devices table, select the devices you want to modify. To select rows, use the scroll bar to make the desired device visible in the *LCU Devices* list, then click the row to select it. It will highlight in grey. Additional rows can be selected using CTRL- left click. A range of rows can be selected by clicking at one end of the range, then SHIFT- left clicking at the other end. A row can be deselected using CTRL- left click.
2. Once you have selected the devices, select the action:
   - **Add PAVs to each selected volume**

     Specify the **Number of PAVs to add** for each selected volume. Note that you cannot enter a non-zero value here if PAVs are disabled for the LCU. The maximum number is limited by the total of 256 devices an LCU can support.
   - **Delete selected PAVs**

     This will delete the selected alias devices.

3. Click **Perform Configuration Update** to modify the PAVs.

You can select any combination of base devices and assign any number of PAVs to them, provided the total number of base devices and aliases for the LCU does not exceed 256. To add different number of aliases for different base devices, you need to repeat steps 1-3. You can assign more aliases to a base which already has aliases assigned. A single base can have up to 255 aliases assigned.

The upper part of the *LCU Devices* table in the example in Figure 8-20 shows base devices to which aliases have already been assigned. The Base/Alias column indicates how many aliases are currently assigned to a base. The lower part of the list shows existing alias device entries. For the PAV aliases, the Base/Alias column indicates the base device to which the alias is currently assigned to. The device table can be viewed in either this panel, or the *S/390 Storage* panel.

Aliases are assigned unit addresses in a specific order. See details in 8.4.6, “Configuring logical control units” on page 124.

When dynamic PAV support is enabled, WLM can reassign PAV aliases from one base to another. The alias devices retain their unit addresses, just the assignment to bases will change. With dynamic PAV the initial assignment of aliases to bases on the *Modify PAV Assignments* panel essentially only serves as a means of allocating a pool of aliases which the WLM can then work with and reassign to achieve the best performance.

WLM can only reassign aliases to base devices within the same LCU. You have to configure aliases for each LCU.

8.4.10 Configure FICON ports

The same ESS Fibre Channel host adapters are used for FICON attachments on S/390 and zSeries hosts or for open systems Fibre Channel protocol (FCP) attachments, depending on how they are configured. You must configure the ports for the intended attachment type using the *Configure Host Adapter Ports* panel (Figure 8-21).

You do not need to configure ESCON ports.

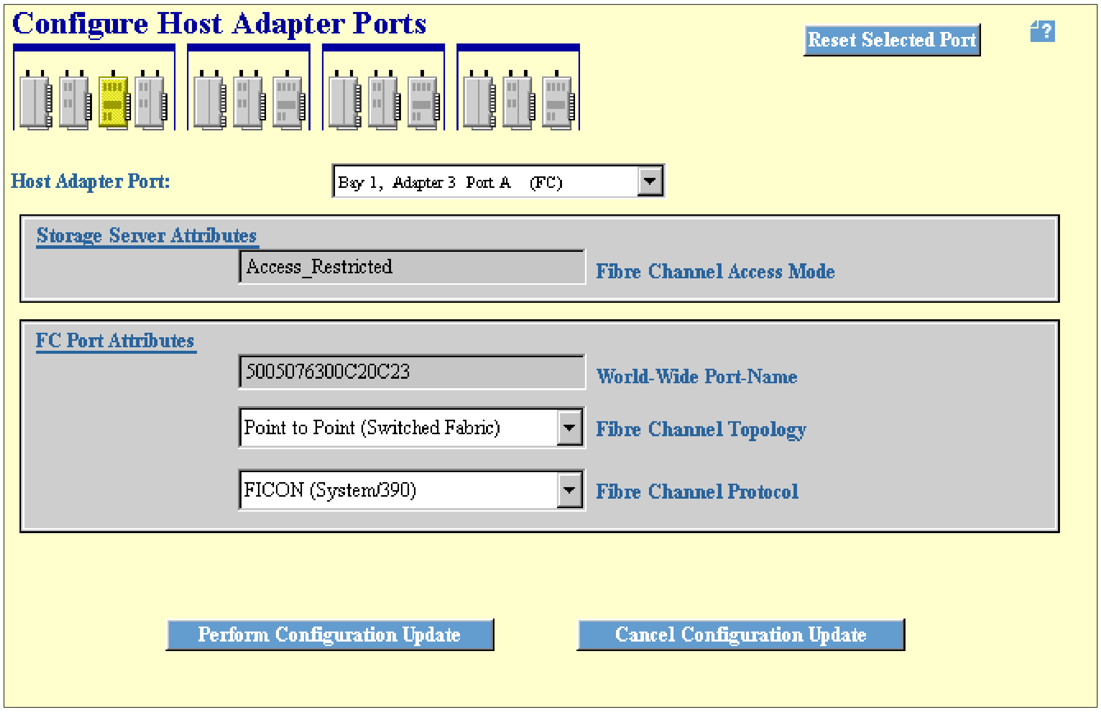

Figure 8-21 ESS Specialist: Configure Host Adapter Ports panel

1. Select a port to configure, either from the **Host Adapter Port** pull-down, or by clicking its icon.
2. The **Fibre Channel Topology** field shows the current Fibre Channel topology for the port you selected. If it is undefined, set the topology to **Point to Point (Switched Fabric)**. This is the only choice when the panel is accessed from the *S/390 Storage* panel.
3. The **Fibre Channel Protocol** field shows the current Fibre Channel protocol for the port you selected. If the topology is undefined, set it to **FICON (System/390)**. For unconfigured ports, this is the only choice if you accessed this panel from the *S/390 Storage* panel.

   If the topology is defined, you must first change the setting to **Undefined** before the ESS can make an alternate setting available for configuration.
4. Select the next port and repeat the previous steps. At any time, you can click **Reset Selected Port** to cancel any pending configuration changes made to the currently selected port.
5. When all ports have been defined, click **Perform Configuration Update** to apply or **Cancel Configuration Update** to cancel the configuration change for all modified ports.

The **Fibre Channel Access Mode** field in the *Storage Server Attributes* box shows the current Fibre Channel access mode for the ESS, either **Access_Any** or **Access_Restricted**. The mode is initially set during installation of the ESS, and can only be changed by a service representative during a service action. The mode only applies to open systems FCP connections, it is not relevant for FICON.

The World-Wide Port Name field shows the WWPN of the ESS host adapter port.

8.4.11 A step-by-step example

In this section we provide an example of how to configure CKD storage on a newly installed ESS subsystem using the ESS Specialist. Our example ESS is a mid-size FICON attached subsystem with 28 36GB 8-packs for a 6.3 TB total capacity. In this configuration two SSA loops only have two 8-packs, the others have four. As a result, some LCUs will have less capacity than the others.

Design guidelines for the target logical configuration:

- Configure 16 LCUs.
- Configure disk groups as RAID arrays.
- Assign two arrays to LCUs 0-B, one array to the remaining four LCUs.
- Assign one smaller (6+P) and one larger (7+P) array for LCUs 0-B for them to have equal capacities.
- Fully populate the arrays with 3390-9 logical volumes..
- Configure three aliases per each base.
- Use Starting PAV Address of 255 to assign the aliases at the high end address range.

Before starting to work with the ESS Specialist, complete the S/390 Storage requirements worksheet in the *IBM TotalStorage Enterprise Storage Server Configuration Planner,* SC26-7353, or fill in the information online at the ESS Web site. In the sample worksheet (Table 8-10) we have added information for the host I/O configuration.

Table 8-10 S/390 storage requirements worksheet

Logical control unit	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Cluster-Device Adapter	**1-1**	**2-1**	**1-2**	**2-2**	**1-3**	**2-3**	**1-4**	**2-4**
SSID	2200	2201	2202	2203	2204	2205	2206	2207
LCU emulation type	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6
Track mode for LCU	3390	3390	3390	3390	3390	3390	3390	3390
Device type for LCU	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9
Number of base devices	44+8	44+8	44+8	44+8	44+8	44+8	44+8	44+8
Number of PAV per base	3	3	3	3	3	3	3	3
Starting PAV address	255	255	255	255	255	255	255	255
Base device unit addr	0x00-33	0x00-33	0x00-33	0x00-33	0x00-33	0x00-33	0x00-33	0x00-33
Alias device unit addr	0x64-FF	0x64-FF	0x64-FF	0x64-FF	0x64-FF	0x64-FF	0x64-FF	0x64-FF
IOCP/HCD information								
CUNUMBR	900	910	920	930	940	950	960	970
CUADD	0	1	2	3	4	5	6	7
ADDRESS (3390B)	9000,80	9100,80	9200,80	9300,80	9400,80	9500,80	9600,80	9700,80
ADDRESS (3390A)	9050,176	9150,176	9250,176	9350,176	9450,176	9550,176	9650,176	9750,176

Logical control units	8	9	A	B	C	D	E	F
Cluster-Device Adapter	**1-1**	**2-1**	**1-2**	**2-2**	**1-3**	**2-3**	**1-4**	**2-4**
SSID	2208	2209	220A	220B	220C	220D	220E	220F
LCU emulation type	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6	3990-6
Track mode for LCU	3390	3390	3390	3390	3390	3390	3390	3390
Device type for LCU	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9	3390-9
Number of base devices	44+8	44+8	44+8	44+8	20+4	20+4	20+4	20+4
Number of PAV per base	3	3	3	3	3	3	3	3
Starting PAV address	255	255	255	255	255	255	255	255
Base device unit addr	0x00-33	0x00-33	0x00-33	0x00-33	0x00-17	0x00-17	0x00-17	0x00-17
Alias device unit addr	0x64-FF	0x64-FF	0x64-FF	0x64-FF	0xB8-FF	0xB8-FF	0xB8-FF	0xB8-FF
IOCP/HCD information								
CUNUMBR	980	990	9A0	9B0	9C0	9D0	9E0	9F0
CUADD	8	9	A	B	C	D	E	F
ADDRESS (3390B)	9800,80	9900,80	9A00,80	9B00,80	9C00,80	9D00,80	9E00,80	9F00,80
ADDRESS (3390A)	9850,176	9950,176	9A50,176	9B50,176	9C50,176	9D50,176	9E50,176	9F50,176

Notice that on the host we define full 256 devices for each LCU, even though we do not configure them all on the ESS. Figure 8-22 shows the layout of the disk groups. In the array boxes, the first number is the number of logical 3390-9 volumes in the array, followed by their unit addresses.

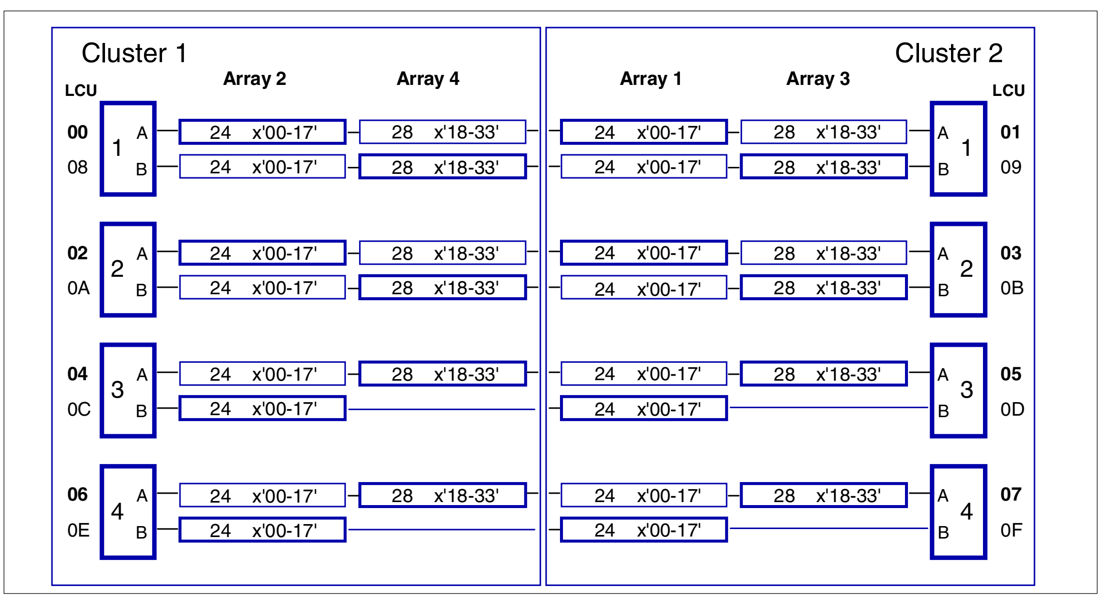

Figure 8-22 Sample configuration

This is the procedure:

1. Configure LCUs 0 to 7.

   Start by configuring the first eight LCUs using the *Configure LCU* panel. Initially only the first eight LCU are displayed on the panel. Select 3990-6 emulation mode, give the SSID, and select Starting PAV Address 255.

2. Configure the first disk group for LCUs 0 to 7.

   Configure Disk Group 2 for the even LCUs (0,2,4,6) in Cluster 1, and Disk Group 1 for the odd LCUs (1,3,5,7) in Cluster 2 using the *Configure Disk Groups* panel. Groups 1 and 2 have smaller capacities than Groups 3 and 4, which helps to locate and select them on the panel. In each LCU you have two Group1s or Group2s to select from, one in each SSA loop. Select the disk group for LCUs 0-3 from loop A, for LCUs 4-7 from loop B. Select a group from loop B here first because we want LCUs 4-7 to have their two arrays from different loops.

   Configure the disk groups as RAID arrays, in 3390 track format, and auto-allocate 20 type 3390-9 standard volumes on each. These will all be 6+P+S arrays.

3. Configure LCUs 8 to F.

   Once you have assigned the first disk group to LCUs 0-7, LCUs 8-F appear on the *Configure LCU* panel and become configurable. Define them as in step 1.

4. Configure the first disk group for LCUs 8 to F.

   Configure Disk Group 2 for the even LCUs (8,A,C,E) in Cluster 1, and Disk Group 1 for the odd LCUs (9,B,D,F) in Cluster 2 using the *Configure Disk Groups* panel. At this point you only have one Group1 or Group2 to select from.

   Configure the disk groups as RAID arrays, in 3390 track format, and auto-allocate 20 type 3390-9 standard volumes on each.

5. Add custom volumes for LCUs 0 to F.

   One array has now been assigned for each LCU with 20 standard volumes on them. In each array there is still room for four full-size 3390-9 custom volumes. Add them now in each LCU using the *Add Volumes* panel. Select volume size of 10017 cylinders.

6. Configure the second disk group for LCUs 0 to 3.

   Configure Disk Group 4 for LCUs 0 and 2 in Cluster 1, and Disk Group 3 for LCUs 1 and 3 in Cluster 2 using the *Configure Disk Groups* panel. In these LCUs you have two Group3s or Group4s to select from, one in each SSA loop. For these LCUs, select the disk group from loop B, that is, from the other loop than the first disk group in the LCU. This will provide maximum SSA loop bandwidth for the LCU.

   Configure the disk groups as RAID arrays, in 3390 track format, and auto-allocate 24 type 3390-9 standard volumes on each.

7. Configure the second disk group for LCUs 4 to B.

   Configure Disk Group 4 for the even LCUs (4,6,8,A) in Cluster 1, and Disk Group 3 for the odd LCUs (5,7,9,B) in Cluster 2 using the *Configure Disk Groups* panel. At this point you only have one disk group in each LCU to select from.

   Configure the disk groups as RAID arrays, in 3390 track format, and auto-allocate 24 type 3390-9 standard volumes on each.

   All disk groups have now been configured. LCUs C-F will not get a second array assigned.

8. Add custom volumes for LCUs 0 to B.

   In steps 6 and 7 you just added a second array with 24 standard volumes for LCUs 0-B. In each of these arrays there is room for four full-size 3390-9 custom volumes. Add them now using the *Add Volumes* panel. Select volume size of 10017 cylinders.

9. Configure PAVs for all LCUs.

   All 16 LCUs are now fully configured with logical volumes. LCUs 0-B have 52 logical volumes, LCUs C-F have 24 volumes. Now assign three PAV aliases to each base address (there would be room for even more) using the *Modify PAV Assignments* panel.

10. Configure FICON ports.

   Finally, if not already done so by the SSR, configure the FICON host adapter ports using the *Configure Host Adapter Ports* panel.

Our experience is that you need roughly two working days to perform this configuration process. Step 9 alone takes about four hours (about 10 PAVs per minute on an ESS model F20 at G4 LIC level).

Except for the last step, this example is fully applicable to ESCON connected systems also. However some considerations apply. If you were to use ESCON channels to connect this system, you should preferably have at least 16 ESCON channels, configured in four path groups with four LCUs on each channel. With less than 16 ESCON channels, you probably need to configure more LCUs on each channel and reduce the number of alias devices in each LCU to keep the total number of devices on a channel below 1024. For example, if each channel was to address eight LCUs, you would have 360 base devices (6 x 52 + 2 x 24) on a channel, leaving 664 addresses available for aliases, less than two aliases per base.

Example 8-1 shows part of the **DISPLAY M=CHP** command output for this configuration.

Example 8-1 DISPLAY M=CHP output

```
D M=CHP(91)

IEE174I 11.07.04 DISPLAY M 790
CHPID 91:  TYPE=1A, DESC=FICON POINT TO POINT, ONLINE
DEVICE STATUS FOR CHANNEL PATH C1
    0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
900 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
901 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
902 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
903 +  +  +  +  $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@
904 $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@
905 UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
906 UL UL UL UL AL AL AL AL AL AL AL AL AL AL AL AL
907 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
908 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
909 AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90A AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90B AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90C AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90D AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90E AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
90F AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL AL
************************ 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
AL DEVICE IS AN ALIAS       UL DEVICE IS AN UNBOUND ALIAS
```

9

S/390 and zSeries host setup tasks

This chapter examines the S/390 and zSeries hosts configuration tasks that you will need to perform to successfully implement the ESS. We describe the IOCP/HCD definitions required to set up the host. We also discuss such topics as how to implement the dynamic alias management function and the ESS Large Volume Support (LVS).

9.1 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 definitions are dependent on the ESS logical configuration. The IOCP/HCD definitions and ESS logical configuration must match each other. See section Section 8.2, “Designing the LCU address layout” on page 112 for a discussion on how to design the ESS configuration.

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 to the host multiple logical control units (LCUs) and logical devices rather than defining the physical 2105 machine configuration.

Due to the hardware internal switching functions, the enhanced recoverability, and internal logical structure within the ESS, you only need to code in HCD/IOCP one 2105/3990 CNTLUNIT statement per LCU rather than the 3990 convention of defining two CU images.

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 FICON port, on the other hand can address 16,384 devices, making it possible for a FICON channel to address all the 4096 devices you can have on an ESS. This allows you to logically “daisy chain” all 16 LCUs in the ESS on one FICON channel.

Depending on your zSeries operating system and the levels of software, you may be restricted to a subset of available new functions. See Chapter 7, “zSeries systems support” on page 95 to review the exploitation levels.

9.2 OS/390 software configuration

This section describes how to implement the ESS in the z/OS and OS/390 systems.

9.2.1 IOCP/HCD

For the currently supported releases of OS/390 and z/OS the ESS runs in what is called Exploitation mode (See Section 7.1.1, “OS/390 and z/OS support” on page 96).

In Exploitation mode the z/OS I/O subsystem recognizes the new ESS control unit type 2105. Up to 16 logical control units (CNTLUNIT UNIT=2105) can be defined per ESS. The host recognizes the ESS as a 2105 control unit and thus can exploit the PAV capability.

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 (See Section 9.2.4, “WLM dynamic alias management” on page 154 for discussion on Static and Dynamic PAVs).

The total number of devices defined (3390, 3380, 3390B, 3380B, 3390A, and 3380A) must not exceed 256 per CNTLUNIT statement.

Here are some additional considerations:

- Make sure the IOCP/HCD control unit CUADD, device type UNIT and device quantities match with 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 DFSMS device support message `IEA435I PHYSICAL DEVICE INCONSISTENT WITH LOGICAL DEFINITION`. This message is generated during the vary on-line process for the device.
- The Multiple Allegiance function does not require specific definitions within OS/390 or z/OS, since it is provided by ESS internal logic.
- Do not define esoteric unit names for PAV alias devices. Alias devices are not used for device allocation.

> **Note:** Previous releases of OS/390 and DFSMS/MVS, now not supported, will have to run the ESS in either Transparency or Toleration mode. These modes of operating the ESS must be considered as transitional because the recommendation is to have all your systems at the currently supported levels of software and therefore capable to run the ESS in Exploitation mode.

If you run an MVS in Transparency mode, then define the control unit in IOCP/HCD as UNIT type 3990. You only need one CNTLUNIT statement for each LCU in IOCP/HCD rather than the 3990 convention of defining two CU images. No Parallel Access Volume (PAV) device types, 3390B, 3390A, 3380B, 3380A are permitted in Transparency mode.

If the OS/390 I/O subsystem recognizes the new ESS control unit type 2105 (3990-6 emulation), but no ESS native functions are supported, then you are running the ESS in Toleration mode. The new Parallel Access Volume (PAV) device types of base (IODEVICE UNIT=3390B or 3390B) and alias (UNIT=3390A or 3380A) are 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.

9.2.2 FICON host connectivity

The ESS now supports FICON channels. FICON channel connectivity brings some differences and provides benefits over the ESCON channel connectivity. These are among the benefits you should consider:

- Increased channel device-address support. FICON channel supports 16,384 devices on one channel, compared to 1,024 devices on an ESCON channel. This makes possible to any FICON channel connected to the ESS, to address all the 4096 devices you can have within the ESS. This extra flexibility will simplify your configuration setup and management.
- FICON provides an increased number of channel-to-control unit concurrent I/O connections. ESCON supports one I/O connection at any one time while FICON channels support multiple concurrent I/O concurrent connections. While an ESCON channel can have only one I/O operation at a time, the FICON channel can have I/O operations to multiple LCUs at the same time, even to the same LCU, by using the FICON protocol frame multiplexing.
- FICON has more than 4 times the effective channel bandwidth for the initial implementation (70 MB/sec, compared to 17 MB/sec for ESCON).
- Reduced number of channels and required fibers with increased bandwidth and I/O rate.
- Greater un-repeated fiber link distances (from 3 km. for ESCON to up to 10 km., or 20 km. with an RPQ, for FICON).

All these factors allow you to lay out more simple redundant configurations using FICON, accessing more data with better performance than what is possible with ESCON.

Table 9-1 summarizes the maximum number of ports, devices, and logical paths supported by FICON and ESCON attachments.

Table 9-1 FICON and ESCON comparison

	FICON	ESCON
Host adapter ports per ESS	16	32
Devices per channel	16384	1024
Logical paths per host adapter port	256	64
Logical paths per LCU (ESS limitation)	128	128

The ESS limitation of 128 logical paths per LCU limits the total number of logical paths per ESS to 2048 for both FICON and ESCON.

The intermixing of ESCON and FICON channels to the same LCU from the same operating system image is only supported for migration purposes. It is not a recommended configuration for the production environment. The coexistence is useful during the transition period from ESCON to FICON channels. The mixture allows you to dynamically add FICON native channel paths to a control unit while keeping its devices operational. A second dynamic I/O configuration change can then remove the ESCON channels while keeping devices operational.

For more information on ESS FICON support, see the *IBM TotalStorage Enterprise Storage Server,* SG24-5465. You may also refer to the document at `http://www.storage.ibm.com/hardsoft/products/ess/support/essficonwp.pdf` for further considerations on FICON system attachment.

9.2.3 IOCP/HCD examples

Following are two IOCP/HCD examples.

Example 1: ESCON with ESCON Director, 8 LCUs

Sample IOCDS

Example 9-1 shows 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 devices 3390A.

Example 9-1 IOCDS for ESCON attached ESS

```
CHPID PATH=(8A,9A,A0,B0,C0,D0,E2,E6),SWITCH=30,TYPE=CNC
CNTLUNIT CUNUMBR=1A00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=0
CNTLUNIT CUNUMBR=1A80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=1
CNTLUNIT CUNUMBR=1B00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=2
CNTLUNIT CUNUMBR=1B80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C3,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=3
CNTLUNIT CUNUMBR=1C00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=4
CNTLUNIT CUNUMBR=1C80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=5
CNTLUNIT CUNUMBR=1D00,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
```

```
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=6
CNTLUNIT CUNUMBR=1D80,PATH=(8A,E2,A0,C0,9A,B0,D0,E6),UNIT=2105,       *
       LINK=(C0,C1,C2,C3,C0,C1,C2,C3),UNITADD=((00,128)),CUADD=7
IODEVICE ADDRESS=(1A00,016),CUNUMBR=(1A00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1A10,112),CUNUMBR=(1A00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1A80,016),CUNUMBR=(1A80),STADET=Y,UNIT=3390B,       *
       UNITADD=00
IODEVICE ADDRESS=(1A90,112),CUNUMBR=(1A80),STADET=Y,UNIT=3390A,       *
       UNITADD=10
IODEVICE ADDRESS=(1B00,016),CUNUMBR=(1B00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1B10,112),CUNUMBR=(1B00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1B80,016),CUNUMBR=(1B80),STADET=Y,UNIT=3390B,       *
       UNITADD=00
IODEVICE ADDRESS=(1B90,112),CUNUMBR=(1B80),STADET=Y,UNIT=3390A,       *
       UNITADD=10
IODEVICE ADDRESS=(1C00,016),CUNUMBR=(1C00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1C10,112),CUNUMBR=(1C00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1C80,016),CUNUMBR=(1C80),STADET=Y,UNIT=3390B,       *
       UNITADD=00
IODEVICE ADDRESS=(1C90,112),CUNUMBR=(1C80),STADET=Y,UNIT=3390A,       *
       UNITADD=10
IODEVICE ADDRESS=(1D00,016),CUNUMBR=(1D00),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(1D10,112),CUNUMBR=(1D00),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(1D80,016),CUNUMBR=(1D80),STADET=Y,UNIT=3390B,       *
       UNITADD=00
IODEVICE ADDRESS=(1D90,112),CUNUMBR=(1D80),STADET=Y,UNIT=3390A,       *
       UNITADD=10
```

If the LCU will be shared with another system image, then FEATURE=SHARED must be specified for the IODEVICE macro.

HCD panels

The following are the HCD panels used to generate the IOCDS shown in Example 9-1. The devices defined are 16 PAV base devices (unit type 3390B) and 112 PAV aliases (unit type 3390A), a total of 128 devices on the LCU.

Begin by defining the 2105 logical control unit as shown in Figure 9-1.

```
 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 9-1 HCD panel control unit 2105

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 5-3 on page 84). The control unit descriptions are shown in the ESS *Specialist S/390* Storage panel (Figure 9-2). The LCU number (CUADD), logical control unit emulation type and the SSID assigned are also detailed. Native 2105 support is shown as 3990-6 emulation.

> **Figure Note:** Figure 9-2 is for illustration purposes only; the values on the figure do not match the example configuration.

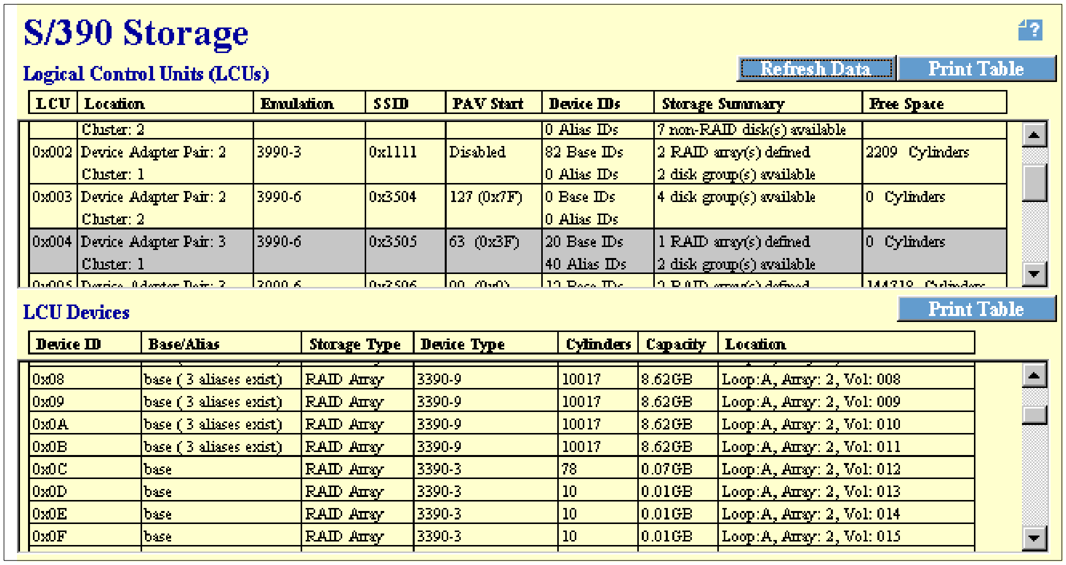

Figure 9-2 ESS Specialist LCU and device details

The CUADD (see Figure 9-3) must match the LCU number specified in the ESS logical configuration. CUADD=0 associates the CU image with ESS LCU 0, CUADD=1 with LCU 1, and so on up to CUADD=F.

```
ssssssssssssssssssssss Select Processor / Control Unit ssssssssssssssssssss
                                                  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 9-3 HCD panel CUADD

Select the device address and range required, as shown in the HCD panel in Figure 9-4. If you are defining PAV devices, the device range match the aliases defined on the ESS.

```
ssssssssssssssssssssss Select Processor / Control Unit ssssssssssssssssssssss
                                                 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 9-4 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 *LCU Devices* table on the ESS Specialist *S/390 Storage* panel (Figure 9-2). These are “real” base volumes.

Start at the lowest device address available and define your PAV base and non-PAV (standard 3390) devices.

Figure 9-5 shows the definition for 16 PAV base devices of type 3390B. Valid devices for bases are 3390B and 3380B.

```
sssssssssssssssssssssssssssssss Add Device ssssssssssssssssssssssssssssssss

 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 9-5 HCD Add Device panel for 16 base devices

The new WLMPAV parameter is shown on the *Define Device Parameters / Features* panel in Figure 9-6. It should be set to **NO** for Static PAVs or non-PAV devices, and **YES** for dynamic PAVs. See Section 9.2.4, “WLM dynamic alias management” on page 154 for more details.

```
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 9-6 HCD Define Device Parameters panel with WLMPAV parameter

Figure 9-7 shows the HCD Define Device Parameters / Features panel for alias device definition. Notice that the only device attribute you can set on an alias device is WLMPAV. Valid devices for aliases are 3390A or 3380A.

```
 ssssssssssssssssssssssssssssssss Add Device ssssssssssssssssssssssssssssssss

  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 9-7   HCD panel define alias device

Example 2: Direct FICON attachment, 16 LCUs

Example 9-2 shows the IOCP statements for a fully configured FICON attached ESS with 16 LCUs. We have four FICON channels directly connected. All channels address all 16 LCUs. Each LCU is defined with 80 base devices and 176 alias devices. The number of base volumes, 80, is the number of full-size 3390-9 volumes in an LCU in a fully configured 11.2TB ESS (36GB DDMs).

A native FICON channel is defined as TYPE=FC. A FICON channel connected through a FICON Bridge card on the ESCON director is defined as TYPE=FCV.

The IOCP in Example 9-2 corresponds to the sample configuration in Section 8.4.11, “A step-by-step example” on page 140.

Example 9-2 IOCP for FICON attached ESS

```
CHPID PATH=(91,92,93,94),TYPE=FC,SHARED
*==============================================================
* 16 LCU'S
CNTLUNIT CUNUMBR=900,PATH=(91,92,93,94),UNITADD=((00,256)),       *
         UNIT=2105,CUADD=0
CNTLUNIT CUNUMBR=910,PATH=(91,92,93,94),UNITADD=((00,256)),       *
         UNIT=2105,CUADD=1
     :
CNTLUNIT CUNUMBR=9F0,PATH=(91,92,93,94),UNITADD=((00,256)),       *
         UNIT=2105,CUADD=F
*==============================================================
* 80 BASE + 176 ALIAS DEVICES PER LCU
IODEVICE ADDRESS=(9000,080),CUNUMBR=(900),STADET=Y,UNIT=3390B,    *
         FEATURE=SHARED,UNITADD=00
IODEVICE ADDRESS=(9050,176),CUNUMBR=(900),STADET=Y,UNIT=3390A,    *
         FEATURE=SHARED,UNITADD=50
IODEVICE ADDRESS=(9100,080),CUNUMBR=(910),STADET=Y,UNIT=3390B,    *
         FEATURE=SHARED,UNITADD=00
IODEVICE ADDRESS=(9150,176),CUNUMBR=(910),STADET=Y,UNIT=3390A,    *
         FEATURE=SHARED,UNITADD=50
     :
IODEVICE ADDRESS=(9F00,080),CUNUMBR=(9F0),STADET=Y,UNIT=3390B,    *
         FEATURE=SHARED,UNITADD=00
IODEVICE ADDRESS=(9F50,176),CUNUMBR=(9F0),STADET=Y,UNIT=3390A,    *
         FEATURE=SHARED,UNITADD=50
```

9.2.4 WLM dynamic alias management

The IBM TotalStorage Enterprise Storage Server's implementation of PAVs introduces the concept of base addresses (base devices) and alias addresses (alias devices).

Base address This is the actual unit address of the volume. There is only one base address for any volume.

Alias address Alias addresses are mapped back to a base device address. I/O scheduled for an alias is physically performed against the base by the ESS. No physical disk space is associated with an alias address, however, they do occupy operating system storage. Alias UCBs are stored above the 16 MB line.

The link between base address and alias addresses occurs during IPL processing. A base address and its aliases must exist in the same ESS logical subsystem (LSS). Together the base and alias addressed are termed exposures. The number of exposures for a volume is base plus number of current aliases.

While the base address is the actual unit address of a given volume, there can by many alias addresses assigned to a base address, and any or all of those alias addresses can be reassigned to a different base address.

Alias devices assigned to a base device allow the host operating system to start multiple I/O requests to the base device in parallel and thus reduce UCB queueing in IOS. Depending on the system configuration, the total number of aliases can be limited by addressing or other constraints. Therefore, aliases can be a relatively scarce resource that needs to be managed to ensure the aliases are assigned to the base devices that need them at the moment.

For best I/O management, the aliases should be assigned to the busiest base devices and those handling the most important work. You can reassign aliases manually using the ESS Specialist. However, this process is labor intensive and, to achieve optimum results, must be repeated whenever there are workload changes. The process of alias reassignment needs to be automatic and be sensitive to the importance and level of I/O activity on the device.

The OS/390 and z/OS Workload Manager (WLM) and I/O Supervisor (IOS) components allow the automatic management of aliases. This function is called *dynamic alias management.* We also refer to this as *dynamic PAV* support. With dynamic alias management, WLM can automatically perform alias device reassignments from one base device to another to help work meet its goals and to minimize IOS queueing as workload requirements change. By default, the function is turned off, and must be explicitly activated for the sysplex through an option in the WLM service definition. A device-level option is available through HCD.

Dynamic PAV is supported on systems running OS/390 2.7 and DFSMS 1.5 with enabling maintenance. To exploit the use of dynamic PAVs, you need to be running your sysplex WLM in *goal mode*.

Static PAV support is available on systems running OS/390 1.3 and DFSMS 1.3 and above. With static PAV, alias devices remain assigned to the base devices you have assigned them using the ESS Specialist. Alias reassignment can be done manually using the ESS Specialist. PAVs are static if you are running your sysplex WLM in compatibility mode.

Controls for dynamic workload management

The following three options control the use of WLM dynamic alias management:

- Use the **Dynamic alias management** option on the *WLM Service Coefficient/Service Definition Options* panel (Figure 9-8) to enable dynamic alias management globally in a sysplex. The default value is **NO** (dynamic alias management is disabled).
- The **I/O priority management** option on the *WLM Service Coefficient/Service Definition Options* panel has on effect on which dynamic alias management algorithms will be used. The default value is **NO**.
- Use the **WLMPAV** parameter in a device's HCD definition to individually enable or disable dynamic alias management on a given device. The default is **YES**.

```
  Coefficients/Options  Notes  Options  Help
--------------------------------------------------------------
              Service Coefficients/Service Definition Options

Enter or change the Service Coefficients:

CPU  . . . . . . . . . . . . . ______ (0.0-99.9)
IOC  . . . . . . . . . . . . . ______ (0.0-99.9)
MSO  . . . . . . . . . . . . . ______ (0.0000-99.9999)
SRB  . . . . . . . . . . . . . ______ (0.0-99.9)

Enter or change the service definition options:

I/O priority management . . . . . . . . YES  (Yes or No)
Dynamic alias management  . . . . . . . YES  (Yes or No)
```

Figure 9-8 Activate WLM dynamic alias management

When you specify **Yes** for the Dynamic alias management option, you enable dynamic alias management globally throughout the sysplex. WLM will keep track of the devices used by different workloads and broadcast this information to other systems in the sysplex. If WLM determines that a workload is not meeting its goal due to IOS queue time, then WLM attempts to find alias devices that can be moved to help that workload achieve its goal. Even if all work is meeting its goals, WLM will attempt to move aliases to the busiest devices to minimize overall queueing.

Note that the WLM Dynamic alias management option is a global setting for the entire sysplex. If any system in the sysplex specifies YES for Dynamic alias management, it is enabled for all systems in the sysplex. There is no consistency checking for this parameter. It is an installation's responsibility to coordinate definitions consistently across a sysplex. WLM will not attempt to enforce a consistent setting of this option.

WLM Service definition options for dynamic alias management

There are two algorithms controlling the distribution of aliases within a subsystem - the goal algorithm and the efficiency algorithm. The goal algorithm recognizes when a service class is missing its goal due to IOS queue delay and moves aliases to the base devices serving the most important work that needs help. The efficiency algorithm simply moves aliases to the base devices that have the greatest IOS queueing without regard to the importance or goals of the work. These two algorithms work together to first help service classes meet goals and beyond that to minimize overall IOS queueing in the subsystem.

You control the activation of these algorithms through the I/O priority management and Dynamic alias management options in the WLM service definition as shown in Table 9-2:

Table 9-2 Dynamic alias management algorithms

I/O priority management	Dynamic alias management	Dynamic alias algorithm in effect
YES	YES	both efficiency and goal
NO	YES	efficiency only
YES	NO	none
NO	NO	none

If you enable dynamic alias management, you should also enable I/O priority management. So you need to specify YES for both of these options on the panel. If I/O priority management is set to NO, you will get only the efficiency part of dynamic alias management and not the goal-oriented management. This means that WLM will make alias moves that minimize overall IOS queueing, but these moves will not take service class goals into consideration. (Note that you still need to be in WLM goal mode, even if you wish to use the efficiency function only.)

Device eligibility for dynamic management

While you can globally enable or disable dynamic alias management on the WLM ISPF panel, you can also individually enable or disable dynamic alias management on a given device via HCD. You can do this by specifying WLMPAV=YES or NO in that device's definition on the *HCD Define Device Parameters / Features* panel. See examples in section Section 9.2.3, "IOCP/HCD examples" on page 148.

The relationship of the WLM Dynamic alias management option and the HCD WLMPAV parameter is following:

- If the WLM Dynamic alias management option for the sysplex is set to NO, then all PAVs are static regardless of their WLMPAV setting.

- If the WLM Dynamic alias management option for the sysplex is set to YES, then PAVs are dynamic or static depending on their WLMPAV setting.

You can specify WLMPAV=YES or NO on both base and alias devices in HCD. The WLMPAV settings on an alias device, however, is only meaningful when the alias device is assigned to a base device that is offline, as follows:

- If the base device is *offline*, then only alias devices with WLMPAV set to YES are eligible to be reassigned to other base devices.

  The WLMPAV setting on the base device itself is irrelevant when the base device is offline, for either “giving” or “receiving” aliases. (Even if WLMPAV was set to YES on the base device, it cannot have new aliases assigned to it, as it is offline.)

- If the base device is *online*, then the WLMPAV settings on the aliases are ignored, and WLM can reassign aliases as follows:
  - If WLMPAV is set to YES on the base device, then the aliases can be reassigned regardless of their WLMPAV settings.
  - If WLMPAV is set to NO on the base device, then no aliases can be reassigned, regardless of their WLMPAV settings.

For a WLMPAV=YES base device, the aliases initially assigned to it should be allowed to default to WLMPAV=YES. The only situation where you might want to change an alias to WLMPAV=NO is if the alias is initially assigned to a WLMPAV=NO base device. Because the base is set to NO, the aliases initially assigned to it will not be moved to other bases by WLM. Then, because the aliases are set to NO, if the base is ever varied offline, the aliases remain assigned to that base and cannot be reassigned by WLM to other bases. Certain combinations of WLMPAV settings are not recommended, as described in Table 9-3:

Table 9-3 Effects of WLMPAV settings on base and alias devices in HCD

Base Device WLMPAV Setting	**Alias Device WLMPAV Setting**	**Effects/Recommendations**
YES	YES	**If base is online**: Base is WLM-managed. Aliases can be freely moved to and from the base device by WLM. **If base is offline**: Aliases become unbound and are available to WLM to assign to other WLM-managed bases.
YES	NO	**Not recommended**. If base is WLM-managed, then it is not predictable which aliases will remain assigned to that base when the base goes offline. If the base device is set to YES, then you should set the aliases to YES as well. (See previous option.)
NO	YES	**Not recommended**. If the base is not WLM-managed, then you risk losing all of its aliases when the device goes offline. (See next option.)
NO	NO	**If base is online**: Base is not WLM-managed. The initial aliases assigned to this base remain there. **If base is offline**: Aliases remain bound to the offline base device and are not available to WLM for reassignment. When the base comes back online, it retains its initial alias assignments.

Aliases of an offline device are considered unbound and WLM will consider unbound aliases 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 alias management.

Scope of management

WLM's scope of management is sysplex wide. WLM's decisions about reassigning aliases is based on sysplex-wide awareness of workload and whether the workload is meeting its goals.

It is recommended not to use dynamic alias management for a device unless all systems sharing that device have dynamic alias management enabled. Otherwise, WLM will be attempting to manage alias assignments without taking into account the activity from the non-participating systems. If systems in the sysplex are at levels that do not support dynamic PAV, but do support static PAVs, a potential performance problem may exist. Static PAV systems will recognize PAV reassignment instituted by dynamic PAV systems, however, the I/O contributed by the back level systems will not be considered by WLM during alias reassignment.

Dynamic PAVs should not be defined for volumes that are shared outside the sysplex as WLM has insufficient information about volume usage to correctly manage these devices. In a configuration where two sysplexes share dynamic PAVs, the WLM on sysplex1 would make alias/base decisions based on the workload on sysplex1, and the WLM on sysplex2 would make alias/base decisions based on the workload on sysplex2. As neither WLM knows anything about the workload on the other sysplex, the two could end up moving the same alias back and forth between different bases, just causing overhead and not really helping either sysplex.

Preparing for dynamic alias management

Following is a checklist of the steps required to prepare for using dynamic alias management and the steps to activate it:

1. **Upgrade OS/390:** Install OS/390 V2R7 or higher on all systems that will be sharing the ESS volumes. The support is integrated in z/OS.
2. **Switch to goal mode:** Dynamic alias management is only active in goal mode.
3. **Turn on I/O priority management:** Set I/O priority management to YES in the WLM service definition. This function is necessary to provide the I/O delay data needed for dynamic alias management. Check velocity goals before switching to I/O priority management because this function factors I/O using and delay samples into the velocity calculation. To aid in this migration step, the RMF *Workload Activity* report shows what the velocities will be when I/O priority management is set to YES.
4. **Configure devices on the ESS:** Using the ESS Specialist, define the logical volumes. They will be assigned base unit addresses. Then configure PAV aliases for the bases. When you initially define the aliases, you assign them to a base device. WLM can then dynamically reassign them. Refer to Section 8.4, “ESS custom configuration for CKD storage” on page 118.

   Our recommendation is to configure one alias for a 3390-2/3 type volume, three for a 3390-9 type volume. You may configure as many aliases as your addressing constraints allow. Base devices plus aliases cannot exceed 256 for each logical subsystem.

   In order for dynamic alias management to be most effective from the beginning, try to spread out your aliases in the initial definition. WLM will only take one PAV alias at a time from a given base. If the aliases are spread across bases, WLM can make multiple moves at one time. If aliases are assigned to the same base, WLM will still move them to the busy devices, but it will take some more time for WLM to reassign the aliases appropriately.

5. **Define device addresses on the host:** Using HCD, define base and alias device addresses to z/OS to match the unit address specified in ESS Specialist. Ensure the IODFs for the base devices and aliases are the same across the sysplex. Leave the WLMPAV parameter set to YES (the default) for the bases and aliases. This assumes a sysplex where all the systems using the device are capable of dynamic alias management.

Enabling dynamic alias management

1. **Install service:** Install enabling function for IOS and WLM together. These are the PTFs for APARs OW39854 for WLM and OW40669 for IOS (please check the latest PSP bucket information). They are corequisites. Roll across all OS/390 V2R7 and V2R8 systems in the sysplex. Support is included in OS/390 2.10 and in z/OS. The default for dynamic alias management in the sysplex is NO.
2. **Turn on SMF Type 99 recording if desired:** This data can be used by IBM for analysis of alias management activity if that becomes necessary.
3. **Activate function:** Switch on alias management for the sysplex by editing the WLM service definition and changing **Dynamic alias management** and **I/O priority management** to YES. Install the modified service definition and activate a policy.

   The dynamic alias management option is NO by default to allow the enabling service to be installed across the sysplex before activating the function. Overriding the default changes the functionality level of the service definition to LEVEL008. This means that once the option is set to YES, you cannot edit, install, activate, or extract the service definition from a pre-R7 system in the sysplex.

   This is the usual functionality level restriction in a mixed-release sysplex. There are also restrictions when sharing service definitions between sysplexes. For a complete discussion of migration considerations when using a new functionality level, see Chapter 16: “Workload Management Migration” in *z/OS MVS Planning: Workload Management*, SA22-7602.

   After you install the service definition for dynamic alias management, DO NOT manipulate the service definition with a back-level R7 or R8 system that does not have the service installed. The results will be unpredictable.
4. **Monitor results:** Monitor dynamic alias management using operator commands and the RMF device reports. The command **`D M=DEV(devnum-devnum)`** displays the current number of exposures (base + aliases) for a base device. The new PAV column in the RMF device reports gives the number of exposures observed for a device at the end of the RMF interval. An asterisk (*) next to the value in the PAV column indicates that the number of exposures changed either up or down during the last interval.

Additional considerations

- To invoke the Parallel Access Volume (PAV) 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 base devices (3390B/3380B) and alias devices (3390A/3380A) on the same LCU.
- The number of base and aliases may vary in different logical control units. On the HCD you must ensure that the base and alias definitions for each LCU are consistent with the ESS logical configuration.
- Non-PAV device types do not have access to the alias device pool.
- PAVs are static if you are running your sysplex WLM in compatibility mode. You must be sure to set WLMPAV=NO in your IOCP/HCD device attributes panel.

- Once WLM has assigned an alias to a base, the alias will remain associated with that new base until WLM decides to move it again.

Displaying PAV information

MVS system commands **DISPLAY M=DEV** and **DISPLAY M=CHP(xx)** provide new PAV related information.

The **DEVSERV QPAVS** (**DS QP**) command can be used to display the status information on a device, group of devices, or an LCU, including PAV devices.

Figure 9-9 shows a **DEVSERV QPAVS** command and resulting display. The `UNIT NUM` heading is the host device addresses and the `UNIT ADDR` heading details the ESS internal device address. The `TYPE` heading shows `BASE`, `ALIAS`, or `NON-PAV` as the device type. It also displays base and alias affiliations. The `STATUS` heading contains information to highlight discrepancies between the IODF and the ESS configuration, for example an invalid alias `INV-ALIAS`.

```
DS QP,D222,VOLUME

IEE459I 08.20.32 DEVSERV QPAVS 591
    HOST                             SUBSYSTEM
 CONFIGURATION                      CONFIGURATION
 -------------                   --------------------
 UNIT                                  UNIT    UA
 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 9-9 DEVSERV QPAVS command

See 11.5, “MVS system commands” on page 187 for more examples.

9.2.5 Custom volumes

Custom volumes are defined as 3380 or 3390 format devices types, usually with a non-standard number of cylinders. There are no special parameters to set when defining them to the host I/O Subsystem. Prior to using the volume, you must first format the volume as all other volumes via ICKDSF minimal INIT. A sample job is included in Example 9-3 on page 165.

When you initialize custom volumes, adjust the VTOC and VTOC Index size to correspond to the volume size. The same also applies to the VVDS.

9.2.6 Large volumes

The ESS initially supported custom volumes up to 10017 cylinders, the size of the largest standard volume, the 3390 model 9. This was the limit set by the operating system software. The ESS Large Volume Support (LVS) enhancement has now increased the upper limit to 32760 cylinders, approximately 27.8 GB. This enhancement is provided as a combination of ESS Licensed Internal Code (LIC) changes and system software changes, available for z/OS, OS/390, and z/VM. Refer to Section 7.4, “Large Volume Support” on page 99 for details on software support.

Benefits of using large volumes can be briefly summarized as follows:

- They reduce storage management tasks by allowing you to define and manage smaller configurations.
- They reduce the number of multivolume data sets you have to manage.
- They relieve architectural constraints by allowing you to address more data within the existing 64K subchannel number limit.

Large volume considerations

The size of the logical volume defined does not have an impact on the performance of the ESS subsystem. The ESS does not serialize I/O on the basis of logical devices so an increase in the logical volume size does not affect the ESS backend performance. Host operating systems, on the other hand, serialize I/Os against devices. As more data sets reside on a single volume, there will be greater I/O contention accessing the device. With Large Volume Support (LVS) it is more important than ever to try to minimize contention on the logical device level. To avoid potential I/O bottlenecks on devices

- Exploit the use of Parallel Access Volumes to reduce IOS queuing on the system level.
- Eliminate unnecessary reserves by using WLM in GOAL mode.
- Multiple allegiance will automatically reduce queuing on sharing systems.

Parallel Access Volume (PAV) support is of key importance when implementing large volumes. PAV enables one MVS system to initiate multiple I/Os to a device concurrently. This keeps IOSQ times down and performance up even with many active data sets on the same volume. PAV is a practical “must” with large volumes. We discourage you from using large volumes without PAV. In particular, we recommend the use of dynamic PAV.

As the volume sizes grow larger, more data and data sets will reside on a single S/390 device address. Thus, the larger the volume, the greater the multi-system performance impact will be of serializing volumes with RESERVE processing. You need to exploit a GRS Star Configuration and convert all RESERVE's possible into system ENQ requests.

Implementation steps

When planning for your logical volume configuration, you will notice that you cannot fully utilize the array capacity with all 32760 cylinder volumes. For example, a 6+P+S array of 72.8GB disks will contain 14 maximum size Large Volumes, but some 26.6 GB, three times the capacity of a 3390-9 will be left over. To use all the space, configure one or more smaller custom volumes to fill the array. Alternatively, if you prefer all of your volumes to be of the same size, select a smaller uniform volume size that will better utilize the array capacity.

- On the ESS, you configure large volumes like any other custom volume using the ESS Specialist *Add Volumes* panel.
- On the host IOCP/HCD, define a large volume as 3390, or 3390B like any other 3390 type volume.

- A large volume is reported by the host system as a 3390-9 device containing up to 32760 cylinders. They have 3390 track format. The maximum size for 3380 track format volumes is 3339 cylinders.

The large volume coexistence support for DFSMS/MVS 1.4 and 1.5 will allow these system levels to coexist in the same Sysplex with LVS systems. You must install the support in order to prevent unpredictable results that may arise from systems without Large Volume Support (LVS) accessing volumes that have more than 10017 cylinders. The coexistence support will

- prevent a device with more than10017 cylinders from being varied online to the system
- prevent a device from coming online during an IPL if it is configured with more than 10017 cylinder.

On an MVS system with coexistence support installed, the following message will be issued if you try to vary online a device with more than 10017 cylinders:

```
IEA434I DEVICE ONLINE IS NOT ALLOWED, GREATER THAN 10017 CYL
```

In IOCP/HCD, define large volumes as UNIT=3390 type devices, or preferably as PAV base devices with UNIT=3390B.

Product specific considerations

ICKDSF

ICKDSF has been enhanced to support larger volume sizes. With the necessary support installed, ICKDSF is able to initialize large volumes.

When you initialize a large volume, be sure to define the VTOC size in proportion to the volume size. The VTOC is to reside within the first 65,535 tracks of the logical volume as it does today. It is highly recommended that you use Indexed VTOC on large volumes.

JES spool data sets

JES spool data set maximum size is 65535 tracks or 4369 cylinders as before. On 3390-3 volumes, the spool could occupy the whole volume, but as you move to larger volumes you will have room on the spool volumes for other data sets. You need to decide whether you want to allow data set allocations on large spool volumes. Review your SMS constructs to allow/disallow these allocations. If you do allow them, check what implications this has on your operational procedures. For example, if your spool volumes are excluded from DFSMShsm migration and incremental backup processing, you need to modify the relevant definitions to have them included.

Note that the spool data set must be allocated within the first 65535 tracks of a volume. With z/OS 1.2 JES spool can reside anywhere on the volume.

DFSMSdss

Large Volume support does not degrade the performance of previously existing functions when the amount of tracks, data sets, or extents being processed is equivalent. For those DFSMSdss functions that process all of the tracks on a volume, physical `FULL` volume dump for example, all of the data sets on a volume, logical copy with `DS(INC(**))` specified for example, or all of the extents on a volume, `DEFRAG` for example, the processing time will increase due to the larger number of tracks, data sets, or extents on large volumes.

Due to the larger number of control blocks that DFSMSdss will need to build in order to represent the larger number of tracks, data sets, and extents on large volumes, the amount of virtual memory that is required will increase whenever DFSMSdss processes any of the following:

- All of the tracks on a volume, as in a physical **FULL** volume **DUMP**
- All of the data sets on a volume, as in a logical **COPY** with **DS(INC(**))**
- All of the extents on a volume, as in a **DEFRAG**.

Stand-alone Restore

The DFSMSdss Stand-alone Restore program has been enhanced to support large volumes. You need generate with the DFSMSdss **BUILDSA** command a new copy of the Stand-alone Restore program (probably on tape) to have the appropriate support included in it. Your old versions of Stand-alone Restore do not support large volumes.

DFSMShsm

DFSMShsm space-management, backup, ABARS, and reporting functions recognize and support the use of large volumes. DFSMShsm supports large volumes for

- SMS storage groups
- Level-0 Non-SMS volumes
- Level-1 (ML1) volumes
- Level-2 (ML2) DASD volumes
- Backup DASD volumes
- Control data sets (MCDS, BCDS, OCDS)
- Journal

DFSMShsm control data set record formats have been updated to include new larger fields for the larger volumes. Since an HSMplex can involve several DFSMShsm releases sharing one MCDS, BCDS, OCDS, and Journal, coexistence PTFs will be required by low level DFSMShsm systems to share control data sets with DFSMShsm systems that support large volumes:

- If a large volume is used for any CDS or the Journal, all systems must have full Large Volume Support (LVS) installed
- Any older releases need the coexistence support to recognize the new Volume (V) record in the MCDS, for purposes of the **LIST** and **REPORT** commands
- The coexistence support will prevent any large volume used for level-0 DASD to be online to older releases.

Large volume, with its larger number of tracks and potentially larger number of data sets, can extend the processing time needed for:

- Volume space management (migration)
- Volume backup and recover
- Volume dump and restore
- AUDIT MEDIACONTROLS

However, total migration, backup, and dump times of your data storage should not change much, as it depends more on the amount of data being processed, and less on the number of volumes.

If planning to use large volumes as ML1 volumes, consider that DFSMShsm recommends using at least one more ML1 volume than the maximum number of concurrent migration tasks. (In the case of Small Dataset Packing, DFSMShsm provides for one SDSP data set per ML1 volume.) If one large volume replaces some number of smaller ML1 volumes, and the resulting number of ML1 volumes is now fewer than the number of migration tasks, performance contention could occur.

Other considerations

The large volume enhancement does not change the size limitations that today exist on certain data set types. Thus, for example:

- The DADSM component of DFSMS/MVS will continue to limit the number of tracks for a data set to 65535 per volume. This applies for example to SAM, PAM (PDS), IMS GSAM, and BDAM data sets. This does not apply to VSAM, extended format, PDSE and HFS data sets.
- DFSORT supports large volumes for input, output, and work devices. The number of records in any DFSORT SORTWK data set cannot exceed 1,000,000 records. The restriction does not apply to DFSORT input and output data sets.
- PDSE data sets can only be single volume. With LVS, you can now allocate larger PDSE than before.
- PPRC, FlashCopy, and XRC support large volumes.
- Small custom volumes may be required for those applications that must still use hardware reserve/release processing, if these are not converted. These custom volumes can be defined in the non-interleaved partition of the rank.

Performance

The performance of large volumes when implemented with PAVs should allow equivalent performance to existing volumes sizes today. With larger volumes configured more PAV alias devices become available for use by z/OS WLM. WLM functions will help to eliminate hot spots on logical volumes by dynamically managing PAV-alias's in order to meet the specified work load goals. ESS functions of Multiple Allegiance, I/O Priority Queuing and striping data across raid ranks will grow in importance in maintaining subsystem performance.

Larger volume sizes will require longer backup and restore times per volume. Physical volume dumps to tape for large volumes will need to be restored to a volume that can accommodate the large volume capacity.

Implementation tasks

Wei15.8(a)-24.7(n)24.2(now r)17.7(e)23.5(v)-8.3(i)3.2(ew)24.2(t)-13.2(he t)11(a)-24.7(si15.8(k)-8.3(s)-8.3(

- Plan PAV alias pool size.
- Confe PAV-alias and PAV-base devices in HCD and ESS.
- Configure large voe in ESS.
- Implement WLM Goal Mode.
- Enable WLM Dynamic Alia Tuning.
- Assess the impact of having large volumes in your disaster recovery / contingency plan.
- Install required vendor products support for large volumes.
- Install coexistence support on pre OS/390 R10 systems.
- Install large volume software support on all sharng systems in the Sysplex.
- Initialize lare volumes for use with ICKDSF INIT function.
- Mirate data to the ESS large volumes based on your storage need.
- Monitor configuration performance.

9.2.7 MIH interval

Th reomnded Mising Interrupt Handler (MIH) value fo ESS logical vous is 30 seonds.

The ESS devices are self describing so entries are not required for base addresses. The ESS will supply a recommended interval of 30 seconds as part of the Read Configuration Data during IPL. This applies for both ESCON and FICON attachment. OS/390 and z/OS will use this information to automatically set its MIH value.

If you need to set a value, ensure you set the same MIH value for all images in your Sysplex that have access to the ESS. The MIH is defined in the IECIOSxx PARMLIB member. The setting should appear as follows: `MIH DEV=(1A00-1A5F),TIME=30.` The operator command to achieve the same is **`SETIOS MIH,DEV=(1A00-1A5F),TIME=30`**.

If any value is defined, this will override the system's determined value. Make sure you don't have in your IECIOS member a global entry such as MIH DASD=mm:ss which applies to all DASD devices.

MIH values are not supported for alias UCBs. Setting an MIH value will cause an error message to be generated: `IOS090I dev IS AN INCORRECT DEVICE`

9.2.8 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 ESS volumes to a new storage group, with associated ACS routines, to take advantage of the higher performance functions available to high activity volumes (Multiple Allegiance and PAV).

DFSMSdss will now invoke the FlashCopy function as described in Section 9.2.10, "ESS Copy Services" on page 166.

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.

9.2.9 Volume initialization

ESS logical volumes need to be initialized with an ICKDSF minimal initialization. A sample job is included in Example 9-3.

Example 9-3 ICKDSF minimal init

```
//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) VOLID(MVS203) OWNER(SYSPROG) -
       VTOC(0,1,29) INDEX(2,0,15) NOVALIDATE NOCHECK
 //
```

A PAV alias cannot be initialized. Aliases are only known to the I/O Supervisor routines of the z/OS.

When you initialize custom volumes, adjust the VTOC and VTOC Index size to correspond to the volume size. You can check the volume size on the ESS Specialist *S/390 Storage* panel (Figure 9-2). The same also applies to the VVDS.

9.2.10 ESS Copy Services

In this section we briefly describe OS/390 and z/OS support for ESS Copy Services. ESS Copy Services implementation tasks are described to a greater extent in Chapter 12, "ESS Copy Services for S/390" on page 197 and in the IBM Redbook *Implementing ESS Copy Services on S/390,* SG24-5680.

FlashCopy

The DFSMSdss **COPY FULL** command will automatically invoke the ESS FlashCopy service. DFSMSdss performs the copy using FlashCopy service if the following conditions are met:

- The source and target volumes support the FlashCopy service.
- The source and target volumes reside on the same LSS.
- The source and target volumes must be online.
- The source and target volumes are the same device type.
- The source volume does not have a FlashCopy relationship with another volume.
- The target volume is not in a FlashCopy session with another volume.
- The target volume is the same size or larger than the source volume.
- The target volume is not an XRC primary volume.
- The target volume is not a PPRC primary or secondary volume.
- The target volume is not the source of an in-progress, concurrent copy operation.

If these conditions are not met, but the CC keyword is specified, DFSMSdss provides the standard concurrent copy service. If these conditions are not met and the CC keyword is not specified, DFSMSdss uses the standard I/O service to copy the volume.

A sample job stream to invoke FlashCopy is shown in Example 9-4.

Example 9-4 DFSMSdss COPY FULL job

```
//COPYFULL JOB . . . . .
//*
//INSTIMG  EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//SRCVOL   DD
//TRGVOL   DD
//SYSIN    DD *
  COPY FULL INDYNAM((SRCVOL)) OUTDYNAM((TRGVOL)) COPYVOLID
/*
```

PPRC and XRC

The PPRC and XRC functions are defined and operated much the same way as was done with the previous IBM 3990 Model 6 and RAMAC storage controllers. There are some additional features to consider, for example, XRC support of the suspend/resume commands.

To invoke the PPRC or XRC function, you need to have ordered the appropriate features (XRC FC 181x, PPRC FC 182x) on your ESS.

The establishing and managing of the PPRC and FlashCopy functions is done using TSO commands or using the ESS Copy Services Web interface.

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 the conditions for using the FlashCopy service are met, FlashCopy will be automatically invoked instead of concurrent copy.

9.3 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 (LSS).

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.

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 (LCU) and volume/device definitions within the ESS. A sample IOCP deck is shown in Example 9-1 on page 148.

Mismatched definitions between the ESS internal logical configuration and VM/ESA can be displayed via the **QUERY PAV** command.

9.3.1 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 Section , “Additional considerations” on page 159.

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.

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
```

9.3.2 ESS 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 managed in the ESS in the same way as with previous storage controllers using the ICKDSF utility. With the ESS, in a VM/ESA environment you can also establish and manage the FlashCopy and PPRC functions using the ESS Copy Services Web 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.

9.4 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.

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.

9.4.1 ESS 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 managed in the ESS in the same way as with previous storage controllers using the ICKDSF utility. With the ESS, in a VSE/ESA environment you can also establish and manage the FlashCopy and PPRC functions using the ESS Copy Services Web interface.

9.5 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 (LSS).

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.

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 (LCU) and volume/device definitions within the ESS.

9.5.1 PPRC and FlashCopy

PPRC and FlashCopy are optional features on the ESS Models F10 and F20 and require ESS LIC level 1.5.0. or later. These functions are optional features on the ESS (FC 182x for PPRC and 183x for FlashCopy). TPF must be Version 4 Release 1.

In a TPF environment the FlashCopy and PPRC functions are established and managed using the ESS Specialist Copy Services Web interface.

9.6 Linux software configuration

On zSeries 900 and on S/390 servers you can run Linux natively as a stand-alone or as a logical partition (LPAR). In addition, the S/390 Virtual Image Facility for Linux, and z/VM V4 enable you to run more Linux images than can be deployed using LPARs and provide capabilities to help create and manage these images.

Some basic considerations about the different configurations that can be used when running Linux on a zSeries server are shown in Table 9-4.

Table 9-4 Linux implementations in zSeries servers

Linux for S/390 implementation	Reasonable number of Linux images	Considerations
Native	1	With the exception of the smaller platforms (as the P/390) this is not economically feasible
LPAR	15	Useful if a very small number of Linux images is needed for testing and development
VIF	Tens to hundreds	Management of Linux images becomes difficult when greater than a few tens
z/VM	Hundreds to thousands	Probably the most cost-effective way to run and manage Linux on the mainframe

Current Linux for S/390 distributions supported with the ESS are SuSE Linux Enterprise Server for S/390 and Turbo Linux Server 6 for zSeries and S/390. For the most current information on the Linux for S/390 releases and distributions that are supported with the ESS, refer to: `http://www.storage.ibm.com/hardsoft/products/ess/supserver.htm`

Native hardware

Linux can be run natively on a zSeries server, or a S/390 server that has *Relative* and *Immediate* instruction support. This is not the most recommend because only one Linux system can be run at any one time.

Logical partition

Linux images can be run in S/390 logical partitions (LPARs) — up to 15 different logical partitions. This configuration has the ability to run and manage up to 15 individual Linux systems. Perhaps more compelling is the ability to run a number of Linux partitions concurrent with z/OS or VSE/ESA partitions providing high value services such as transactions and access to existing applications and data.

z/VM guest support

Linux runs as a guest operating system in one or more VM Virtual Machines. z/VM uses architectural hardware functions in S/390 to virtualize the S/390 instruction set, making each guest think it has its own dedicated server processor. Any number of Linux or other guests, including z/OS, VSE/ESA and z/VM CMS Interactive System can run concurrently. The number of concurrent guests is limited by the resources available to virtualize them.

This configuration has the ability to use devices, both virtual and real, supported by z/VM and not directly supported by Linux, and this is a plus. You will be able to share more efficiently the resources among the z/VM users instead of having to dedicate them. For communication between guest virtual machines, this can be done at extremely high speed.

Running IBM's industry leading z/VM virtualization technology, the zSeries offering for Linux offers an ideal platform for server consolidation, utilizing the mainframe's ability to create as few as 20 and up to hundreds of virtual Linux servers on a single physical box, saving server users substantially on energy, floor space, and maintenance expense.

Virtual Image Facility (VIF)

The S/390 Virtual Image Facility for Linux is an entry level product that offers a complete server environment for multiple Linux systems on a single zSeries (or S/390) server.The Virtual Image Facility is an easy-to-use, high-performance environment that operates within a logical partition or in native S/390 mode and provides the capability to create a significant number of Linux images. An internal network provides high-speed communication among Linux images.

Integrated Facility for Linux (IFL)

z/VM V4R2 users can make use of the IBM Integrated Facility for Linux (IFL) to run multiple Linux images on an IFL-configured z/VM system without affecting the software fees of their existing zSeries 900 or S/390 environment.

The heart of zSeries and S/390 hardware is the Multi-Chip Module (MCM) which contains up to 20 Processing Units (PU), commonly referred to on other platforms as CPUs or engines. All PUs are identical, but can take on different roles in the system. The IFL engine is a new type of PU which allows additional engines to be dedicated to z/VM or Linux workloads. It is lower priced than standard engines and allows traditional S/390 software charges to remain the same. IFL is based on PR/SM LPAR technology and it can be dedicated to a single logical partition or shared by multiple logical partitions.

Note: Installations should note that the Linux workspace enabled by the IFL facility will not support any of the S/390 traditional operating systems (z/OS, OS/390, TPF, VSE, or VM). Only Linux applications or Linux operating in conjunction with z/VM or VIF are supported by the IBM Integrated Facility for Linux.

For detailed step-by-step guidelines to implement Linux for S/390 we recommend that you read the following IBM publications: *Implementing Linux with IBM Disk Storage,* SG24-6261; *Linux for S/390,* SG24-4987; and *Linux for IBM zSeries and S/390: Distributions,* SG24-6264.

10

Data migration in the zSeries environment

This chapter gives you the information necessary to plan the methods and tools your installation will be using when doing the migration of the existing data into the IBM TotalStorage Enterprise Storage Server. The information presented here is basically oriented to S/390 users who must move data from non-ESS storage subsystems into a ESS storage subsystems.

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 plan for non-disruptive methods of migration rather than the disruptive ones. In addition, it does not matter what method will be used for the migration, it must guarantee total data integrity.

Storage consolidation

The ESS is a very scalable subsystem that ranges from 420 GB (eight arrays of 9.1 GB DDMs) to 22.4 Terabytes (48 Arrays of 72.8 GB DDMs). The 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 data (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 methods

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 using IBM migration services offerings.

10.2 Data migration for MVS environments

There are two major considerations when planning data migration for MVS environments.

Systems-Managed Storage (SMS) environments

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 environments

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`**.

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 3390). To move a volume logically, use the DFSMSdss **COPY DATASET** command.

DFSMSdss **DUMP** and **RESTORE** commands move data from source volumes to tape 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.

See the *z/OS DFSMShsm Storage Administration Guide,* SC35-0421, and *z/OS DFSMShsm Storage Administration Reference,* SC35-0422.

DFSMShsm ABARS

Although its primary purpose is disaster recovery, you can also use DFSMShsm ABARS to migrate volumes to cartridges and restore them in an ESS. In the Aggregate Group you have to define the data sets you 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 details about XRC characteristics and implementation, see *Planning for IBM Remote Copy,* SG24-2595 and *Implementing ESS Copy Services on S/390,* SG24-5680

10.3 Data migration for VM environments

DFSMS/VM contains services which include a data mover, an automated move process, and 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 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. Back up and restore VSAM files.
3. Back up and restore ICCF libraries.
4. Back up and restore the system history file.
5. Bac kup 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 these capabilities:

- 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 Data migration across ESSs

The information in the previous sections is oriented to the data migration scenarios where data from non-ESS storage subsystems is moved (or copied) to ESS storage subsystems.

If your data migration has to do with moving data from one ESS to another ESS you still can use the techniques described so far. But for this ESS to ESS migration scenarios you also have the additional alternative of using the ESS Copy Services functions. More precisely you can take advantage of the optimized ESTABLISH and RESYNCH operations over PPRC pairs of logical volumes at long distances using the INRANGE Fibre Channel Directors (see Section , “INRANGE” on page 22).

10.6 IBM migration services

This is the easiest way to migrate data, because IBM will assist you throughout the complete migration process. In several countries IBM offers a migration service. Check with your IBM sales representative about migration services for your specific environment and needs.

11

Managing and monitoring the ESS

In this chapter we present the tools you will be using for managing and monitoring the IBM TotalStorage Enterprise Storage Server in the S/390 storage environments. Because for S/390 environments, the ESS emulates 3990 storage controls with 3390 (or 3380) storage devices attached, the basic management and monitoring processes and tools that have been used with previous 3990 and RAMAC storage controllers are also applicable to the ESS. However, there are some changes relating to the specific characteristics and functions of the ESS which are covered in this chapter.

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 set up 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 Managing tools

Typically the management of the ESS consist of activities like configuration changes, storage capacity re-allocations, enablement-disablement of ESS resources and functions. The IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) program and the Control Unit Initiated Reconfiguration (CUIR) function will assist you in these activities.

11.2.1 IBM TotalStorage Enterprise Storage Server Specialist

The tool you will be using most of the time for managing the ESS is the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist). Any configuration related activity, like changes related to server attachment, or changes related to storage capacity assignment, will need the use of the ESS Specialist. Chapter 5, "IBM TotalStorage Enterprise Storage Server Specialist" on page 79, describes this management tool.

11.2.2 Control Unit Initiated Reconfiguration (CUIR)

The Control Unit Initiated Reconfiguration (CUIR) function automates channel path quiesce/resume actions to minimize manual operator intervention during selected ESS service or upgrade actions.

CUIR allows the IBM Service Support Representative (IBM SSR) to request that all attached system images set all paths associated with a particular service action offline. System images with the appropriate level of software support will respond to such requests by varying off the affected paths, and either notifying the ESS subsystem that the paths are offline, or that it cannot take the paths offline. CUIR reduces manual operator intervention and the possibility for human error during maintenance actions, at the same time reducing the time required for the maintenance. This is particularly useful in environments where there are many systems attached to an ESS.

CUIR is supported on zSeries and S/390 platforms and requires ESS Licensed Internal Code (LIC) level 1.5.0, or later. CUIR support is not available for open systems. Refer toSection 7.5, "CUIR support" on page 100 for information regarding CUIR software support.

The CUIR function can be enabled/disabled on each ESS subsystem selectively. The IBM SSR will enable or disable CUIR during the initial ESS install sequence based on information provided by the customer on the Communication Resources work sheet (option *Allow CUIR to automatically vary paths off and on*). See Section 8.3.3, "Communication Resources worksheet" on page 117. The options are:

- **Enable:** The ESS will initiate the reconfiguration for service.
- **Disable:** The operator will set the paths offline for service.

The IBM SSR can later change the original configuration option setting.

When CUIR is enabled, the ESS invokes CUIR as part of selected service processes. When the service operation reaches a stage where it becomes necessary to take a disruptive action that would involve paths being disabled, the IBM SSR can request a quiesce of the channel paths from the ESS subsystem. ESS sends a reconfiguration request to all attached operating systems requesting that they take the appropriate paths offline.

When an operating system receives the request from the ESS, it determines the paths affected by the request, waits until in-process I/O operations are complete, and then takes the appropriate paths offline, if not so already, and marks them as in use by CUIR. CUIR will not take the last path offline from any processor image to an online device.

The following MVS messages are issued in response to a successful quiesce request (see Figure 11-1).

```
IOS275I  C.U.I.R. REQUEST TO QUIESCE THE FOLLOWING PATH(S):
           CHPID A0 TO DEVICE(S) 2E00-2E3F
IOS281I  C.U.I.R. REQUEST SUCCESSFUL
```

Figure 11-1 MVS messages for quiesce request

When the paths are offline, the host system acknowledges to the ESS that the paths have now been quiesced and may be disabled for the service action. After ESS has received positive response from all the affected systems, the IBM SSR can perform the service action. If the quiesce request was unsuccessful, the process notifies the SSR to request the system operator to issue the host commands to vary the necessary channel paths offline. Possible reasons for a quiesce request to fail are:

- No response from the host system
- The host does not support CUIR
- The request would take the last path offline to a device.
- The vary path process failed

When the paths are no longer required for the service action, the IBM SSR resumes the I/O components that were quiesced. The ESS sends a resume request to all attached host systems to bring the paths back online. The operating system determines which channel paths are affected and varies them online automatically.

MVS issues the following messages in response to a successful resume request (see Figure 11-2):

```
IOS278I  C.U.I.R. REQUEST TO RESUME THE FOLLOWING PATH(S):
           CHPID A0 TO DEVICE(S) 2E00-2E3F
IOS281I  C.U.I.R. REQUEST SUCCESSFUL
```

Figure 11-2 MVS messages for resume request

A path quiesced by CUIR is logically offline to the host. Use the **DEVSERV PATHS** command to display path status. Figure 11-4 on page 188 shows the output of the **DSERV PATHS** command.

You cannot vary online channel paths that are taken offline by CUIR. In MVS, an attempt will fail with message IEE169I (see Figure 11-3).

```
VARY PATH(2E00,A4),ONLINE

IEE169I VARY REJECTED, PATH(2E00,A4) OFFLINE DUE TO C.U.I.R.
```

Figure 11-3 Vary path online rejected

If you know that the path is no longer being serviced, use MVS command **VARY PATH((dev1-dev2),xx),ONLINE,FORCE** to vary the path online. It is possible that the path was never varied back online by the CUIR service after it completed a service action involving the path. This situation could occur if the service representative neglected to resume the path after finishing the service action or if the resume failed.

CUIR supports both ESCON and FICON channel paths.

CUIR does not have automatic quiesce support for ESCON links used for PPRC. If a service action requires a PPRC port to be disabled, the service representative will disable it manually using the service terminal. PPRC will continue to operate with the remaining paths. When the service action is complete, the IBM SSR will enable the port and PPRC restarts using the link. If you only have a few PPRC paths defined, quiescing one of them may impact performance. If possible, add a temporary link with the PPRC **ESTPATH** command in place of the affected link. Otherwise, schedule the service action at a time of low write activity.

Example 11-1 is a list of MVS messages related to CUIR. For details, please see *z/OS MVS System Messages, Vol 9,* SA22-7639.

Example 11-1 MVS messages related to CUIR

```
IOS275I C.U.I.R. REQUEST TO QUIESCE THE FOLLOWING PATH(S): CHPID xx TO
         DEVICE(S) dev,dev1-dev2, ... CHPID yy TO DEVICE(S) dev,dev1-dev2, ...
IOS278I C.U.I.R. REQUEST TO RESUME THE FOLLOWING PATH(S): CHPID xx TO
         DEVICE(S) dev,dev1-dev2, ... CHPID yy TO DEVICE(S) dev,dev1-dev2, ...
IOS281I C.U.I.R. REQUEST SUCCESSFUL
IOS283I C.U.I.R. VARY PATH(dev,xx) REJECTED, text
IOS284I C.U.I.R. REQUEST REJECTED - VARY COMMAND PROCESSOR FAILED
IOS290I C.U.I.R. REQUEST UNSUCCESSFUL
```

11.3 Monitoring tools

Several tools help you in the activity of monitoring the ESS. As already mentioned, most of these tools work basically in the same manner as they did with the previous IBM disk storage controllers (IBM 3990s, RAMACs). In the following sections we describe these tools, highlighting the ESS specifics:

- RMF
- CRR and DFSMS Optimizer
- SMF
- SIM and EREP
- IDCAMS
- MVS system commands

IBM StorWatch Enterprise Storage Server Expert

The IBM StorWatch Enterprise Storage Server Expert (ESS Expert) is a managing and monitoring tool specifically oriented for the ESS administration. The ESS Expert is covered in Chapter 6, “IBM StorWatch Enterprise Storage Server Expert” on page 87 of this redbook.

11.3.1 RMF

RMF reports regarding ESS performance are almost the same as for previous IBM 3990 storage controllers. Enhancements have been provided for PAV reporting, and FICON channel reporting.

Device Activity report

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 that have been used by a device.

Example 11-2 shows an extract from an RMF Device Activity report. The PAV column shows the number of PAV *exposures* (base plus aliases) for the device at the end of the reporting interval. If the number has been changed during the reporting interval, it is followed by an asterisk (*).

Example 11-2 RMF Device Activity report

```
                              D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

             z/OS V1R2                  SYSTEM ID XXXX              START 11/13/2001-09.00.00
                                        RPT VERSION V1R2 RMF        END   11/13/2001-09.30.00

TOTAL SAMPLES =  3,600   IODF = 02   CR-DATE: 09/20/2001   CR-TIME: 15.48.19     ACT: POR

                                           DEVICE   AVG  AVG   AVG  AVG  AVG   AVG  AVG  AVG
STORAGE  DEV  DEVICE    VOLUME PAV  LCU  ACTIVITY RESP IOSQ  DPB  CUB  DB    PEND DISC CONN
 GROUP   NUM  TYPE      SERIAL           RATE     TIME TIME  DLY  DLY  DLY   TIME TIME TIME

SYSTEM   9000 33909     SYS201   4  001C   4.887   1.5  0.0        0.0  0.0   0.3  0.2  1.0
GROUP01  9001 33909     PUB201   4  001C   0.661   1.3  0.0        0.0  0.0   0.3  0.3  0.7
GROUP01  9002 33909     PUB202   4  001C   0.206   1.5  0.0        0.0  0.0   0.3  0.5  0.7
GROUP01  9003 33909     PUB203   4  001C  11.675   1.8  0.0        0.0  0.0   0.3  0.8  0.7
GROUP01  9004 33909     PUB204   4  001C   6.689   1.1  0.0        0.0  0.0   0.3  0.3  0.5
GROUP01  9005 33909     PUB205   4  001C   0.132   2.4  0.0        0.0  0.0   0.3  1.3  0.8
GROUP01  9006 33909     PUB206   4  001C   0.066   4.3  0.0        0.0  0.0   0.3  2.3  1.7
WORK     9007 33909     WRK201   4  001C   0.001   0.8  0.0        0.0  0.0   0.4  0.0  0.4
GROUP01  9008 33909     PUB207   4  001C   0.359   7.6  0.0        0.0  0.0   0.4  3.4  3.8
GROUP01  9009 33909     PUB208   4  001C   0.004   3.3  0.0        0.0  0.0   0.3  2.7  0.4
GROUP01  900A 33909     PUB209   4  001C   0.214   2.8  0.0        0.0  0.0   0.3  2.0  0.5
```

Additional information about usage of alias addresses can be found in SMF record, type 74, subtype 5. RMF does not use this information.

Besides the Device Activity report showing the number of PAV exposures in its PAV column, also the following RMF reports have an additional column to show the number of aliases:

- Mon I / Postprocessor: DASD Activity report, Shared DASD Activity report
- Mon II: DEV report, DEVV report
- Mon III: DEVR report, DSNV report

Workload Activity report

The RMF Workload Activity report has been enhanced to show the I/O Priority Management status and WLM Dynamic Alias Management status on the Service Policy Page (see Example 11-3).

Example 11-3 RMF Workload Activity report

```
                             W O R K L O A D   A C T I V I T Y

z/OS V1R1    SYSPLEX PLEX2       START 11/14/2001-09.00.00 INTERVAL 000.30.01  MODE = GOAL
    CONVERTED TO z/OS V1R2 RMF   END   11/14/2001-09.30.00

                            POLICY ACTIVATION DATE/TIME 11/09/2001 12.15.48
                                        - SERVICE POLICY PAGE -

SERVICE DEFINITION: servdef1  WLM definitions            -SERVICE DEFINITION COEFFICIENTS-
  INSTALL DATE: 11/09/2001 12.15.35  INSTALLED BY: XXXXX      IOC     CPU     SRB     MSO
  POLICY: FULLTIME  Standard WLM definition
  I/O PRIORITY MANAGEMENT:  YES                                0.5     1.0     1.0  0.0000
  DYNAMIC ALIAS MANAGEMENT: YES
```

FICON Channel Path reporting

RMF support for FICON channels (both FCV and FC) includes the reporting of five new measurements:

- Bus utilization
- Percentage of bus cycles the bus has been found busy for this channel in relation to the theoretical limit.
- Read bandwidth for an image in MB/sec
- Data transfer rates from the control unit to the channel for this individual logical partition.
- Total read bandwidth (for all images on the system) in MB/sec
- Data transfer rates from the control unit to the channel for the entire system.
- Write bandwidth for an image in MB/sec
- Data transfer rates from the channel to the control unit for this individual logical partition.
- Total write bandwidth (for all images on the system) in MB/sec
- Data transfer rates from the channel to the control unit for the entire system.

Sample RMF Channel Path Activity report information is shown in Example 11-4.

Example 11-4 RMF Channel Path Activity report

```
                      RMF V1R2   CHANNEL PATH ACTIVITY         LINE 25 OF 47
COMMAND ===>                                                SCROLL ===> CSR

SAMPLES: 119     SYSTEM: XXXX  DATE: 11/12/01  TIME: 12.44.00  RANGE: 120   SEC

 CHANNEL PATH       UTILIZATION(%)   READ(B/S) WRITE(B/S)    MSG  MSG SEND RECV
ID NO  G  TYPE  S   PART  TOT  BUS   PART  TOT  PART  TOT   RATE SIZE FAIL FAIL

90        FC    Y    3.4  7.0 15.4   872K   2M   88K 406K
91        FC    Y    3.5  7.0 15.3   870K   2M   96K 415K
92        FC    Y    3.4  7.0 15.4   875K   2M   94K 412K
A0        FC    Y    2.6  7.0 15.7   386K   2M   93K 635K
A1        FC    Y    2.6  7.2 15.8   430K   2M   95K 632K
A2        FC    Y    2.6  7.2 15.8   426K   2M   93K 639K
A3        CNC_P Y    1.1  1.1
B0        FC    Y    0.2  1.7 13.5    48K 988K   56K 235K
B1        FC    Y    0.2  1.7 13.5    56K 995K   54K 223K
B2        FC    Y    0.2  1.6 13.5    56K   1M   42K 213K
C0        FC    Y    0.1  1.1 13.0    21K 665K   27K 115K
C1        FC    Y    0.1  1.0 13.0    23K 654K   17K  73K
C2        FC    Y    0.1  1.1 13.0    19K 665K   27K  84K
```

11.3.2 CRR and DFSMS Optimizer

Two tools also have available for monitoring the ESS performance are the Cache RMF Reporter (CRR) that collects data and produces its own reports, and the DFSMS Optimizer. CRR also collects certain type of records for DFSMS Optimizer reporting.

Cache RMF Reporter

The Cache RMF Reporter (CRR), that is part of RMF, has been enhanced to produce a new RAID Rank Activity report, which shows I/O statistics for the RAID ranks in the ESS. The sample output in Example 11-5 shows measurements for two RAID ranks, 0000 and 0001.

Example 11-5 CRR RAID Rank Activity report

```
-------------------------------------------------------------------------------------------
                               RAID RANK ACTIVITY
-------------------------------------------------------------------------------------------
ID    RAID     DA   HDD    -------- READ REQ -------     ------- WRITE REQ -------
      TYPE                 RATE  AVG MB  MB/S  RTIME     RATE  AVG MB  MB/S  RTIME
*ALL                14       94   0.037   3.5     17       31   0.023   0.7    101
0000  RAID-5   01    7       85   0.039   3.3     18       20   0.025   0.5    108
0001  RAID-5   01    7        9   0.018   0.2     12       11   0.020   0.2     87
```

DFSMS Optimizer

The DFSMS Optimizer collects and stores historical system activity data in a database. The extraction process captures data from SMF, RMF, HSM, and the cache RMF reporter, or CRR (the data collection component of CRR is distributed with the DFSMS Optimizer), and compresses this data by 80% to 90% into the DFSMS Optimizer database, allowing you to store months of data at a time. The database is used for DFSMS Optimizer analyses and reports.

A major capability of the DFSMS Optimizer is its ability to analyze and simulate your storage environment based on its database. The analyses possible are:

- Hardware performance analysis
- Data set performance analysis
- DFSMShsm performance analysis
- Management class analysis

The two first reports can help you when analyzing the storage subsystem performance:

- **Hardware performance analysis:** The DFSMS Optimizer provides performance information about your storage subsystem in the form of reports and charts to help you isolate the cause of hardware performance problems such as device skew, I/O bottlenecks, or I/O intensive periods. This information is used to balance the I/O workloads across your subsystem.
- **Data set performance analysis:** The DFSMS Optimizer uses SMF and RMF data to provide data-set-level statistics on the use of cache. Using cost components that you provide, it reports the cost of cache resources associated with the storage class policies and the data sets to which those policies apply.

11.3.3 SMF

SMF record type 99 subtype 7 contains summary information on alias usage for an ESS with Parallel Access Volume (PAV) definitions enabled. A subtype 7 record is written every third policy interval. Turn on SMF type 99 recording if desired; this data can be used for analysis of alias management activity if it becomes necessary. Refer to the *z/OS System Managemet Facilities (SMF),* SA22-7630 for a description of the record layout.

11.3.4 SIM and EREP

The Environmental Recording, Editing, and Printing (EREP) program provides problem incident reports for the new device type 2105. However, the ESS maintenance strategy does not rely on the analysis of data in EREP reports. Service Information Messages (SIMs) and problem records in the ESS internal problem log contain information necessary for the IBM SSR to start a service action.

Sense data records for some ESS temporary and all permanent errors are sent from the ESS to the system to give information necessary to perform needed system error recovery procedures. The ESS sense data is logged in the error-recording data set (ERDS) in the system, but it is not used for ESS problem determination by the IBM SSR. The ESS generates and send to the host SIMs whenever ESS service is needed. The SIMs summarize the service information necessary to isolate and repair ESS error conditions. SIMs are presented as host system console messages and are logged in the ERDS. The SIM ID is the same as the problem number in the ESS problem log.

Even if EREP reports are not generally required for ESS maintenance, you may still want to run the EREP Service Information Message report (see Example 11-6) on a regular basis as it gives a compact summary of SIMs generated during the reporting period and serves as a checklist of problem incidents. Other system exception reports might contain ESS information. The other reports may occasionally be required by the IBM SSR so you should be prepared to run the reports if necessary.

Example 11-6 Service Information Messages report

```
SERVICE INFORMATION MESSAGES                        REPORT DATE 024 99
                                                    PERIOD FROM 021 99
                                                             TO 022 99
*****************************************************************************************
COUNT   FIRST OCCURRENCE              LAST OCCURRENCE
1       021/99 17:44:27:78            021/99 17:44:27:78
        MODERATE ALERT 2105-E20     S/N 0113-10473 REFCODE C211-1060-A00A ID=03
        DASD EXCEPTION ON SSID 0011
        ADDITIONAL ANALYSIS REQUIRED TO DETERMINE REPAIR IMPACT.
        SEE PROBLEM NUMBER 03 FOR DETAILS

2       021/99 19:24:19:56            021/99 19:24:19:56
        SERVICE ALERT 2105-E20     S/N 0113-30224 REFCODE 4320-0000-5284 ID=06
        MEDIA EXCEPTION ON SSID 00D2, VOLSER 380050 DEV 0E12, 0D
        REFERENCE MEDIA MAINTENANCE PROCEDURE 2

3       021/99 19:24:04:67            022/99 03:29:01:65
        SERIOUS ALERT 2105-E20     S/N 0113-10473 REFCODE C211-1060-A00A ID=09
        DASD EXCEPTION ON SSID 00D2
        ADDITIONAL ANALYSIS REQUIRED TO DETERMINE REPAIR IMPACT.
        SEE PROBLEM NUMBER 09 FOR DETAILS
```

11.4 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.

The **LISTDATA** command now generates a new RAID Rank Counters report and a modified Subsystem Counters report. **SETCACHE**, on the other hand, has been updated to produce messages indicating that ESS does not support the 3990 Dual Copy function or the disabling of Cache or DASD Fast Write (DFW).

11.4.1 IDCAMS LISTDATA

IDCAMS **LISTDATA** command generates reports on activity within all IBM caching models of 3880, 3990, and RAMAC storage controllers.

With the ESS you can also use IDCAMS **LISTDATA** to request activity information. You will receive a modified Subsystem Counters report and a new RAID Rank Counters report when you issue the **LISTDATA COUNTS SUBSYSTEM** or the **LISTDATA COUNTS ALL** command. **LISTDATA COUNTS DEVICE** command generates the Subsystem Counters report, but not the RAID Rank Counters report.

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. Example 11-7 presents a Subsystem Counters report including this identifier. The command used in this case was **LISTDATA COUNTS VOLUME(volser) UNIT(3390) SUBSYSTEM.**

Example 11-7 IDCAMS Subsystem Counters report

```
IDCAMS  SYSTEM SERVICES
                         2105 STORAGE CONTROL
                       SUBSYSTEM COUNTERS REPORT
                    VOLUME SYSXXX DEVICE ID X'20'
                         SUBSYSTEM ID X'5100'
                                      CHANNEL OPERATIONS
                   ......SEARCH/READ.....  ..............WRITE...............
                        TOTAL  CACHE READ       TOTAL      DASDFW CACHE WRITE
REQUESTS
  NORMAL              5882213     5811266       17158       17058       17058
  SEQUENTIAL          4370514     4364583      194108      194050      194050
  CACHE FAST WRITE          0           0           0         N/A           0
TOTALS               10252727    10175849      211266      211108      211108
REQUESTS                   CHANNEL OPERATIONS
  INHIBIT CACHE LOADING                   0
  BYPASS CACHE                            0
TRANSFER OPERATIONS             DASD/CACHE   CACHE/DASD
  NORMAL                            101440       851628
  SEQUENTIAL                        926668          N/A
DASD FAST WRITE RETRIES                   0
DEVICE STATUS     CACHING:         ACTIVE
                  DASD FAST WRITE: ACTIVE
                  DUPLEX PAIR:     NOT ESTABLISHED
RAID RANK ID X'0001'

                       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
:
: (Legend removed)
:
RAID RANK ID       - THE RAID RANK ID ASSOCIATED WITH THIS VOLUME
```

IDCAMS RAID Rank Counters report

The RAID Rank Counters report is produced to support ESS. You can generate this report by issuing the **LISTDATA COUNTS SUBSYSTEM** or the **LISTDATA COUNTS ALL** command. Generated after the Subsystem Counters report, the RAID Rank Counters report contains summary data for the whole RAID rank. Example 11-8 shows a RAID Rank Counters report (***Note***: due to hardware restrictions, RAID Rank data is only available when the device with the lowest unit address in a RAID rank is online).

Example 11-8 IDCAMS RAID Rank Counters report

```
IDCAMS  SYSTEM SERVICES
                         2105 STORAGE CONTROL
                       SUBSYSTEM COUNTERS REPORT
                          RAID RANK COUNTERS
                         SUBSYSTEM ID X'5100'
RAID RANK ID X'0001'
DEVICE ADAPTER ID X'01'
NUMBER OF HDDS IN RAID RANK       8
HDD SECTOR SIZE                 524
                                     RAID RANK OPERATIONS
REQUESTS                  I/O REQUESTS    RESP/TIME    SECTOR REQUESTS
```

```
  READ                        25770871        0.012         2342659927
  WRITE                        7982700        0.111         1315104518

                       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
```

11.4.2 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, on the ESS Cache and DFW are ON by default, and you are not allowed to modify them. In addition, ESS does not support the Dual Copy function. You will receive message IDC31562I with return code 12 indicating that a parameter is not available if you issue any of the following Cache/DFW or Dual Copy commands:

- SETCACHE SETSECONDARY
- SETCACHE SUSPENDPRIMARY
- SETCACHE SUSPENDSECONDARY
- SETCACHE RESETTODUPLEX
- SETCACHE REESTABLISHDUPLEX
- SETCACHE RESETTOSIMPLEX
- SETCACHE DFW OFF
- SETCACHE NVS OFF
- SETCACHE DEVICE OFF
- SETCACHE SUBSYSTEM OFF

Example 11-9 displays the message given in response to the **SETCACHE SETSECONDARY** command.

Example 11-9 SETCACHE response message

```
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
```

11.5 MVS system commands

MVS system commands **DEVSERV** and **DISPLAY** have been enhanced in support of ESS.

11.5.1 DEVSERV command

In z/OS, the **DEVSERV** command can be used to obtain status information about DASD and tape subsystems. The command 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.

DEVSERV PATHS

The **DEVSERV PATHS** command (**DS P**) displays useful information on the state of a device and the paths used to access it from the host. Figure 11-4 shows the output (IEE459I message) for the **DEVSERV PATHS** command for three devices in an ESS, beginning at address 2E00 (**DS P,2E00,3**).

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 a path in the example indicates successful communication. The dynamic legend at the bottom defines the symbols that appear in the output.

```
DS P,2E00,3

IEE459I 16.00.12 DEVSERV PATHS 340
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
     RTYPE   SSID CFW TC   DFW   PIN  DC-STATE CCA  DDC   ALT  CU-TYPE
2E00,33903 ,O,000,MC2E00,A4=* A5=+ A6=+ A7=+ F8=+ F9=< FA=+ FB=+
     2105    6290  Y  YY.  YY.    N   SIMPLEX   00   00       2105
2E01,33903 ,O,000,MC2E01,A4=* A5=+ A6=+ A7=+ F8=+ F9=< FA=+ FB=+
     2105    6290  Y  YY.  YY.    N   SIMPLEX   01   01       2105
2E02,33903 ,A,000,MC2E02,A4=* A5=+ A6=+ A7=+ F8=+ F9=< FA=+ FB=+
     2105    6290  Y  YY.  YY.    N   SIMPLEX   01   01       2105
************************ SYMBOL DEFINITIONS ************************
A = ALLOCATED                  O = ONLINE
+ = PATH AVAILABLE             * = LOGICALLY OFF, PHYSICALLY ON
< = PHYSICALLY UNAVAILABLE
```

Figure 11-4 DEVSERV PATHS output

An important attribute of the **DEVSERV PATHS** command is that the state of the subsystem is 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 DASD device. Figure 11-5 shows the output of a **DS QD** to an ESS subsystem.

```
DS QD,9000,3

IEE459I 13.31.27 DEVSERV QDASD 461
UNIT VOLSER SCUTYPE DEVTYPE   CYL  SSID SCU-SERIAL DEV-SERIAL EF-CHK
9000 SYS201 2105F20 2105    10017  2100 0175-12345 0175-12345 **OK**
9001 PUB201 2105F20 2105    10017  2100 0175-12345 0175-12345 **OK**
9002 PUB202 2105F20 2105    10017  2100 0175-12345 0175-12345 **OK**
****      3 DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

Figure 11-5 DEVSERV QUERY DASD output

The **DEVSERV QUERY DASD** command with parameter UCB displays the device UCB as shown in Figure 11-6.

```
DS QD,9000,UCB

IEE459I 13.32.00 DEVSERV QDASD 505
UNIT VOLSER SCUTYPE DEVTYPE   CYL  SSID SCU-SERIAL DEV-SERIAL EF-CHK
9000 SYS201 2105F20 2105    10017  2100 0175-12345 0175-12345 **OK**
  UCB AT V00E58FA0
00A8FF8C90000000 0000000008E4C3C2 3030200F00E58F78 000F0100E2E8E2F2
F0F1104800A00197 00E58DA0020209C8 0080030300000000
  UCB PREFIX AT V0211BDB0
000E804000000000 0000000000010583 289C001FE00080E0 9091920000000000
0108000000000001
  UCB COMMON EXTENSION AT V00E58F78
0000094020AA0008 0211BDB00000011D 0000000000FBE93C 00E58F4000000000
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

Figure 11-6 DEVSERV QUERY DASD UCB output

Examples of other data areas that can be displayed using the **DEVSERV QUERY DASD** command are DCE, DCPT, and SSSCB. Other **QD** parameters such as RDC (read device characteristics), RCD (read configuration data), and SNSS (sense subsystem status) return information directly from the hardware. This can assist in determining the hardware configuration of a logical device.

DEVSERV QPAVS

DEVSERV is a z/OS system command used to request basic status information on a device, a group of devices, or storage control units. You can use **DEVSERV** with the **QPAVS** parameter to:

- Describe how a logical subsystem configuration is defined to z/OS
- Highlight the inconsistencies, if any, between IODF and the LSS configuration
- Display unbound alias device types with the UCB parameter and, if necessary, unbox a boxed alias device with the **UNBOX** parameter
- Show information on both a PAV-base address and its PAV-aliases by specifying the **VOLUME** parameter
- Display information on devices

Based on your service needs, you can issue the **DEVSERV QPAVS** command in the following forms:

- **DEVSERV QPAVS,*dddd,nn***
- **DEVSERV QPAVS,*dddd,tttt***
- **DEVSERV QPAVS,SSID=*ssss***

The parameters in these commands have the following values:

- **QPAVS** is a required positional parameter
- ***dddd*** is a 3 or 4-hex-digit device number
- ***nn*** is a decimal number from 1-256 with 1 as the default
- ***tttt*** can be **UCB**, **VOLUME**, or **UNBOX**
- ***ssss*** is the specified SSID value

You can refer to the *z/OS MVS System Commands* for additional information on the **DEVSERV** command. Also the *IBM TotalStorage DFSMS Software Support Reference*, SC26-7440 manual can be referenced for additional information on the **DEVSREV** command.

Figure 11-7 and Figure 11-8 show the command output for two different **DEVSERV QPAV** commands. The first displays the requested device and all its base and alias affiliations. The second displays the requested devices in order.

The form of the command, **DS QP,SSID=ssid** displays the information for all devices in the requested LCU. The **SSID=** parameter can be a long running command as it issues I/O to every DASD drive in the host's hardware configuration and then only prints the drives with the SSID. This allows for checking of duplications of the SSID which must be unique to each LCU.

```
DS QP,D222,VOLUME

IEE459I 08.20.32 DEVSERV QPAVS 591
    HOST                             SUBSYSTEM
 CONFIGURATION                     CONFIGURATION
 -------------                  --------------------
 UNIT                                 UNIT    UA
 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 11-7 DEVSERV QPAVS VOLUME output

Explanation of the output fields:

UNIT NUM.	MVS device number (address)
UA	Unit address from host configuration
TYPE	Device type from host configuration: BASE, ALIAS-dddd, or NON-PAV
SSID	SSID where the device belongs to
UNIT ADDR	Unit address (device ID) from ESS subsystem configuration
UA TYPE	Device type from subsystem configuration: BASE, ALIAS-bb, or NC (not a base not an alias)

STATUS Indicates any discrepancies between the host and the ESS subsystem configurations. Possible values are:

- **INV-ALIAS** On the host side, the unit is defined as an ALIAS whose BASE is different from the one on the subsystem side.
- **NOT-BASE** On the host side, the unit is a BASE, while on the subsystem side, it is not.
- **NOT-ALIAS** On the host side, the unit is an ALIAS, while on the subsystem side, it is not.
- **NON-NPAV** On the host side, the unit is not a BASE nor an ALIAS, while on the subsystem side, it is an ALIAS.

```
DS QP,D123,3

IEE459I 08.20.32 DEVSERV QPAVS 591
     HOST                                 SUBSYSTEM
 CONFIGURATION                          CONFIGURATION
 -------------                       --------------------
 UNIT                                      UNIT    UA
 NUM. UA  TYPE        STATUS      SSID  ADDR.   TYPE
 ---- --  ----        ------      ----  ----    ------------
 D123 23  BASE                    0101   23     BASE
 D124 24  BASE                    0101   24     BASE
 D124 25  BASE                    0101   25     BASE
****      3 DEVICE(S) MET THE SELECTION CRITERIA
```

Figure 11-8 DEVSERV QPAVS output

Figure 11-9 shows a **DEVSERV QPAVS** output with devices for which the device information was not listed, and the corresponding reason code.

```
DS QPAVS,C15F,4

IEE459I 09.18.55 DEVSERV QPAVS 844
     HOST                                 SUBSYSTEM
 CONFIGURATION                          CONFIGURATION
 -------------                       --------------------
 UNIT                                      UNIT    UA
 NUM. UA  TYPE        STATUS      SSID  ADDR.   TYPE
 ---- --  ----        ------      ----  ----    ------------
**** UNLISTED DEVICE(S) AND REASON CODES :
  C15F(07) C160(0A) C161(0A) C162(0A)
**** (07) - DEVICE I/O ERROR
**** (0A) - DEVICE IS AN UNBOUND PAV-ALIAS
****      0 DEVICE(S) MET THE SELECTION CRITERIA
```

Figure 11-9 Unlisted devices

The following list shows the reason codes for unlisted devices:

```
(01) - DEVICE NOT CONFIGURED,UCB NOT FOUND
(02) - UCB NOT CONNECTED
(03) - DEVICE UNAVAILABLE,SCP ROUTINE IN CONTROL
(04) - SUBCHANNEL ERROR
(05) - DEVICE BOXED
(06) - UCB NOT A DASD
(07) - DEVICE I/O ERROR
(08) - DEVICE IS NOT A DASD
(09) - DSE-1 CCW BUILD FAILED
(0A) - DEVICE IS AN UNBOUND PAV-ALIAS
```

11.5.2 DISPLAY MATRIX command

An enhancement to the MVS **DISPLAY MATRIX** (**D M**) command causes information on PAV aliases to be displayed for an ESS.

DISPLAY M=CHP

The **DISPLAY M=CHP** command displays the online and offline status of channel paths. Refer to Figure 11-10 for a sample output from **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
AL DEVICE IS AN ALIAS        UL DEVICE IS AN UNBOUND ALIAS
```

Figure 11-10 Display Channel path status

Symbol `AL` indicates an *active* alias device. Symbol `UL` indicates an *unbound* alias device. The alias device is inactive, not bound to any base device. For example, when a base device is varied offline, its alias devices become unbound.

DISPLAY M=DEV

The **DISPLAY M=DEV** command displays the number of online channel paths to devices (see Figure 11-11).

```
DISPLAY M=DEV(C100)

IEE174I 11.30.43 DISPLAY M 704
DEVICE C100   STATUS=ONLINE
CHP                   C0 C1 C2
DEST LINK ADDRESS     0D 0D 0D
DEST LOGICAL ADDRESS  01 01 01
PATH ONLINE           Y  Y  Y
CHP PHYSICALLY ONLINE Y  Y  Y
PATH OPERATIONAL      Y  Y  Y
ND         = 002105.   .IBM.75.000000012345
DEVICE NED =   2105.   .IBM.75.000000012345
PAV BASE AND ALIASES  4
```

Figure 11-11 D M=DEV - Base device

The `PAV BASE AND ALIASES` field shows the number of exposures for an active PAV base device. The number of exposures is obtained by adding the base device plus all aliases that are currently bound to the base device.

When issued for an alias device, the command shows the base device the alias is currently bound to (see Figure 11-12)

```
D M=DEV(D2FF)

IEE174I  21.02.20  DISPLAY M 504
DEVICE  D2FF     STATUS=ALIAS OF BASE D205
```

Figure 11-12 D M=DEV - Alias device

11.6 Messages and codes

This section documents new and changed z/OS messages and codes in support of ESS. This is not an all inclusive list of ESS related messages.

11.6.1 IEA434I

IEA434I DEVICE ONLINE IS NOT ALLOWED, GREATER THAN 10017 CYL

Explanation:

> During **VARY ONLINE** command processing, the device service exit determined that the device is not allowed to become online for the reason specified.

This message is displayed if you try to vary online a device with more than 10017 cylinders on a z/OS system without Large Volume Support (LVS) coexistence support installed.

Also, devices with more than 32760 configured cylinders will get the message:
IEA434I DEVICE ONLINE IS NOT ALLOWED, GREATER THAN 32760 CYL

11.6.2 IEA435I

The existing message IEA435I generated by DFSMS Device Support is updated to identify the following error condition:

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:

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.

System Programmer Response:

Correct the HCD device definitions to match the physical device type attached to the system. Issuing the DSERV QDASD command in this situation will provide information from the hardware which can be cross checked with the HCD definitions.

11.6.3 IEA480E SIM message

The IEA480E Service Information Message (SIM) console message is a summary message prepared and sent to a host asynchronous to any given error event or events. It reflects the result of error event collection and analysis in the storage control subsystem and indicates that some kind of action needs to be taken, that is, a service action point has been reached.

The SIM message contains a description of the problem the subsystem has identified (the impact of failure) and a description of what resources the repair action affects (the impact of repair). For media SIMs, a description of what action should be taken to resolve the problem is given as well.

The SIM message is not new to the ESS. It is used with other products such as the 3990 or the Magstar 3590 tape subsystem, but its role as a service notification facility for the ESS has changed, as documented in the *IBM TotalStorage Enterprise Storage Server Service Guide 2105 Models E10, E20, F10 and F20, and Expansion Enclosure, Volume 1,* SY27-7605-08:

> “SIM generation by the ESS family of products is not intended to be the primary notification for service, as it was for the 3390, 3990, 9340, and 9390 product families. SIM generation for ESS is a complement to the existing problem notification process, and is used to support previous system attachments to S/390 hosts. The strategy for SIM presentation differs from previous products. Instead of directing a SIM to the failing device and system, hardware SIMs will be presented to all S/390 hosts attached to the storage subsystem. Exception Class 0 and Media SIMs will still be off-loaded against the failing device and system.”

IEA480E *yyyy*, **{SCU | CACHE | DASD | MEDIA }, {SERVICE | MODERATE | SERIOUS | ACUTE} ALERT, MT=***machine type/model***, SER=MMPP-SSSSS, REFCODE=** *nnnn-nnnn-nnnn***, VOLSER=***volser***, ID=***id***, CCHH=x***'cccc hhhh'*, **REPEATED [CONTINUATION OF IEA480E SNS=***sense*]

Explanation:

IEA480E is a Service Information Message (SIM) console message. The specified device or storage control has detected an abnormal condition that requires operator or service attention.

The subsystem component affected by the failure is categorized as **SCU** (storage control hardware), **CACHE** (storage control cache or nonvolatile storage), **DASD** (storage device hardware), or **MEDIA** (device data storage media).

The severity level of the SIM event being reported is classified as **SERVICE, MODERATE, SERIOUS,** or **ACUTE**. ACUTE is the most severe and SERVICE is the least severe.

You can select the severity level of SIM messages being reported. The IBM service representative will set the reporting level at ESS installation time based on information you provide on the Communication Resources work sheet (option **Service information messages (SIMs) for S/390**). See 8.3.2, “Configuration worksheets” on page 116.

Figure 11-13 shows a SIM message that was issued on the system console to notify the operator of a DDM failure on a RAID array.

```
*IEA480E A70D,DASD,SERIOUS ALERT,MT=2105,
SER=0175-12345,REFCODE=E100-0191-0909,VOLSER=/UNKN/,ID=09
CONTINUATION OF IEA480E SNS=00101B4000358F09
8F80400191090904E50049C95201E10005100200FE000000
```

Figure 11-13 Service Information Message (SIM)

11.6.4 IEE169I

IEE169I VARY REJECTED, PATH(*dddd,xx*) **OFFLINE DUE TO C.U.I.R.**

Explanation:

A request was made to place online a path to a device through a channel path. However, the path is being kept offline by a system service.

This message is issued if you try to vary online a channel path that has been taken offline by CUIR.

11.6.5 IOS090I

IOS090I {*xx,rrrr.* **| SETIOS}** *dev* **IS AN INVALID DEVICE**

Explanation:

An error occurred when the operator entered the `SET IOS=xx` or `SETIOS MIH,..` command.

This message is issued if the IECIOSxx PARMLIB member sets a MIH value to a PAV alias device, or the operator enters the `SETIOS MIH,...` command to set a MIH value to a PAV alias device.

11.6.6 CUIR messages

Refer to 11.2.2, “Control Unit Initiated Reconfiguration (CUIR)” on page 178 for a summary of CUIR initiated messages.

11.7 IBM StorWatch Enterprise Storage Server Expert

The IBM StorWatch Enterprise Storage Server Expert (ESS Expert) is an optional software product which provides storage resource monitoring and management functions for the ESS and IBM tape libraries. The ESS Expert provides asset management, capacity management, and performance management capabilities.

The performance reporting function for the ESS provides cluster, device adapter, disk groups, and logical volume level statistics. The information includes number of I/O requests, number of bytes transferred, read and write response times, and cache use statistics reported on an hourly basis. The reports can be viewed in tabular and graphical format. With this information, you can make informed decisions about volume placement and capacity planning, as well as isolate I/O performance bottlenecks.

The ESS Expert is described in Chapter 6, “IBM StorWatch Enterprise Storage Server Expert” on page 87. For detailed information on how to use the ESS Expert, you can consult the IBM redbook *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102.

12

ESS Copy Services for S/390

This chapter overviews the ESS Copy Services available for the S/390 environments. The copy services functions include Peer-to-peer Remote Copy, FlashCopy, Extended Remote Copy, and Concurrent Copy.

12.1 S/390 Copy Services: introduction

ESS Copy Services for S/390 is composed of four functions which serve as tools for disaster recovery, data migration and data duplication. These functions are the following:

- Peer-to-Peer Remote Copy (PPRC)
- FlashCopy
- Extended Remote Copy (XRC)
- Concurrent Copy (CC)

Extended Remote Copy and Concurrent Copy rely on a high speed data engine, the System Data Mover (SDM). Peer-to-Peer Remote Copy and FlashCopy are functions of the ESS which do not rely on the SDM for data movement.

An entire IBM Redbook, *Implementing ESS Copy Services for S/390*, SG24-5680, is devoted to topics related to implementing these four functions. In this chapter, we do not attempt to go into that same level of detail regarding ESS Copy Services. Rather, we simply remind you about the various planning considerations involved.

12.1.1 Peer-to-Peer Remote Copy (PPRC)

PPRC is designed to maintain a current copy of your application data at a remote site. PPRC is a hardware-based disaster recover solution. PPRC is a feature of the ESS which maintains copies of the data using a synchronous method.

PPRC is a hardware solution that enables the shadowing of application system data from one site and its associated volumes to a second system at another site. Updates made on the primary volumes are synchronously shadowed to the secondary volumes. The update is only completed when the data being updated is secured on the secondary site.

12.1.2 FlashCopy

FlashCopy is a function of the ESS which provides a point-in-time copy of the data for application usage such as backup and recovery operations. FlashCopy enables you to copy or dump the data while applications are updating the data. Both the source and target volumes must reside on the same LSS.

DFSMSdss automatically invokes FlashCopy when you issue the COPY Full command to a subsystem that supports FlashCopy functions.

12.1.3 Extended Remote Copy (XRC)

XRC is the asynchronous remote copy solution available for the ESS. It is a combination of hardware and software which offers the highest levels of data integrity and data availability in a disaster recover, workload movement and disk migration environment.

The DFSMSdfp component called the System Data Mover (SDM) will copy the writes issued to the primary volumes to the secondary devices. The SDM in conjunction with the XRC support on the ESS are required for an XRC implementation.

12.1.4 Concurrent Copy (CC)

Concurrent Copy (CC) is both an extended function of the ESS and a component of DFSMSdss. CC enables you to copy or dump data while applications are updating the data. Both IMS/ESA and DB2 can use CC for their backups. CC delivers a copy of the data, in a consistent form, as it existed before the updates took place.

Similar to FlashCopy, CC creates a T0 copy of the source. CC will allow you to make copies at the data set level.

12.2 ESS Copy Services: considerations

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

Here are some things to keep in mind as you plan for ESS Copy Services on your ESS:

- 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. For PPRC, this would involve copying from one ESS to another ESS; for FlashCopy, the copies must be within the same LSS of the ESS.
- If you intend to manage PPRC using the ESS Copy Services Web interface, 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 ESS 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.
- There is a maximum of 2000 volumes allowed for a ESS Copy Services server. This number includes all the primary and secondary PPRC volumes plus all the source and target FlashCopy volumes. (Note: for z/OS when using TSO to control PPRC and FlashCopy, then there is no 2000 volumes limitation).
- The source and target volumes must have the same track capacities and the same number of tracks per cylinder. The target volume must have the same volume capacity or larger than the source volume.
- For an XRC implementation the SDM (System Data Mover) placement must be decided with all considerations taken into account. Locating the SDM at the secondary site has good performance implications.

You will also need to check that the CRIT=YES - Light or Heavy options have been set correctly in the ESS VPD by your IBM SSR.

You will need to have the IBM SSR 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 ESS Copy Services primary.

12.3 PPRC

This section describes the configuration rules and invocation processes for PPRC for S/390

12.3.1 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 a zSeries host. It does not need to be dedicated to PPRC (Note: if the direction of the link is swapped, then the port will become unavailable to the host).
- 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 depends on the total write activity to all the primary devices in each ESS.

12.3.2 How to invoke PPRC for S/390

In this section we describe how to invoke PPRC for S/390.

ESS Copy Services Web interface

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 Copy Services Web interface or the OS/390 DEVSERV QDASD command.

To invoke PPRC you need to perform the following tasks from the ESS 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, you could show only:

- S/390 or open systems volumes
- Source or target volumes
- 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 mentioned above can be found in *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7346, the Web Interface Users Guide, SC26-7346, and *Implementing ESS Copy Services on S/390,* SG24-5680.

In z/OS and OS/390 environments, PPRC can also be invoked using TSO commands or the ICKDSF utility. Also, in these environments, the ANTRQST macro supports PPRC requests.

TSO commands

Specific TSO/E commands can be used in the z/OS and OS/390 environments to control PPRC. These commands are extremely powerful, so it is important that they are used correctly and directed to the correct devices. It is recommended that you place the TSO commands in a RACF-protected library to restrict PPRC TSO commands to authorized storage administrators only.

ICKDSF

ICKDSF supports PPRC operations in the z/OS, z/VM, VSE/ESA and stand-alone environments. The PPRC functions are supported through the PPRCOPY command. ICKDSF also provides a subset of the z/OS PPRC commands for stand-alone use when no operating system is available.

12.4 FlashCopy

This section describes the configuration rules and invocation processes for FlashCopy. How to invoke FlashCopy. These are the ways that FlashCopy can be invoked in the z/OS or OS/390 environment:

- DFSMSdss can be used to invoke FlashCopy in batch via the **`ADDRSSU`** program, with the **`Copy Full`** utility function.
- TSO/E commands that are unique to FlashCopy. These commands are: **`FCESTABL`**, **`FCQUERY`**, and **`FCWITHDR`**.
- DFSMSdfp Advanced Services ANTRQST macro which calls the System Data Mover API.
- The ESS Copy Services Web interface presents you the panels needed to invoke FlashCopy.

To invoke FlashCopy with ESS Copy Services panels, 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 Copy Services Web interface, 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 can be terminated (WITHDRAW) using the ESS Copy Services Web interface. If NOCOPY was not specified, the relationship ends automatically once the background source copy has been completed.

If you have the appropriate DFSMSdss PTF, you may also use the FCWITHDRAW for full volume DUMP operations. If the source volume of the full volume dump operation is the source of a FlashCopy relationship, specifying FCWITHDRAW will cause DFSMSdss to withdraw the FlashCopy relationship when the dump has completed successfully.

Part 3

Implementation in the open systems environment

In this part we discuss the processes for implementing the IBM TotalStorage Enterprise Storage Server in open systems environments. We first cover software requirements for the different open platforms. Then we examine in length how to design and implement an ESS logical configuration that meets your needs, with specific implementation guidelines for the different servers. Finally, we discuss monitoring and managing, the migration of data to the ESS, and considerations for implementing copy services.

The open systems environment is very broad and expands across many different types of platforms. The guidelines presented in this part of the book must be complemented with the detailed technical bibliography of the specific platforms you will be implementing.

13

Open systems support

This chapter guides you through the process of verifying that the hardware and software requirements of the host server are adequate to support the attachment of the IBM TotalStorage Enterprise Storage Server (ESS).

The guidelines discussed in this chapter are for configurations where the host system is an open systems server that is to be attached to the ESS using either SCSI connections or Fibre Channel connections.

13.1 General considerations

Figure 13-1 shows the multiplatform attachment characteristics of the IBM TotalStorage Enterprise Storage Server (ESS).

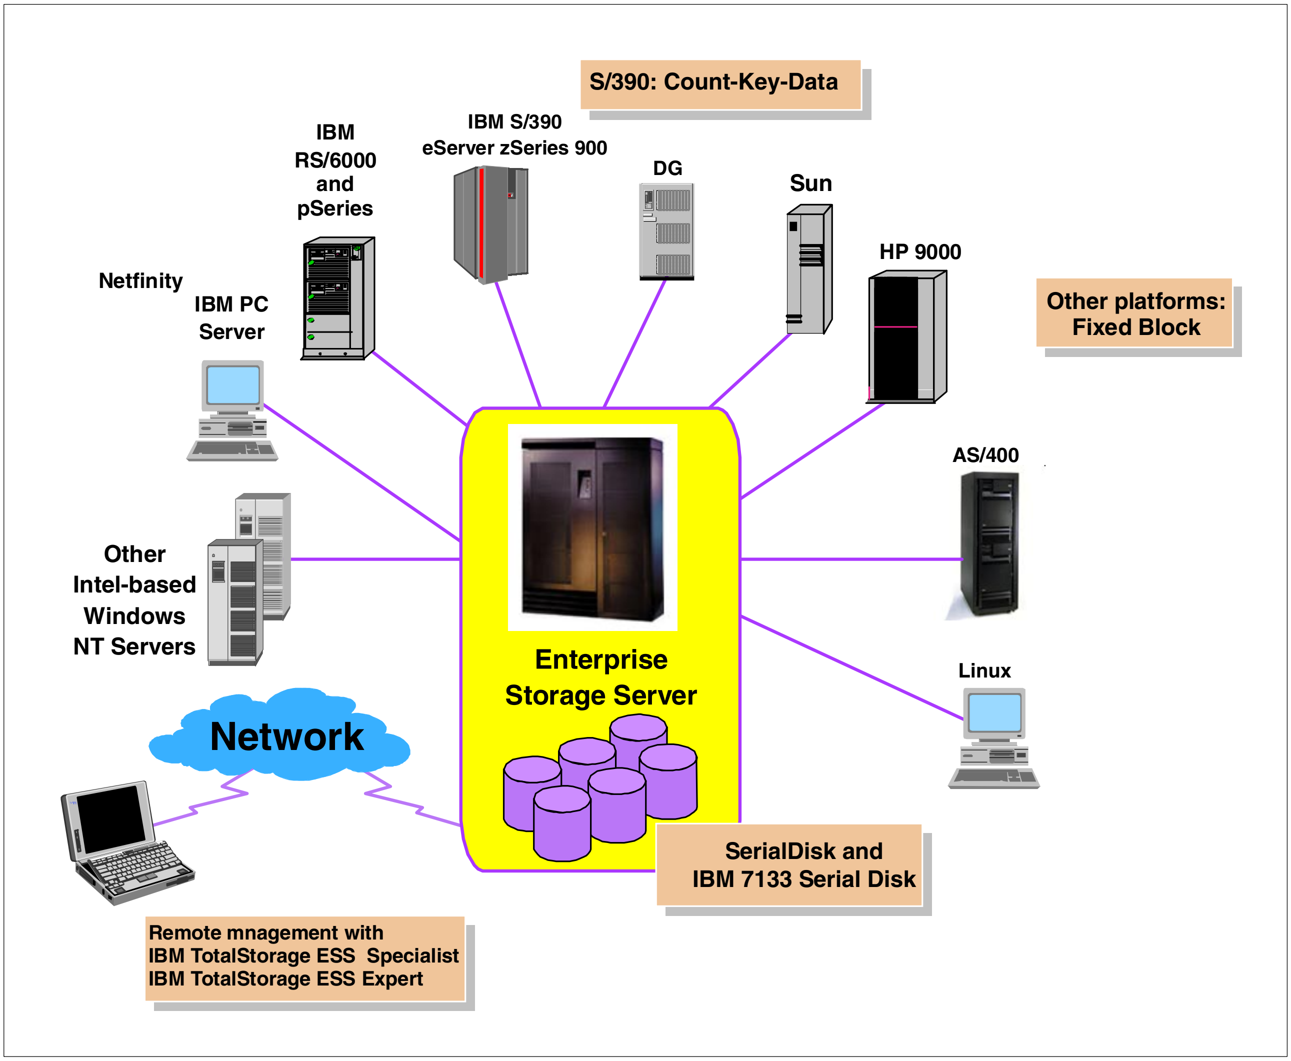

Figure 13-1 ESS - open systems support

When starting to plan for the host systems that are going to be attached to the ESS, you should commence at the following Web site:

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

Here you will find the most current list of open servers that can be attached to the ESS. In Table 13-1 we present the information available at the time of writing our redbook.

Table 13-1 Open systems support summary

Platforms support	O/S Levels support	SDD support	CLI support	Load balancing	SCSI channel support	Fibre channel support
pSeries	AIX 4.2.1 - 5L version 1	Yes	Yes	Yes	Yes	Yes
Compaq	Open VMS	No	No	No	Yes	No
	Tru 64 4.x & 5.x	Yes	Yes	No	Yes	Yes
Hewlett Packard	HP/UX 10.20 - 11.00	Yes	Yes	Yes	Yes	Yes
iSeries	AS/400 V4 release 4.0	No	No	No	Yes	Yes
Intel PC based servers	Windows NT	Yes	Yes	Yes	Yes	Yes
	Windows 2000	Yes	Yes	No	Yes	Yes
Linux	Redhat 7.1	No	Yes	No	Yes	Yes
	SuSE 7.2	No	Yes	No	Yes	Yes
Novell	NetWare 4.2 -5.1	Yes	Yes	No	Yes	Yes
Data General	DG/UX 4.2	No	No	No	Yes	No
Sun	Solaris	Yes	Yes	Yes	Yes	Yes
xSeries/NUMA-Q	DYNIX/ptx	No	Yes	No	No	No
SGI	RPQ	RPQ	RPQ	No	No	RPQ

13.2 Operating system requirements

During the preinstallation planning you should go through the following considerations in respect to the operating system (O/S) running in the servers you plan to attach to the ESS:

- Is the host O/S at the minimum software level required to support the ESS?
- Has all the recommended software maintenance and fixes for the O/S been installed and tested?

The information needed to check the preceding considerations about the O/S requirements, is at the following Web site:

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

Once there, you click, or download, the PDF at **ESS Supported Servers** (see Section 1.5.1, “ESS Web site” on page 9).

This document is updated with the latest support data. You will find in this document, detailed information on the O/S requirements for each open systems platform connectable to the ESS, for both SCSI and Fibre Channel attachment.

13.3 I/O adapters: considerations

During the preinstallation planning you will also go through the following considerations in respect to the input/output (I/O) adapters your servers will use to attach to the ESS:

- What type of connectivity will be used? SCSI and/or Fibre Channel?
- What SCSI or Fibre Channel I/O adapters does your server support, to connect to the ESS?

The information needed to check the preceding considerations about the host system I/O adapter requirements, is at the following Web site:

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

Once there, you click, or download, the PDF at **ESS Supported Servers** (see Section 1.5.1, “ESS Web site” on page 9).

This document is updated with the latest support data. You will find in this document detailed information on the I/O adapters that are supported for each host system server, to attach to the ESS. The information is provided for both SCSI and Fibre Channel attachment.

13.4 Subsystem Device Driver (SDD)

The Subsystem Device Driver (SDD) program, that runs in the host system, provides path failover/failback for increased availability, as well as path load balancing for improved I/O performance.

If you plan to use the SDD program you will have to check how it is supported in your planned configuration. This information is at the URL:

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

Once there, you click, or download, the PDF at **ESS Supported Servers** (see Section 1.5.1, “ESS Web site” on page 9).

This PDF contains a section that documents the SDD support in a matrix format, both for clustered and for non-clustered configurations.

Also, from this same URL, you can click over the entry **SDD (Subsystem Device Driver)**. This will take you to:

`http://ssddom01.storage.ibm.com/techsup/swtechsup.nsf/support/sddupdates`

This last Web page has the latest support information and downloads for SDD installation. At this Web page you can access the *IBM Subsystem Device Driver User’s Guide* as PDF or HTML.

13.5 Command Line Interface (CLI)

The ESS Copy Services Web interface is available for all supported platforms. The ESS Copy Services Web interface function is a Java/CORBA based application that runs on the host system and requires a TCP/IP connection to each ESS under it. ESS Copy Services also includes a Command Line Interface feature (CLI) for communicating with the ESS Copy Services server from the host's command line. The CLI code is supplied with the ESS microcode and is specific to the code level.

If you are planning to use the CLI facility for the ESS Copy Services functions of the ESS, then you should check the O/S environments that do support the CLI interface. This information is at:

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

Once there, you click, or download, the PDF at **ESS Supported Servers** (see Section 1.5.1, "ESS Web site" on page 9).

This PDF contains a section on ESS Copy Services that documents the host O/S requirements for supporting the Command Line Interface.

13.6 Additional considerations

When reviewing the software requirements, you may also have to consider the following:

- Have the Independent Software Vendors (ISV) been contacted regarding their support of the ESS?
- Have any required ISV fixes and prerequisites been installed and tested?
- Are there any considerations for volume manager software (Veritas VxVm, HP SAM, or Sun Solaris DiskSuite)? These are software products that provide disk administration tools to allow you to configure and manage the ESS defined devices, once presented to the host for some open systems platforms.

14

Open systems host setup tasks

The IBM TotalStorage Enterprise Storage Server is designed to handle open systems hosts such as the IBM pSeries and RS/6000 with AIX, IBM iSeries and AS/400 with OS/400, xSeries Servers with Windows NT, Windows 2000, Linux, Novell Netware, NUMA-Q, and many other UNIX based systems. All of these systems use SCSI and/or Fibre Channel generic devices, or specific disk emulations, to address the ESS disk storage.

This chapter discusses the setup tasks you need to perform and the parameters you need to know to successfully connect the ESS to your open systems host, as well as to access the disk devices defined during the logical configuration of the ESS. We discuss how to get these parameters and the commands you will be using for the various open platforms.

14.1 General

The parameters discussed in this chapter are necessary for you to enter in as values, required to complete the ESS open system host definitions and other ESS internal settings. We recommend that before you commence to set up the logical configuration of the ESS, that you gather this information together, so that you have it available to ensure smooth logical configuration of the ESS to your open system platform(s).

You should obtain the following general information for each host platform you will be connecting.

- The number of host SCSI or Fibre Channel I/O adapters for each host.
- The WWPN of each Fibre Channel I/O adapter card used to connect the host to the ESS. This is necessary to complete the ESS logical setup planning.
- How to get the host to recognize the disks, assigned to each SCSI or Fibre Channel I/O adapter.
- How to list the disks on each host and identify the ESS LUN serial number on each host disk.

You will want to know how to get the host to recognize the ESS LUNs, to ensure that your definitions match the ESS internal addressing for the target ID and LUNs. These values are assigned by you using the ESS Specialist, during the logical configuration of the ESS, and can be viewed in the ESS Specialist *Storage Allocation* panel (Tabular view).

The information presented in this chapter can be complemented with the following:

- To work with the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist), refer to the IBM publication *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7346.
- As an additional source of information when implementing FCP on AIX, Sun, HP and Windows NT, refer to the IBM redbook *Implementing Fibre Channel Attachment on the ESS,* SG24-6113
- For the most current and up to date information on the number of supported LUNs per target ID, for the different host servers, please refer to Chapter 1 of the *Host Systems Attachment Guide,* SC26-7296.

See Section 1.5.3, "ESS documentation Web site" on page 10 and Section 1.5.4, "Redbooks" on page 12 for information on how to get ESS related documentation.

As an initial recommendation you should consider not adding more than 120 GB per SCSI I/O adapter, and not more than 300GB per Fibre Channel I/O adapter. This recommendation is because of the SCSI and Fibre Channel attachment characteristics, and not because any limitation on the ESS.

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. AIX 4.3.3 and 5.1 support up to 64 LUNs per target ID. The number of target IDs per adapter is not a consideration when you are attaching the ESS to the host with Fibre Channel adapters.

The AIX operating system contains entries in its Object Distribution Manager (ODM), database to identify the ESS. However, the AIX operating system accesses the ESS through its generic SCSI or Fibre Channel Data Distribution Manager (DDM) and, therefore, sees the ESS LUNs as an 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.

14.2.1 I/O adapters in the host server

To obtain the number the SCSI or Fibre Channel of I/O adapters in your server, enter the following commands:

lsdev -Cc adapter | grep scsi for SCSI

lsdev -Cc adapter | grep fcs for Fibre Channel

The example output for SCSI looks like the following:

```
scsi0   Available 10-60    Wide/Fast-20 SCSI I/O Controller
scsi1   Available 10-68    Wide/Fast-20 SCSI I/O Controller
```

The example output for Fibre Channel looks like the following:

```
fcs0    Available 14-08    FC Adapter
fcs1    Available 2A-08    FC Adapter
```

14.2.2 WWPN of the Fibre Channel I/O adapters

To get the WWPN of the Fibre Channel I/O adapters, enter the following command:

lscfg -vl fcs0 | grep Net

The output looks like this

```
Network Address.............10000000C9280712
```

14.2.3 Host recognition of newly assigned disks

To get the host server to recognize the newly assigned disks run configuration manager or reboot the host:

- To bring in the disks connected to this Fibre Channel type: **cfgmgr -l fcs0**
- To bring in the disks assigned to all adapters, type: **cfgmgr**

14.2.4 ESS LUNs assigned as host hdisks

To list how the ESS LUNs have been assigned as hdisks to the host server, enter the following commands:

- To list all the hdisks, type: **lspv**

  The output looks like this:

  ```
  hdisk8        000b7fad1ffac5ff    none
  ```

- To list all the hdisks assigned to Fibre Channel adapter fcs0, type:

  lsdev -Cc disk | grep 14-08

The output looks like this:

```
hdisk8  Available 14-08-01     IBM FC 2105F20
```

- To show the ESS LUNs serial number corresponding to the hdisk, type:

  lscfg -vl hdisk8 | grep Serial

  The output looks like this:

  ```
  Serial Number...............40D15630
  ```

  Where 40D is the LUN ID and 15630 is the ESS serial number.

If you have implemented and configured IBM Subsystem Device Driver (SDD) then enter the following command to obtain the LUN to disk serial number.

lsvpcfg vpath# where # is the vpath number.

If two adapters are configured on the host then the output looks like this:

```
vpath0 (Avail pv data03MQ1vg) 40D15630 = hdisk8 (Avail ) hdisk33 (Avail )
```

The host sees hdisk8 and hdisk33 as two hdisks, but they are the same LUN connected to two paths. The definition of vpath is used to manage the ESS LUN on the host. This way the host sees the vpath as the physical device and associates the two hdisk definitions to it.

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**.

Type in the following command to delete the device definition:

rmdev -dl hdisk# or **rmdev -dl vpath#** where # is the number of the hdisk or vpath.

Once the hdisks are available, you can assign the hdisks to a volume group and put file systems on the ESS LUNs.

14.2.5 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

For detailed procedures on implementing FlashCopy on an AIX hosts, please refer to the IBM redbook *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757.

14.3 Compaq (DEC)

The Compaq Open VMS and Tru64 4.x supports up to eight LUNs per target ID, 0 to 7, for SCSI adapters, and Compaq Tru64 5.x supports sixteen LUNs per target ID, 0-15, for SCSI adapters. The number of target IDs per adapter is not a consideration when you are attaching the ESS to the host with Fibre Channel.

14.3.1 I/O adapters in the host server

Use the following procedures to install and configure the ESS disk drives on a Compaq Tru64 4.x and 5.x systems:

1. Push the halt button to turn on the Alpha server

   At the **>>>** prompt type in the following command:

   ```
   show device
   ```

   The system responds with a list of controllers and disks that are connected to the host. In the description field on the right of the screen you can check for disk devices that begin with the letters `dk`.

2. Type `boot` to restart the system.
3. Once the system has restarted, open two terminal windows and type in each:

   ```
   uerf -R 300 | more
   ```

   The output should show devices that start with `rz` in each window. For example:

   `rz28, rz29, rzb28,` and `rzb29`

Compare the lists to determine which ESS devices you want to add to the host system.

14.3.2 WWPN of the Fibre Channel I/O adapters

To determine the WWPN of the Fibre Channel I/O adapters in your host server, type in the following command:

```
PO>>>wwidmgr -show
```

Here is an example output with `10000000c922d469` as the WWPN:

```
[0] pga0.0.0.7.1 1000-0000-c922-d469 FABRIC FABRIC
```

Or type in the following command and look for the `wwn` number.

```
view /var/adm/messages
```

14.3.3 Host recognition of newly assigned disks

You must reboot.

If you install the devices after the initial operating system installation, you must make special files that create the character devices needed for the file systems.

For Open VMS

1. Type `cd /dev`
2. Type `./MAKEDEV rzxx`, where `xx` is the number portion of the device name. Repeat this for each new drive that you installed.

Operating system device recognition for Tru64 4.x

Use the following procedures to install and configure the IBM ESS disk drive modules on a Compaq Tru64 UNIX Version 4.0x host system.

Console device check

The following procedures tell you how to perform a console device check.

1. Push the Halt button to turn on the AlphaServer.
2. Type **`show device`** at the **`>>>`** prompt to list the devices available to the AlphaServer.

Configuring devices to mount automatically

To enable an AdvFS file system to start automatically, add an entry to the /etc/fstab file that the mount command will issue during startup. Figure 14-1 on page 217 shows an example of a modified /etc/fstab file.

In Figure 14-1 on page 217, the lines that are shown in bold type are the lines that were entered since the initial operating system installation.

```
# root_domain#root / advfs rw,userquota,groupquota 0 0
/proc /proc procfs rw 0 0
usr_domain#usr /usr advfs rw,userquota,groupquota 0 0
/dev/rz8b swap1 ufs sw 0 2
vol1_dom#vol1  /vol1 advfs rw,userquota,groupquota 0 2
vol2_dom#vol1  /vol2 advfs rw,userquota,groupquota 0 2
vol3_dom#vol1  /vol3 advfs rw,userquota,groupquota 0 2
vol4_dom#vol1  /vol4 advfs rw,userquota,groupquota 0 2
vol5_dom#vol1  /vol5 advfs rw,userquota,groupquota 0 2
vol6_dom#vol1  /vol6 advfs rw,userquota,groupquota 0 2
vol7_dom#vol1  /vol7 advfs rw,userquota,groupquota 0 2
vol8_dom#vol1  /vol8 advfs rw,userquota,groupquota 0 2
vol9_dom#vol1  /vol9 advfs rw,userquota,groupquota 0 2
vol10_dom#vol1 /vol10 advfs rw,userquota,groupquota 0 2
```

Figure 14-1 Example of a modified /etc/fstab file

Installing and configuring the Compaq Tru64 UNIX Version 5.0x

Use the following procedures to install and configure the ESS disk drives on a Compaq Tru64 UNIX Version 5.0x.

Console device check

These are the steps for doing a console device check:

1. Push the halt button to turn on the AlphaServer.

   The system performs self-test diagnostics and responds with the console prompt >>>.

2. Type **show device** at the **>>>** prompt to list the devices available to the AlphaServer.

   The system responds with a list of controllers and disks that are connected to the system. In the description field on the right of the screen, you should see a list of all devices assigned by the ESS. Disk devices begin with the letters **dk**. If you do not see a list of devices, verify the SCSI connections, connectors, and terminators on the bus. If you still do not see a list of devices, check the ESS to ensure that the ESS is operating correctly.

3. When the list of devices is displayed, type **Boot** to restart the system.

Operating system device recognition

To view the disk information for installed devices on the Compaq AlphaServer, type:
#/sbin/hwmgr -view devices

Table 14-1 shows an example of what you see when you type the **#/sbin/hwmgr-view devices** command. The /dev/disk/dsk* device special files are the actual devices connected to the Compaq AlphaServer. The default internal disk, **/dev/disk/dsk0a** is the disk the Compaq AlphaServer uses for starting the operating systems.

Table 14-1 Example of disk information for installed devices for a Compaq AlphaServer

Hardware ID	Device Special File	Manufacturing	Model	Location
3	/dev/kevm	N/A	N/A	N/A
28	/dev/disk/floppy0c	N/A	3.5 inch floppy	fdi0-unit-0
30	/dev/disk/dsk0c	DEC	RZ1DF-CB DEC	bus-0-targ-0-lun-0
31	/dev/disk/cdrom0c	DEC	RRD47 DEC	bus-0-targ-4-lun-0

If you install an SSA device after the initial operating system installation, you must make the device special files that create the character devices needed for filesystems. Perform the following steps:

1. Type: **`#cd /dev`**
2. Type: **`#./MAKEDEV /dev/rdisk/dsk1a`**
3. Type: **`#./MAKEDEV /dev/rdisk/dsk1a`** for each new drive that you installed in the SSA device.

Initializing disk devices

After the list of devices have been determined, you must label the disk volume sizes. Perform the following steps to label the disk volume sizes:

1. Write the new label by typing: **`#disklabel –rw dsk1a ess`**
2. Verify label by typing: **`#disklabel dsk1a`**
3. The **`#disklabel dsk1a`** command shows the new partition layouts on the CompaqTru64 and automatically detects the LUNs that are provided by the SSA device.

Configuring AdvFS

Before you create an AdvFS file system you must design a structure by assigning the file sets.Perform the following steps to create an AdvFS file system with file sets:

1. Type: **`#cd/`**
2. Type: **`#mkfdmn –rw /dev/disk/dsk1a vol1_dom`**
3. Type: **`#mkfset vol1_dom vol1`**
4. Type: **`#mkdir /vol1`**
5. Type: **`#mount vol1_dom#vol1 /vol1`**
6. To display all mounted devices, type: **`df -k`**

Configuring devices to mount automatically

To enable an AdvFS file system to start automatically, add an entry to the /etc/fstab file that the mount command will issue during startup. See Figure 14-2 for an example of a modified /etc/fstab file.

In Figure 14-2, the lines that are shown in bold type are the lines that were entered since the initial operating system installation.

```
# root_domain#root / advfs rw,userquota,groupquota 0 0
/proc /proc procfs rw 0 0
usr_domain#usr /usr advfs rw,userquota,groupquota 0 0
/dev/rz8b swap1 ufs sw 0 2
vol1_dom#vol1 /vol1 advfs rw,userquota,groupquota 0 2
vol2_dom#vol1 /vol2 advfs rw,userquota,groupquota 0 2
vol3_dom#vol1 /vol3 advfs rw,userquota,groupquota 0 2
vol4_dom#vol1 /vol4 advfs rw,userquota,groupquota 0 2
vol5_dom#vol1 /vol5 advfs rw,userquota,groupquota 0 2
vol6_dom#vol1 /vol6 advfs rw,userquota,groupquota 0 2
vol7_dom#vol1 /vol7 advfs rw,userquota,groupquota 0 2
vol8_dom#vol1 /vol8 advfs rw,userquota,groupquota 0 2
vol9_dom#vol1 /vol9 advfs rw,userquota,groupquota 0 2
vol10_dom#vol1 /vol10 advfs rw,userquota,groupquota 0 2
```

Figure 14-2 Example of how to label disks with volume sizes

14.3.4 ESS LUNs assigned as host disks

You can also identify the ESS LUN serial number on Fibre Channel attachments with the following procedures:

1. Type **set mode diag** to put the console into diagnostic mode.
2. Type **wwidmgr -show wwid**

   Figure 14-3 shows an example of the output.

```
[0] UDID: -1 WWID:01000010:6000-1fe1-492-4d20-0000-0000-28b1-5660 (ev:none)

[1] UDID: -1 WWID:01000010:6000-1fe1-492-4d20-0000-0000-2881-5660 (ev:none)

[2] UDID: -1 WWID:01000010:6000-1fe1-492-4d20-0000-0000-2821-5660 (ev:none)
```

Figure 14-3 Example listing of LUN assignment for Compaq servers

In the example in Figure 14-3, *282* is the LUN ID and *15660* is the serial number of the ESS.

1. To verify attachment of the ESS LUNs on the host system, type in the following command:

   hwmgr -view dev -cat disk

2. Here is an example of the output:

   `63: /dev/disk/dsk5c IBM 2105F20 bus-0-targ-253-lun-0`

14.4 Hewlett Packard 9000

HP 9000 systems support up to eight LUNs, 0 to 7 per target ID, 0 to 7, for SCSI adapters. The number of target IDs per adapter is not a consideration when you are attaching the ESS to the host with Fibre Channel I/O adapters.

> **Note:** Ensure that you have at least 1 MB minimum of hard disk space available to install the 2105inst file.

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.

14.4.1 I/O adapters in the host server

Once you have identified the host I/O adapters physically connected to the ESS with the system administrator, you then can type in the following commands for SCSI adapters:

ioscan -kfn | grep SCSI

The output should look something like the example in Figure 14-4.

```
ext_bus   0 0/0/1/0 c720    CLAIMED   INTERFACE   SCSI C896 Fast Wide Single-Ended
ext_bus    1  0/0/1/1 c720     CLAIMED     INTERFACE    SCSI C896 Ultra Wide Single-Ended
```

Figure 14-4 SCSI adapter identification - HP-UX

For Fibre Channel adapters, type in the following command:

ioscan -kfnC fc

Figure 14-5 shows an example output.

```
Class     I  H/W Path  Driver S/W State   H/W Type     Description
=================================================================
fc       0  0/1/0/0   td   CLAIMED     INTERFACE    HP Tachyon TL/TS Fibre Channel Mass Storage Adapter /dev/td0
fc       1  0/1/1/0   td   CLAIMED     INTERFACE    HP Tachyon TL/TS Fibre Channel Mass Storage Adapter /dev/td1
```

Figure 14-5 Fibre Channel adapter identification - HP-UX

14.4.2 WWPN of the Fibre Channel I/O adapters

The tool used to discover the WWPN is:

tdutil /dev/tdx

Here, **x** is the instance number.

For example:

/opt/fcms/bin/tdutil /dev/td0

This should return the WWPN along with all the other information of instance 0.

This is an example of output with **50060b000007a01a** as the WWPN.

```
N_Port Node World Wide Name = 0x50060b000007a01b
N_Port Port World Wide Name = 0x50060b000007a01a
```

14.4.3 Host recognition of the newly assigned disks

To see how the host server is recognizing the newly assigned disks, type in the following command:

ioscan -fnC disk | more

When you are initially bringing new disks/LUNs in there won't be any special device files (ex. /dev/dsk/c4t4d1, etc.) you must then run the command:

insf -e

Figure 14-6 shows a sample output.

```
disk      0  0/0/1/1.2.0 sdisk    CLAIMED     DEVICE       IBM DMVS18D /dev/dsk/c1t2d0   /dev/rdsk/c1t2d0
disk      1  0/0/2/0.2.0 sdisk    CLAIMED     DEVICE       IBM DMVS18D /dev/dsk/c2t2d0   /dev/rdsk/c2t2d0
```

Figure 14-6 Host view of newly assigned disks - HP-UX

14.4.4 ESS LUNs assigned as host disks

You can identify the ESS LUN serial number to disk number with the following command-line interface command:

rslist2105s

Figure 14-7 shows a sample output.

```
disk name       2105 serial number
---------       ------------------
c4t4d1            02117005
c5t1d7            10F17005
c6t1d5            20D17005
c8t4d4            02417005
```

Figure 14-7 LUN serial number to disk number - HP-UX

Note that 021 is the LUN ID and 17005 is the ESS serial number.

14.5 AS/400 and iSeries

The iSeries and AS/400 support only six target IDs and eight LUNs ranging from 0 to 7. The AS/400 target address is always six. Fibre Channel attachment in the iSeries supports 32 target IDs, 1-32.

The information presented on this section can be complemented with the IBM redbook *iSeries in Storage Area Network A Guide to Implementing FC Disk and Tape with iSeries,* SG24-6220.

14.5.1 I/O adapters in the host server

Follow this procedure to recognize the I/O adapters in the host server:

1. Enter the command: **WRKHDWRSC *STG**

   Figure 14-8 shows the sample output, with `2766 Type` resource for the iSeries FC IOA and `6501 Type` resource for SCSI adapters.

```
Work with Storage Resources
                                   System:   RCHAST00

 Type options, press Enter.

   7=Display resource detail   9=Work with resource

Opt  Resource          Type-model     Status          Text

     CMB01             284E-001       Operational     Combined function IOP
       DC01            2778-001       Operational     Storage Controller
     CMB03             284B-001       Operational     Combined function IOP
       DC03            2763-001       Not detected    Storage Controller
       DC07            2782-001       Not detected    Storage Controller
       DC04            2749-001       Operational     Storage Controller
     CMB04             2842-001       Operational     Combined function IOP
       DC05            2768-001       Operational     Storage Controller
       DC06            2766-001       Operational     Storage Controller<<<<<< FC IOA <<<<<<<<
     SI03              6501-001       Operational     Storage Controller
       DC32            9337-5A0       Operational     Disk Storage Controller
```

Figure 14-8 Identifying the I/O adapters - iSeries and AS/400

2. You may need resource detail information for later steps.

   For this you choose option 7 to display resource detail on the 2766 or 6501. Figure 14-10 shows the output information you will get.

```
Display Resource Detail

                                          System:   RCHAST27

Resource name . . . . . . . :     DC06
Text . . . . . . . . . . . . :    Storage Controller
Type-model . . . . . . . . . :    2766-001
Serial number  . . . . . . . :    10-22041
Part number  . . . . . . . . :    0000003N2454
Physical location:
  Frame ID                         2
  Card position                   C08
Logical address:
   PCI bus:
   System bus                     26
   System board                    0
   System card                    16
  Storage:
   I/O adapter                     6
```

Figure 14-9 Resource detailed information for I/O adapters - iSeries and AS/400

3. Take note of the Frame ID and Card position, and System bus, card and adapter.

14.5.2 WWPN of the Fibre Channel I/O adapters

These are the steps you need to follow:

1. On the screen for the AS/400 Main Menu panel, type `strsst`
2. Go to the Sign On panel and enter a valid Service Tools user ID and password.
3. On the *System Service Tools* (SST) panel, type `1` to Start a Service Tool.
4. On the *Start a Service Tool* panel, type `7` to select Hardware Service Manager.
5. On the *Hardware Services* panel, type `1` to select Packaging Hardware Resources.
6. On the *Packaging Hardware Resources* panel, type `9` to select the system expansion unit or System Unit, for the tower that your 2766 is installed in - see Frame ID above.
7. On the *Packaging Hardware Resources* panel, type `8` to select your 2766 Multiple Function IOA.
8. On the Logical Resources Associated with the *Packaging Resource* panel, type `5` to select Multiple Function IOA.
9. On the *Auxiliary Storage Hardware Resource Detail* panel, locate the field name for Port name (see Figure 14-10). The WWID is listed as

```
Port worldwide name . . . . . . . . :   10000000C922D228
```

```
            Auxiliary Storage Hardware Resource Detail


Description . . . . . . . . . . . . . . . . . . . . . . :  Storage IOA
Type-Model . . . . . . . . . . . . . . . . . . . . . . . :  2766-001
Status . . . . . . . . . . . . . . . . . . . . . . . . . :  Operational
Serial number . . . . . . . . . . . . . . . . . . . . . :  10-22041
Part number . . . . . . . . . . . . . . . . . . . . . . :  0000003N2454
Resource name . . . . . . . . . . . . . . . . . . . . . :  DC06
Port worldwide name . . . . . . . . . . . . . . . . . . :  10000000C922D228
PCI bus . . . . . . . . . . . . . . . . . . . . . . . . :
System bus . . . . . . . . . . . . . . . . . . . . . . . :  26
System board . . . . . . . . . . . . . . . . . . . . . . :  0
System card . . . . . . . . . . . . . . . . . . . . . . :  16
Storage . . . . . . . . . . . . . . . . . . . . . . . . :
I/O adapter . . . . . . . . . . . . . . . . . . . . . . :  6
I/O  bus . . . . . . . . . . . . . . . . . . . . . . . . :
Controller . . . . . . . . . . . . . . . . . . . . . . . :
Device . . . . . . . . . . . . . . . . . . . . . . . . . :
```

Figure 14-10 Identifying the WWPN - AS/400 and iSeries

14.5.3 Host recognition of the newly assigned disks

The iSeries host should recognize the newly assigned disks automatically. They will need to be configured into an Auxiliary Storage Pool (ASP) before use.

14.5.4 ESS LUNs assigned as host disks

Type in the following command and follow the instructions:

WRKHDWRSC, option 9 for your 2766 will show the 2105 type resources.

Figure 14-11 shows LUNs that are currently assigned along with LUNs that were assigned, but have been since re-assigned elsewhere, hence are not detected.

```
                    Work with Storage Controller Resources

                                                     System:   RCHAST27

 Type options, press Enter.
 5= Work with configuration desciption        7= Display resource detail

Opt  Resource        Type-model     Status           Text
     DC06            2766-001       Operational      Storage Controller
       DD018         2105-A81       Not detetcted    Disk Unit
       DD034         2105-A05       Operational      Disk Unit
       DD021         2105-A81       Not detetcted    Disk Unit
       DD042         2105-A05       Operational      Disk Unit
       DD020         2105-A81       Not detetcted    Disk Unit
       DD044         2105-A05       Operational      Disk Unit
       DD024         2105-A81       Not detetcted    Disk Unit
       DD045         2105-A05       Operational      Disk Unit
       DD019         2105-A82       Not detetcted    Disk Unit
       DD036         2105-A01       Operational      Disk Unit
       DD025         2105-A82       Not detetcted    Disk Unit
       DD035         2105-A01       Operational      Disk Unit
       DD022         2105-A82       Not detetcted    Disk Unit
       DD041         2105-A01       Operational      Disk Unit
```

Figure 14-11 Listing of the assigned LUNs - iSeries

14.6 Intel based servers

The Intel based servers with Windows NT 4.0 Windows 2000 support up to eight LUNs per target.

Before attaching a host with Windows NT and 2000, 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 *IBM TotalStorage Enterprise Storage Server Host Systems Attachment Guide 2105 Models E10, E20, F10 and F20,* SC26-7296, for more details.

14.6.1 I/O adapters in the host server

See the next step to determine how may I/O adapters are on the system.

14.6.2 WWPN of the Fibre Channel I/O adapter

To find the WWPN for a Windows NT and Windows 2000 host systems follow these procedures depending on the I/O adapter.

Qlogic adapter

Perform the following steps:

1. Restart the server.
2. Press **Alt+Q** to get the FAST!Util menu.

   If you have more than one Fibre Channel adapter installed, all the Fibre Channel adapters are displayed. Scroll down to the adapter you want. Press Enter.
3. From the FAST!Util menu, scroll down and choose **Select Host Adapter**.
4. Scroll up and highlight *Configuration Settings*. Press Enter.
5. From the *Configuration Settings* menu, click **Host Adapter Settings**.
6. Write down the host adapter name, for example: `200000E0800C2D5`.

   Note: The nice feature of QLview is that you do not have to reboot the server. If you don't install QLview, the reboot option is the only way to get the WWPN. The QLview utility can be installed on a workstation as well as a server (it must be installed on the Hosts you intend to query).
7. Click the desktop icon (that is a shortcut to the QLview utility).
8. Update the Machine box - using the Host Name of the Windows/NT or Windows 2000 server.
9. Press Enter or click **Connect**.

   A panel is displayed with diagram of Qlogic adapters installed in the host with WWPN displayed.

Emulex LP8000 adapter

Perform the following steps:

1. Click **Start -> Programs -> Emulex Configuration Tool**
2. From the *Emulex Configuration Tool* window in the *Available Adapters* window, double-click the adapter entry for which you want to display the WWPN information.

14.6.3 Host recognition of newly assigned disks

Perform the following steps:

1. Back up disk configuration on Windows host.
2. Assign LUNs to host using ESS Specialist.
3. Reboot host.
4. Use SDD to verify volume numbers displayed are the same as assigned, and that path is available to all volumes.
5. Use Disk Administrator (NT) or Disk Management (2000) to create logical volumes (LVs).
6. Verify that the *writing signature* prompt matches the new Disk Volume numbers.

14.6.4 ESS LUNs assigned as host disks

You can use the Subsystem Device Driver (SDD) that comes with the ESS, utility to verify LUNs assigned to the host system. You can also use the DiskAdmin command to invoke device discovery.

Perform the following steps:

1. Click **Start** -> **Programs** -> **Subsystem Device Driver** -> **Subsystem Device Driver Management**...
2. At the command prompt, enter:

    `datapath query device`

3. Figure 14-12 shows a sample output:

```
Total Devices : 5

DEV#:   0  DEVICE NAME: Disk2 Part0  TYPE: 2105F20   SERIAL: 01219445

===================================================================

Path#             Adapter/Hard Disk    State     Mode     Select     Errors

    0    Scsi Port5 Bus0/Disk2 Part0     OPEN   NORMAL       1057          0

    1    Scsi Port5 Bus0/Disk7 Part0     OPEN   NORMAL       1137          0

DEV#:   1  DEVICE NAME: Disk3 Part0  TYPE: 2105F20   SERIAL: 01A19445

===================================================================

Path#             Adapter/Hard Disk    State     Mode     Select     Errors

    0    Scsi Port5 Bus0/Disk3 Part0     OPEN   NORMAL         94          0

    1    Scsi Port5 Bus0/Disk8 Part0     OPEN   NORMAL        117          0
```

Figure 14-12 Listing of the assigned LUNs - Intel based servers

In the sample shown in Figure 14-12:

TYPE is the ESS Model, in this case 2105F20

SERIAL is the made up of the Volume Number (vvv) and ESS Serial Number (sssss) in the format vvvsssss

For example DEV#1 shows the Volume 01A and ESS Serial number 19445

Device Name is the Windows/NT disk volume displayed in Disk Administrator or the Windows 2000 disk volume displayed in Disk Management

Path# initially only path 0 will be displayed after reboot when assigning new LUNs

Path 1 will be displayed for the “dual” path (if multiple paths are available)

Path 1 displayed after Disk Administrator or Disk Management is used to define Logical Volumes

> **Note:** The disk number is different on the second path. These volumes will be shown as follows:
>
> **`OFFLINE volumes on Windows/NT`**
>
> They do not display in Disk Management on Windows 2000 (“place holder” is maintained).

14.7 Linux

Intel based servers with Linux support up to 128 LUNs per device. The standard Linux kernel (up to 2.4.x) uses a major and a minor number address mechanism. Each disk device, in fact, each partion on a disk device, is represented by a special device file; for example, **`/dev/sda1`**.

By default there are a maximum of 16 partitions per disk. Both the major and minor numbers are 8 bit. There are 8 major numbers reserved for SCSI devices. Note that Fibre Channel attached devices are handled as SCSI devices as well. The major numbers are 8, 65, 66, 67, 68, 69, 70 and 71 (defined from **`SCSI_DISK0_MAJOR to SCSI_DISK7_MAJOR`** in **`/usr/src/linux/include/linux/major.h`**). According to Linus Torvalds there will be no additional major numbers allocated for disk devices in the future!

Furthermore, there are 256 minor numbers available for each of the 8 major numbers. As there are a maximum of 16 partions per device, the following formula will provide the maximum number of devices under Linux:

```
Number of devices = (number of major#'s) x (number of minor#'s) / (number of partions)
Number of devices = 8 x 256 / 16 = 128
```

You will need to do one of the following tasks to add the Fibre Channel I/O adapters to your Linux host system. You have two options for loading the driver:

- Compile it into the kernel.
- Load it as a module.

14.7.1 Kernel configuration

Use the following procedure to configure the kernel:

1. Edit the top-level *Makefile* and specify the new kernel in **`/usr/src/linux`**.

   Change the "*EXTRAVERSION=*" line according to your new kernel.

   Note: this step is to ensure that you preserve the existing kernel. It might not be necessary when installing a new kernel in parallel.

2. Determine the appropriate processor generation and set the default kernel configuration:

   – **`cd /usr/scr/linux`**
   – **`cp -p configs/kernel-*-i686.config arch/i386/defconfig`**
   – **`make mrproper`**
   – **`make oldconfig`**

   > **Note:** You can either use one of the Red Hat template kernel configuration files and modify it or just start the kernel configuration tool, walking through all parameters. After saving the configuration, you will find a **`.config`** file in **`/usr/src/linux`**. Make a backup copy of the **`.config`** file for future reference (**`make mrproper`** deletes this file).

3. Configure the Linux kernel, change parameters as needed using one of the following tools:

   – **`make config`**
   – **`make menuconfig`**
   – **`make xconfig`**

4. Build the new kernel (it will be a compressed image)

   – **`make dep bzImage`**

5. Build the modules:

   – **`make modules`**

   > **Note:** You can run all the compile steps and install the modules with this command:
   >
   > **`make dep bxImage modules modules_install`**

You can get an updated version of the procedure to configure the Linux kernel, from the Web:

`http://osl.storage.sanjose.ibm.com/cgi-bin/docs.cgi?/linux/Distributions`

14.7.2 Loading the I/O adapter card as a module

Refer to the Qlogic Readme on the Qlogic page for detailed information on loading this card as a module, at the following Web site:

`http://www.qlogic.com/support/product_resources.asp?id=113`

Click **Linux** and take a look at the 2.4 kernel readme.

This is the procedure:

1. Install the new kernel, its **`System.map`** and its modules from within **`/usr/scr/linux`**
   - **`cp -p arch/i386/boot/bzImage /boot/vmlinuz-<version>`**

     where <version> is the full kernel version as specified in this step.
   - **`cp p System.map /boot/System.map-<version>`**
   - **`make modules_install`**
2. Create an initial ramdisk image:
   - **`mkinitrd /boot/initrd-<version>.img <version>`**
3. Edit /**`etc/lilo.conf`**. Append the following lines:
   - **`image=/boot/vmlinuz-<version>`**
   - **`lable=new_kernel`**
   - **`initrd=/boot/initrd-<version>.img`**
   - **`read-only`**
   - **`root=/dev/your-root-partition`** where your-root-partition is whatever you designated it to be.

     Note: In addition you might want to change the global **`default=linux`** line to **`default=new_kernel`**
4. Test the modified */etc/lilo.conf* and make the changes take affect by typing in the commands:
   - **`lilo -t`**
   - **`lilo -v -v`**
5. Reboot the system and verify that the new kernel is running by typing in the following command:
   - **`uname -f`**

14.7.3 I/O adapters in the host server

Type in the following command to list all instances of the installed adapters (this means the number of adapters by number).

```
ls -l /proc/scsi/qla2x00
```

14.7.4 WWPN of the Fibre Channel I/O adapter

On Linux systems the required WWN is displayed in the system log file `/var/log/messages` when loading the driver. Type in the command:

```
/var/log/mesages
```

The output will display the WWN along with other information, look for a line similar to the following:

```
Nov 2 08:55:10 skywalker kernel: scsi-qlal-adapter-port=210000e08b02e534
```

This WWN will be used when configuring the storage using the ESS Specialist *Modify Host Systems* panel (See 15.2.4, “Modify host systems” on page 243) and specifying the appropriate **Linux (x86)** host type.

14.7.5 Host recognition of newly assigned disks

All LUNs are automatically recognized after reboot. If you add LUNs later, the easiest way is to reload the module and scan for attached devices.

To reload the module, type in the following command:

```
rmmod qla2x00
```

To scan for attached devices, type in the following command:

```
insmod qla2x00
```

14.7.6 ESS LUNs assigned as host disks

On Red Hat distribution, there are special device file entries available for the 128 devices. On SuSE distribution, there are only special device files available for the first 16 devices. All other devices have to be created manually using the **mknod** command.

The range of devices goes from `/dev/sda` (LUN 0) to `/dev/sddx` (LUN 127). To list, type in the following command:

```
ls -l /dev/sda
```

14.7.7 Partitioning the disks

Before a file system can be created the disk has to be partitioned using the **fdisk** utility. You have to specify the special device file of the disk you want to partition when executing **fdisk**. For example: **fdisk /dev/sdb**

14.7.8 Creating and using a filesystem

After the disk has been partitioned, as described in the previous section, the next step is to create a file system. The file system can be created using either the **mke2fs** or **mkfs** command. Once the file system is created, it can be mounted and is ready to be used.

14.8 Novell NetWare

In this section we discuss the setup tasks for a Novell NetWare server.

14.8.1 I/O adapters in the host server

At the console prompt use the command:

```
list storage adapters
```

Each adapter appears as **A0**, **A1**, **Ax** where **x** is the adapter number.See Figure 14-13 on page 230.

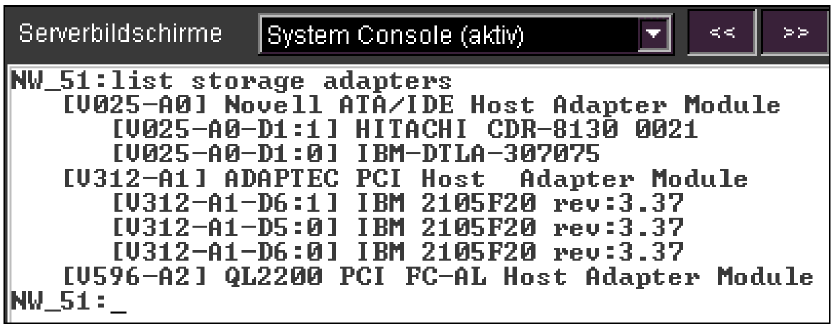

Figure 14-13 Adapter device list

14.8.2 WWPN of the Fibre Channel I/O adapter

Both Qlogic and EMULEX adapters can be used or you can find the WWPN in the BIOS of the FC-Adapter.

Qlogic adapter

For the Qlogic adapter perform the following steps:

1. Restart the server.
2. Press **Alt+Q** to get the FAST!Util menu.

If you have more than one Fibre Channel adapter installed, all the Fibre Channel adapters are displayed. Scroll down to the adapter you want. Press Enter.

1. From the FAST!Util menu, scroll down and choose **Select Host Adapter.**
2. Scroll up and highlight **Configuration Settings**. Press Enter.
3. From the *Configuration Settings* menu, click **Host Adapter Settings**.
4. Write down the host adapter name, for example: `200000E0800C2D5`.

Emulex adapter

For EMULEX adapters, perform the following steps:

1. Boot to MSDOS.
2. Run the Emulex utility **LP6DUTIL**.
3. It will display the WWN, which is the number you are looking for.

14.8.3 Host recognition of newly assigned disks

At the system console, either reboot the host or type the following commands:

`scan all A1` to Scan all LUNs on this SCSI adapter

`scan all luns` to Scan all LUNs on all adapters

A list of all new devices is displayed. You can now partions your new devices and make volume groups.

14.8.4 ESS LUNs assigned as host disks

At the system console, type the following command (also see Figure 14-14.):

`list devices`

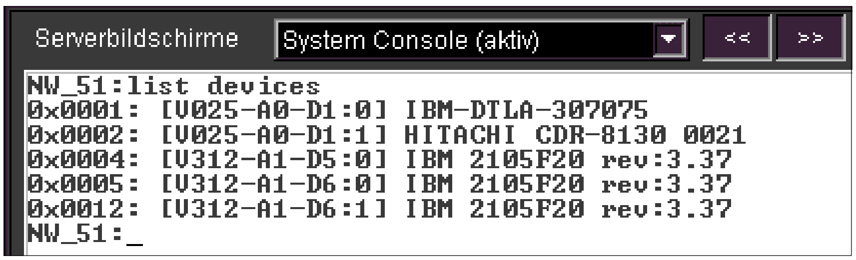

Figure 14-14 list devices output

To get the LUN ID serial number, you can use the command-line interface command:

`rsList2105s`

14.9 Sun

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.

14.9.1 I/O adapters in the host server

Please check with your system administrator and physically identify them. You can also type in the command:

`prtconf`

Here is a sample output:

```
QLGC,isp, instance #4
```

14.9.2 WWPN of the Fibre Channel I/O adapters

Note: If you have multiple I/O adapters installed, you will see more than one WWPN.

After installation of the adapters and restart of the system, view the `/usr/adm/messages` or `/var/adm/messages` log file and search for the line that contains the following:

- For JNI Bus adapters, search for **`fcawx: Fibre Channel WWNN`**, where **`x`** is the adapter number (0, 1,...) You can find the WWNN on the same line following the WWNN.
- For JCI PCI adapters, search for **`fca-pcix: Fibre Channel WWNN`**, where **`x`** is the adapter number (0, 1,...) You can find the WWNN on the same line following the WWNN.
- For the Qlogic QLA2200F adapter, search for **`qla2200-hbax-adapter-port-name`**, where **`x`** is the adapter number (0, 1, ..)
- For the Emulex LP8000 adapter, search for **`lpfcx: Fiber Channel WWNN`**, where **`x`** is the adapter number (0, 1, ..)

14.9.3 Host recognition of newly assigned disks

Boot the host with **`reboot -- -r`**

The host should boot properly and mount all the volumes. VERITAS Volume Manager will identify all of its volumes. If you can not reboot right away you can enter in the series of the next three commands:

```
drvconfig

disks

devlinks
```

14.9.4 ESS LUNs assigned as host disks

Type in the following command:

```
vxdisk list
```

Figure 14-15 shows a sample output:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t2d0s2	sliced	-	-	online
c0t3d0s2	sliced	rootdisk	rootdg	online
c1t1d2s2	sliced	barkydg08	barkydg	online
c1t1d3s2	sliced	barkydg07	barkydg	online
c1t1d4s2	sliced	barkydg06	barkydg	online

Figure 14-15 LUNs listing, sample output -Sun

You can list the ESS LUN serial number to the host with the following command-line interface command:

```
rslist2105s
```

Figure 14-16 show an example output:

```
disk name       2105 serial number
---------       ------------------
c4t4d1            02117005
c5t1d7            10F17005
c6t1d5            20D17005
c8t4d4            02417005
```

Figure 14-16 Output for rslist2105s - LUN serial numbers - Sun

Note that 021 is the LUN ID and 17005 is the ESS serial number.

14.10 NUMA-Q

To complement the information presented in this section, you can refer to the following Web site:

http://techdocs.beaverton.ibm.com

14.10.1 I/O adapters in the host server

Run **devctl -c fabricn**, in where n is the number of the FC fabric connected to the target volume. DYNIX/ptx will then discover the device.

14.10.2 WWPN of Fiber Channel I/O adapters

Type in the following command:

infodev -a

In the following example, the Fibre Channel I/O adapter whose WWPN is sought is `ff0`. The line that begins with `Port WWN` contains the desired value. Perform this step for each Fibre Channel I/O adapter that is to be configured.

/etc/infodev -a ff0

Figure 14-17 shows a sample output:

```
Vendor ID      : Emulex
         Product ID     : Firefly
         Revision Level : FireFly 4
         Device Class   : ctlr
         Node WWN       : 10:00:00:00:c9:20:01:12
         Port WWN       : 10:00:00:00:c9:20:01:12
         Located on     : quad0
```

Figure 14-17 Identifying the WWPN - NUMA-Q servers

14.10.3 Host recognition of newly assigned disks

For each ESS storage device, do the following:

1. In the *NUMA Center Configuration Tool* screen, select **ESS Storage Device 1** (or **2**, **3**, etc.) in the left-hand pane.
2. Update the IP Address 1 and IP Address 2 fields to reflect the subnet and/ or new host ID.
3. Fill in the *Description* box if desired.
4. Save the changes. From the *File* menu, choose **`Save`**.

Configure your system by following the instructions of your host system publication using the **`ConfigApp`** system administration tool. For detailed information you can visit the following site:

http://techdocs.beaverton.ibm.com/docs/ncrnab00/ch_5.htm#SE178140

14.10.4 ESS LUNs assigned as host disks

Use ptx/SVM's **`vxdctl enable`** command to list all ptx/SVM disks known to the operating system's naming database.

The command **`rsList2105s.sh`** is used to list devices recognized by DYNIX/ptx. It is limited by the value of the parameter **`ARRMAX`**, which defines the maximum number of items that an array can contain. In the DYNIX/ptx operating system, that value is 1024.

Thus, one DYNIX/ptx instance can keep track of a maximum of 1024 device files on the system. Since each logical volume within an ESS unit is mapped to a DYNIX/ptx device file, 1024 is the maximum number of logical volumes that can be configured on an ESS unit for access byDYNIX/ptx.

15

ESS configuration for open systems fixed block storage

This chapter describes detailed procedures using the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) to configure fixed block storage in the ESS for use in the open systems environment. At the end of the chapter we present sample scenarios that walk you through the steps in an orderly fashion.

15.1 Introduction

The IBM TotalStorage Enterprise Storage Server Specialist (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, the ESS Specialist might be used to define hosts and add volumes for one of the array configurations that can be ordered for a new ESS.

The ESS Specialist is described in Chapter 5.1, “ESS Specialist” on page 80. To compliment the information presented in the chapter you can refer to the IBM publication *IBM TotalStorage Enterprise Storge Server Web Interface User’s Guide,* SC26-7346 (see Section 1.5.3, “ESS documentation Web site” on page 10).

On the ESS Specialist *Storage Allocation* panel, two buttons appear, see Figure 15-1:

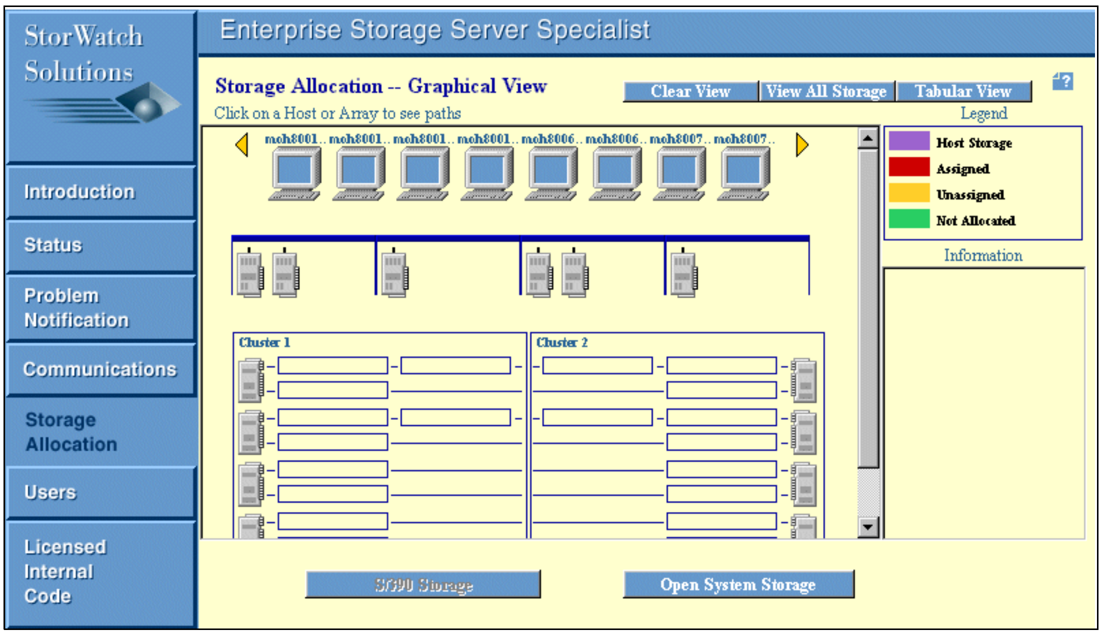

Figure 15-1 ESS Specialist: Storage allocation graphical view

Open System Storage This is used to configure RAID and non-RAID FB storage

S/390Storage This displays configured S/390 storage in tabular form, and leads to further panels that are used to configure CKD storage.

15.1.1 Logical subsystem (LSS)

The ESS uses the concept of the Logical Subsystem (LSS) 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 a SCSI or Fibre Channel host adapter port. 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. Knowing the architecture will help you visualize and plan the configuration. Figure 15-2 shows the relationship between clusters, LSSs, disk groups and LUNs for an ESS with only 8 LSSs configured.

Cluster LSS Disk Group LUN

CLUSTER 1								CLUSTER 2							
LSS 10		LSS 12		LSS 14		LSS 16		LSS 11		LSS 13		LSS 15		LSS 17	
01	09	03	11	05	13	07	15	02	10	04	12	06	14	08	16
16GB 000	16GB 00D	16GB 200	16GB 20D	16GB 400	16GB 40D	16GB 600	16GB 60D	16GB 100	16GB 10D	16GB 300	16GB 30D	16GB 500	16GB 50D	16GB 700	16GB 70D
16GB 001	16GB 00E	16GB 201	16GB 20E	16GB 401	16GB 40E	16GB 601	16GB 60E	16GBb 101	16GB 10E	16GB 301	16GB 30E	16GB 501	16GB 50E	16GB 701	16GB 70E
16GB 002	16GB 00F	16GB 202	16GB 20F	16GB 402	16GB 40F	16GB 602	16GB 60F	16GB 102	16GB 10F	16GB 302	16GB 30F	16GB 502	16GB 50F	16GB 702	16GB 70F
16GB 003	16GB 010	16GB 203	16GB 210	16GB 403	16GB 410	16GB 603	16GB 610	16GB 103	16GB 110	16GB 303	16GB 310	16GB 503	16GB 510	16GB 703	16GB 710
16GB 004	16GB 011	16GB 204	16GB 211	16GB 404	16GB 411	16GB 604	16GB 611	16GB 104	16GB 111	16GB 304	16GB 311	16GB 504	16GB 511	16GB 704	16GB 711
16GB 005	16GB 012	16GB 205	16GB 212	16GB 405	16GB 412	16GB 605	16GB 612	16GB 105	16GB 112	16GB 305	16GB 312	16GB 505	16GB 512	16GB 705	16GB 712
16GB 006	16GB 013	16GB 206	16GB 213	16GB 406	16GB 413	16GB 606	16GB 613	16GB 106	16GB 113	16GB 306	16GB 313	16GB 506	16GB 513	16GB 706	16GB 713
16GB 007	16GB 014	16GB 207	16GB 214	16GB 407	16GB 414	16GB 607	16GB 614	16GB 107	16GB 114	16GB 307	16GB 314	16GB 507	16GB 514	16GB 707	16GB 714
16GB 008	16GB 015	16GB 208	16GB 215	16GB 408	16GB 415	16GB 608	16GB 615	16GB 108	16GB 115	16GB 308	16GB 315	16GB 508	16GB 515	16GB 708	16GB 715
16GB 009	16GB 016	16GB 209	16GB 216	16GB 409	16GB 416	16GB 609	16GB 616	16GB 109	16GB 116	16GB 309	16GB 316	16GB 509	16GB 516	16GB 709	16GB 716
16GB 00A	16GB 017	16GB 20A	16GB 217	16GB 40A	16GB 417	16GB 60A	16GB 617	16GB 10A	16GB 117	16GB 30A	16GB 317	16GB 50A	16GB 517	16GB 70A	16GB 717
16GB 00B	16GB 018	16GB 20B	16GB 218	16GB 40B	16GB 418	16GB 60B	16GB 618	16GB 10B	16GB 118	16GB 30B	16GB 318	16GB 50B	16GB 518	16GB 70B	16GB 718
16GB 00C	16GB 019	16GB 20C	16GB 219	16GB 40C	16GB 419	16GB 60C	16GB 619	16GB 10C	16GB 119	16GB 30C	16GB 319	16GB 50C	16GB 519	16GB 70C	16GB 719

Figure 15-2 Relationship between clusters, LSSs, disk groups and LUNs

In the ESS, logical volumes reside in RAID arrays (disk groups) or non-RAID disks, on SSA loops. Each SSA loop has from two to six disk groups of eight drives (not to be confused with the physical drive enclosures, or 8-packs in which the disks are installed). In Figure 15-2, we show the information for an ESS of only eight LSSs. These eight logical subsystems for FB storage in the ESS are numbered from 10 to 17. You can have up to 16 LSSs for FB storage, the last eight being numbered from 18 to 1F. We are only showing sixteen disk groups, two within each LSS. We hope that this example will help you visualize the architecture of the ESS.

In the next example (Figure 15-3), DA 2 manages up to three 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.

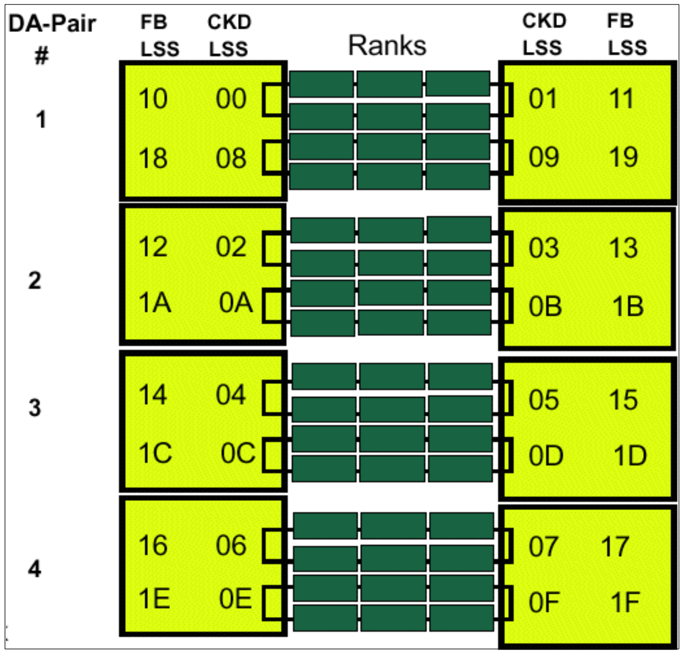

Figure 15-3 Relationship between LSSs DAs and SSA loops

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 0F (see Figure 15-3). 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.

Summary of LSS attributes

An LSS is uniquely associated with one storage cluster, one DA, and half of the 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 disk groups.
- There can be up to 16 FB logical subsystems (LSSs) in an ESS (and from zero to sixteen S/390 LCUs as well).
- An LSS can have up to 256 logical volumes (LUNs) assigned.
- An LSS may have as few as one LUN assigned. That is, a single LUN may contain virtually all the FB capacity in an entire array, and this array may occupy the only disk group in the LSS.

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 ESS Copy Services Web interface, you can specify a group of volumes either by ESS subsystem, or by LSS.
- When 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. Note that Target IDs do not apply to Fibre Channel.

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. (Note that there is an option to make it the same when modifying volume assignments). 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. Up to 14 targets on a given LSS 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).

15.1.3 FB logical volumes and Fibre Channel LUNs

The host sees the FB logical volume as a LUN, because it maps to a single LUN in any Fibre Channel host to which it has been added.

Unlike LUNs assigned to SCSI ports, target IDs, and LUNs per target ID, is not valid with Fibre Channel.

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 15-4. On the ESS Specialist *Storage Allocation* panel there are five buttons in the lower part of the panel. The functions you will use are:

Modify Host Systems Define the open system hosts that will address logical volumes in the ESS. You will define the SCSI or Fibre Channel attached host systems by type and name. You will also select the WWPN and assign Fibre Channel attached host systems to specific ESS fibre ports.

Configure Host Adapter Ports Associate each SCSI or Fibre Channel port with the hosts it attaches. For SCSI, define for each port the target IDs on the SCSI bus. For FCP you will be defining the access mode and also the topology.

Configure disk groups Configure the fixed block RAID (or eventually non-RAID) storage you require.

Add Volumes Create logical volumes (LUNs) and associate them with SCSI ports or Fibre Channel adapters.

Modify Volume Assignments Associate logical volumes with additional SCSI ports or Fibre Channel adapters.

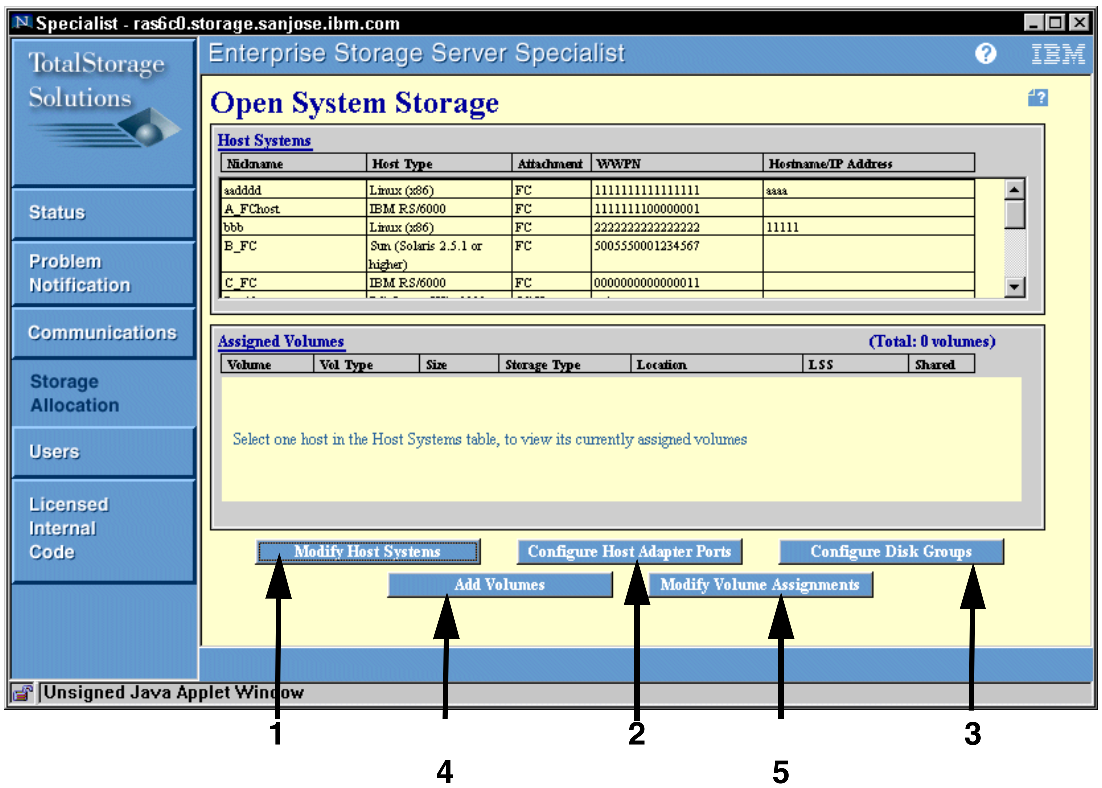

Figure 15-4 Recommended configuration sequence for FB storage

15.2.2 Connecting to ESS Specialist

When you first log in to the ESS, click the *Storage Allocation* panel. Next, click **Open Systems Storage** (see Figure 15-1 on page 236), and then you get to thee screen shown in Figure 15-4. Next, click **Configure Disk Groups**.

Note: When using the ESS Specialist, ensure that the terminal that the SSR uses 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 Configure Disk Groups

The *Fixed Block Storage* panel appears as shown in Figure 15-5. 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 the following:

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 15-5, adapter pair 1, loop A, group 2 has a capacity of 105.24 GB. This is the formatted capacity of a 6+P+S RAID array of 18.2 GB disk drives.

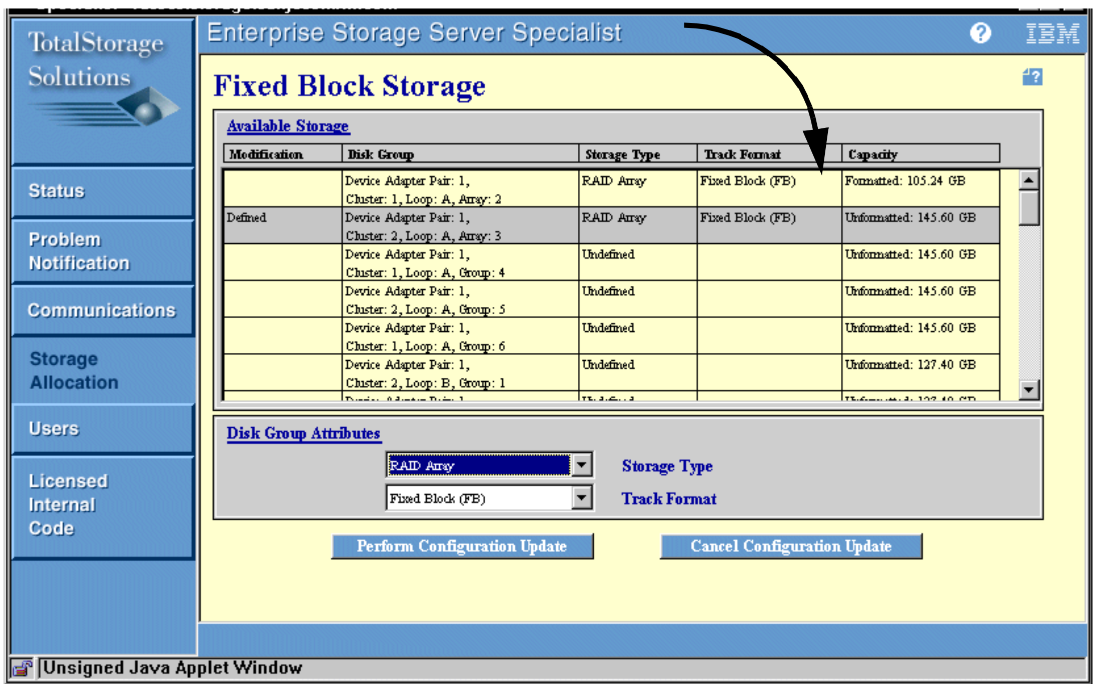

Figure 15-5 Configuring fixed block disk groups

If you select **non-RAID**, the row in the table expands to eight rows, one for each of the disk drives in the group. See Figure 15-6. 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.

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 *Fixed Block Storage* panel (see Figure 15-5) and are available to be defined as FB.

In Figure 15-6, none of the non-RAID disks has been defined as fixed block.

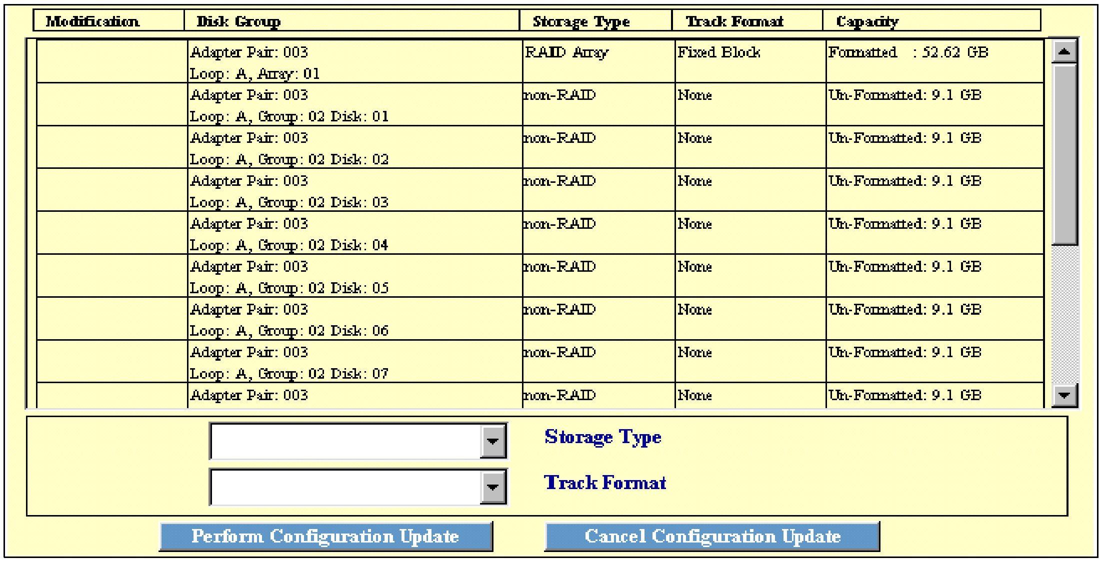

Figure 15-6 Configuring fixed block storage

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. See Figure 15-5 on page 241. Click **Perform Configuration Update**. Message 1402 is presented. Click **Yes** to continue. See Figure 15-7.

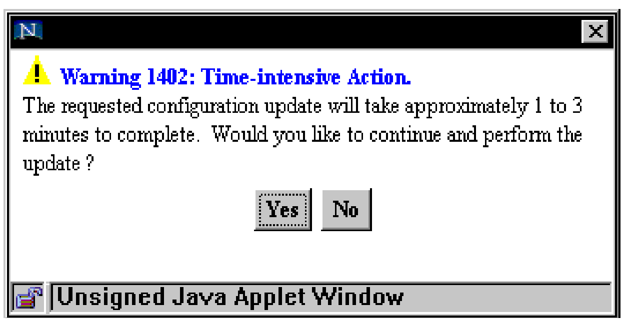

Figure 15-7 ESS Specialist: Message 1402

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 15-8 on page 243.

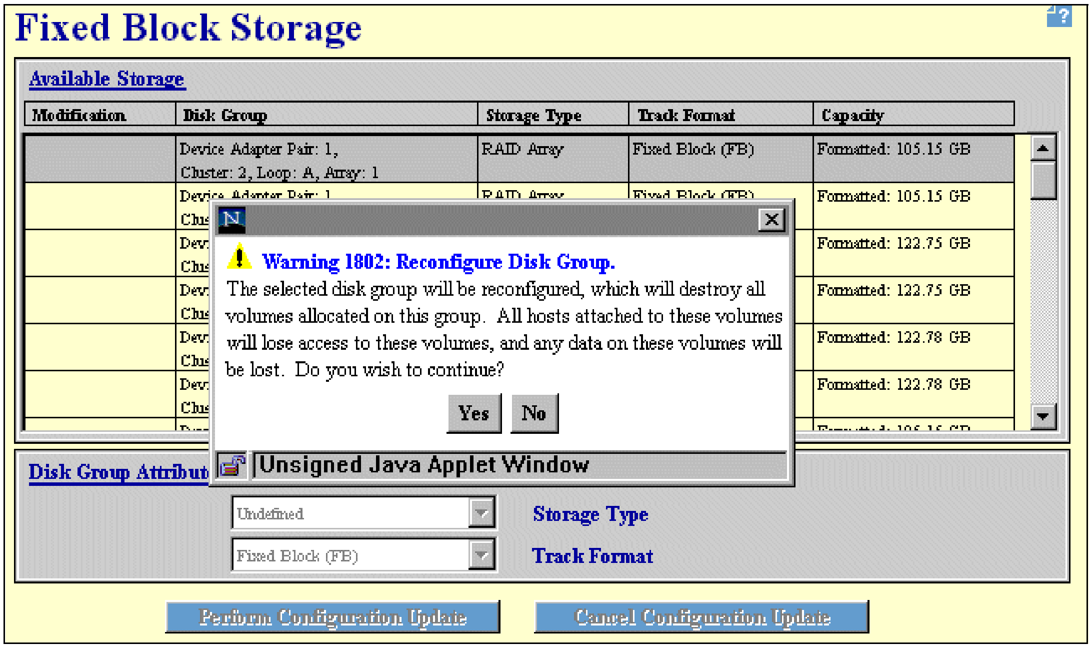

Figure 15-8 ESS Specialist: Warning 1802

Note that all logical volumes and data on the disk group will be discarded if you proceed. To cancel the reconfiguration, click **No** and then click **Cancel Configuration Update** to return to the *Storage Allocation* panel now.

If you want to proceed with the deletion, click **Yes**. The disk group has **Undefined** in the *Modification* column. Click **Perform Configuration Update**. 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 *S/390 Configure 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 or Fibre Channel host to the ESS, on the *Storage Allocation - Graphical View* panel, select **Modify Host Systems**. The panel appears as shown in Figure 15-9. In this example you can see previously created hosts. Enter in the Nickname, Host Type, Host Attachment and Hostname/IP Address fields.

> **Note:** Either the Hostname or the IP Address can be entered. If you know both, you can enter both, but it is not required to enter both.

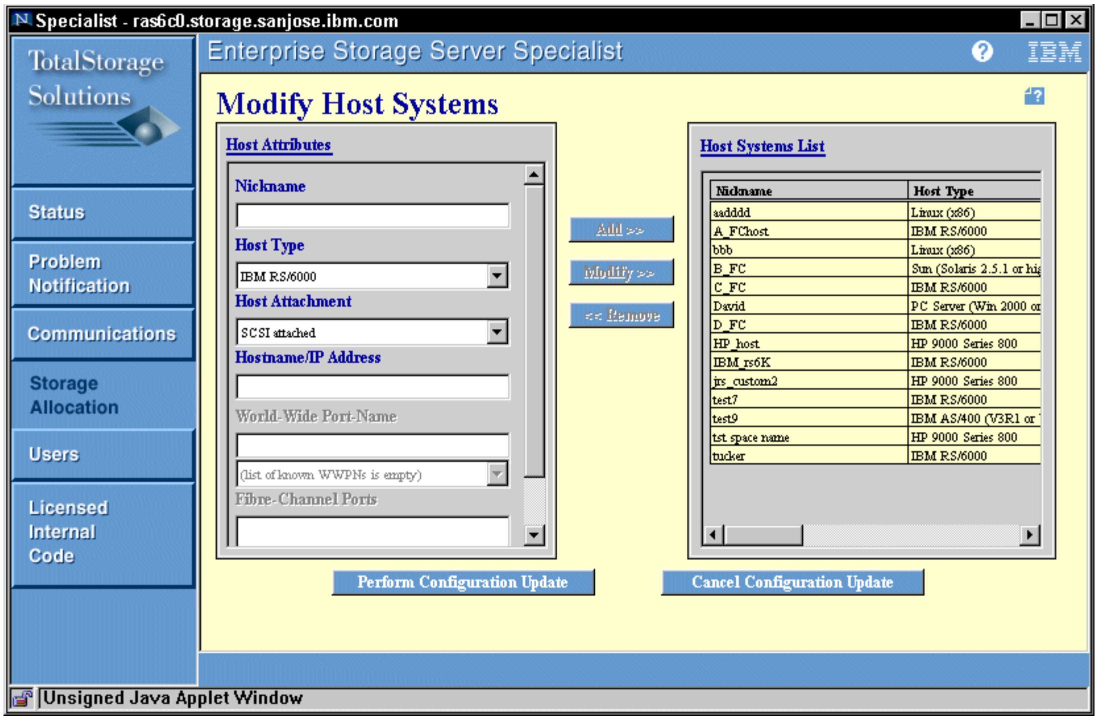

Figure 15-9 ESS Specialist: Modify Host System panel for SCSI

To configure a Fibre Channel host, the Nickname, Host Type, Hostname/IP Address, WWPN and Fibre Channel Ports, are the required fields that you need to define. See Figure 15-10.

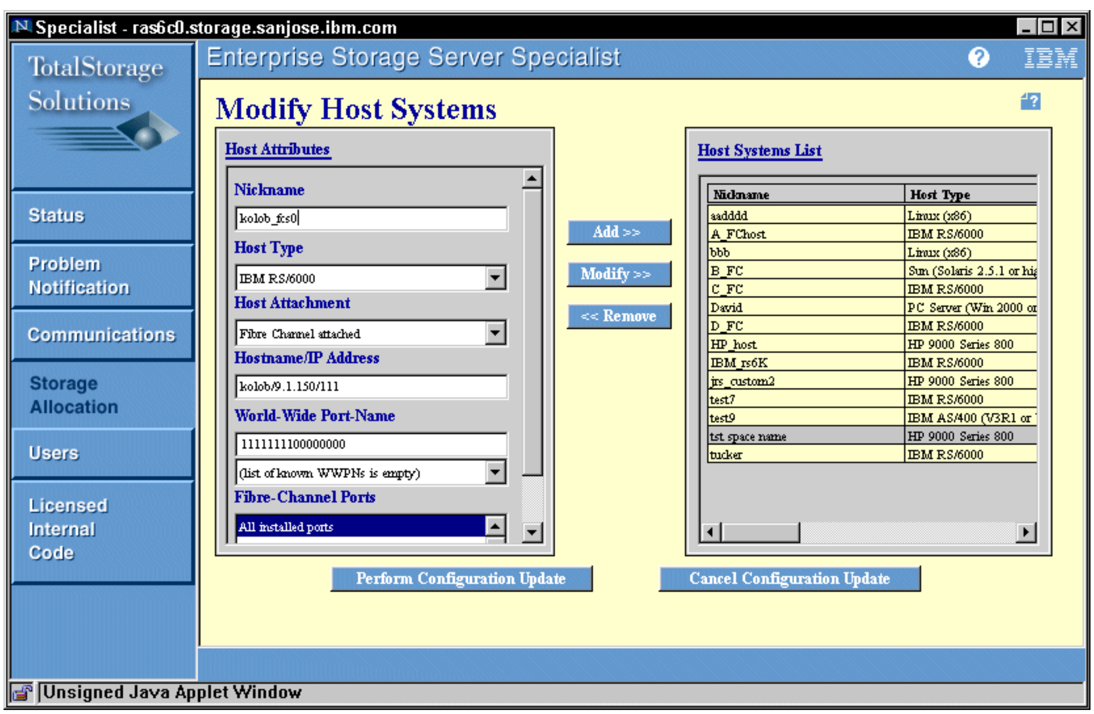

Figure 15-10 ESS Specialist: Modify host systems panel for Fibre Channel

Host Nickname

We suggest that you define a meaningful nickname. For example, if you have two Fibre Channel ports attached to your host (hostname kolob) and the first host port adapter is fcs0, then give the first host port the nickname of kolob_fcs0 and the second port the nickname of kolob_fcs1.

World-Wide Port-Name

Enter in the WWPN for that adapter.To obtain the WWPN refer to the instructions on how to get the WWPN in Chapter 14, “Open systems host setup tasks” on page 211, for your host type. In some cases if the host is physically connected to the ESS then the WWPN will appear in the pull down box.

Fibre-Channel Ports

You will need to determine if you are connecting through a direct cabling scheme or through a switch fabric. If you are connecting through a Switch Fabric, for example a 2109 switch, then we suggest you choose **All installed ports**, and configure the zoning in the switch to make the connection from the ESS to the host. If you are connecting directly without a switch, then you may choose the bay, port and adapter. See the drop down box in Figure 15-11.

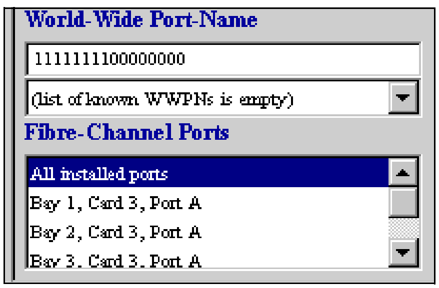

Figure 15-11 ESS Specialist: Fibre-Channel Ports

When you create FB logical volumes and associate them with specific SCSI or Fibre Channel ports, the ESS Specialist uses the host type information to determine how many SCSI and/or LUNs can be created for each target ID in the port. Fibre Channel supports 1 target per host, 16,000 LUNs per target and 128 initiators per FCP. The port uses the defined SCSI or Fibre Channel parameters for the hosts when communicating with a host on the bus.

For SCSI connections the ESS conforms to the SCSI-3 standard, and supports a maximum of 32 SCSI port (two ports per adapter), a maximum of 15 targets per SCSI adapter, a maximum or 64 LUNs per target, depending on the host system type and a maximum or 512 SCSI-FCP host logins or SCSI-3 Initiators per ESS. Hosts using the SCSI-2 standard support 16 target IDs per bus, but only eight LUNs per target ID. Refer to Chapter 13, “Open systems support” on page 205 to determine the maximum number of supported LUNs per target ID for your host type.

You need to plan the storage allocation process carefully to ensure that you can define your required configuration, or modify it as needed.

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).

Host Type

Select the host type from the pull-down list. See Figure 15-12 on page 246.

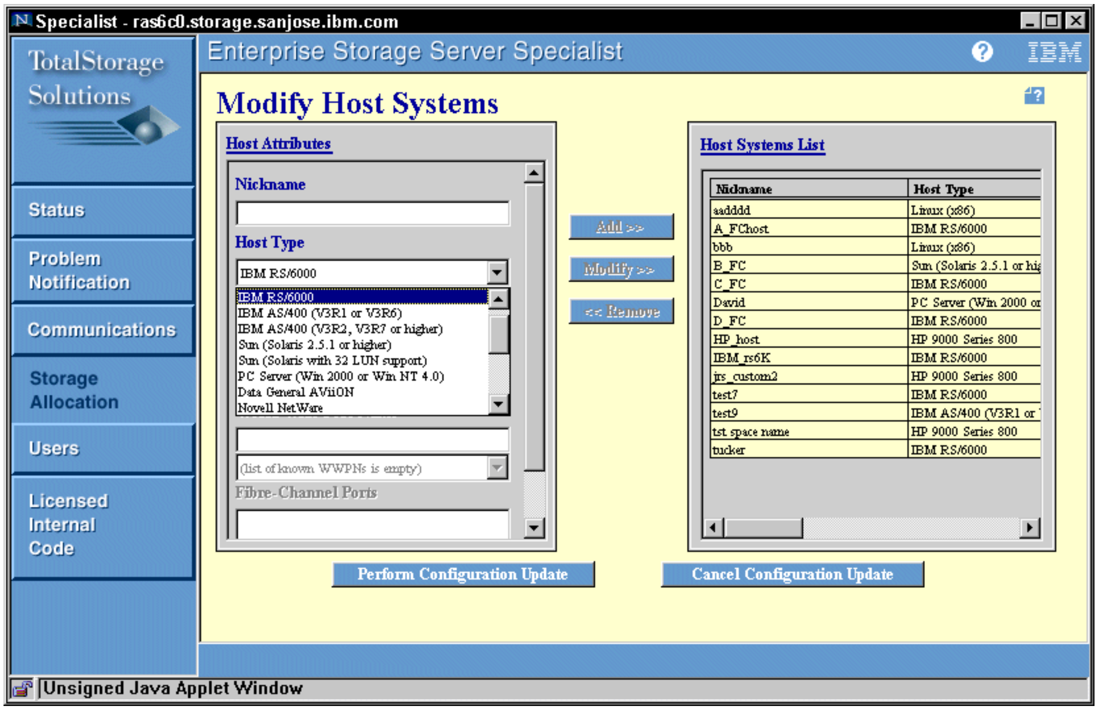

Figure 15-12 ESS Specialist: Host Type drop down box

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/hardsoft/products/ess/supserver.htm`

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 list of hostname appears in the *Hosts Type* pull down box. You can add up to 64 hosts if required.

Removing a host

You can remove any of the previously defined hosts that are no longer required by selecting the *Host Systems List* of the *Modify Host Systems* panel and clicking on the **Remove** button. See Figure 15-12.To activate the changes, click the **Perform Configuration Update** button.

> **Notes:**
>
> - 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.
> - You should remove a host system from the list only after you disconnect it physically from the storage server.

15.2.5 Configure Host Adapter Ports

To define the hosts that will attach to a port, on the *Open System Storage* panel, click **Configure Host Adapter Ports**. See Figure 15-4 on page 240.

SCSI adapter ports

The panel appears as shown in Figure 15-13 for SCSI adapters.In this example, the hostname *tucke*r has been added to the *Map of SCSI IDs on Bus*. Please ignore it for this explanation.

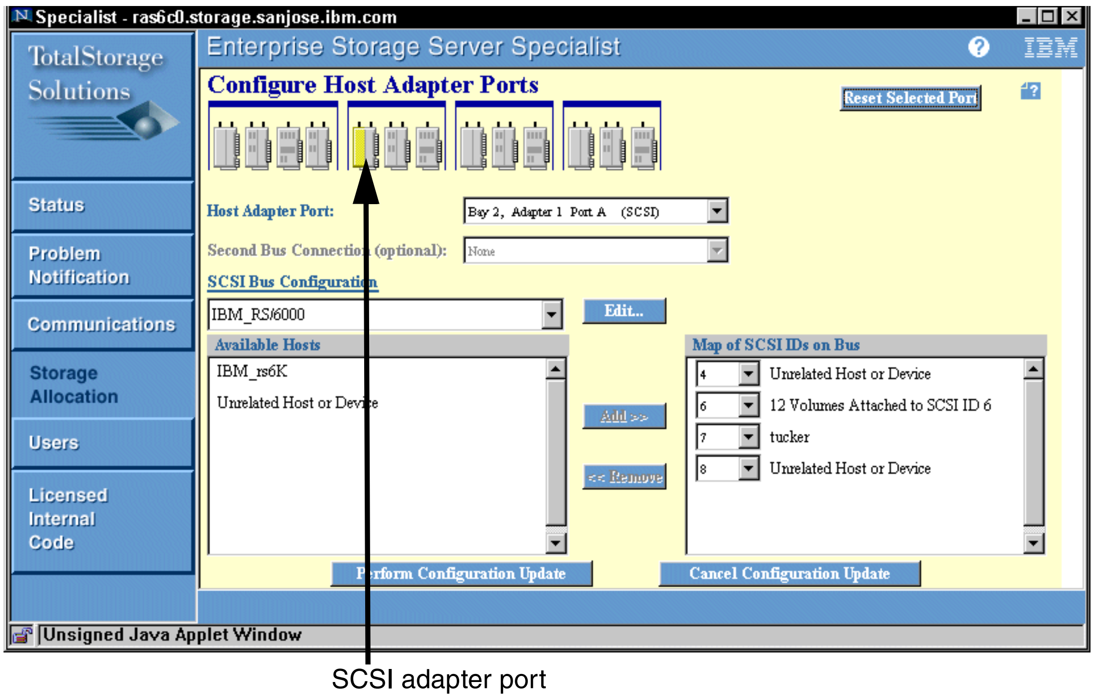

Figure 15-13 ESS Specialist: Configure SCSI host adapter ports

The *Configure Host Adapter Ports* panel identifies the attachment of one or more host systems to an ESS SCSI 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.

Selecting a SCSI port

The first row of the panel (see Figure 15-13) contains icons for the installed host adapters. The ESCON,FICON and Fibre Channel adapter icons have a different appearance than the SCSI adapter icons. In Figure 15-13, the SCSI adapters are adapter that appear to be void of any marking in them. 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 *Host Adapter Port* field.

If a host system is attached to two ports via a common SCSI bus cable, use the *Second Bus Connection (optional)* 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 on BUS* 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.

A list of previously defined hosts (if any) that match the selected bus configuration appears in the *Available Hosts* list.

Editing a SCSI bus configuration

If you want to examine or modify any of the SCSI bus parameters for your host system, click the **Edit** button. The *SCSI Bus Configuration* window is displayed. Refer to Figure 15-14.

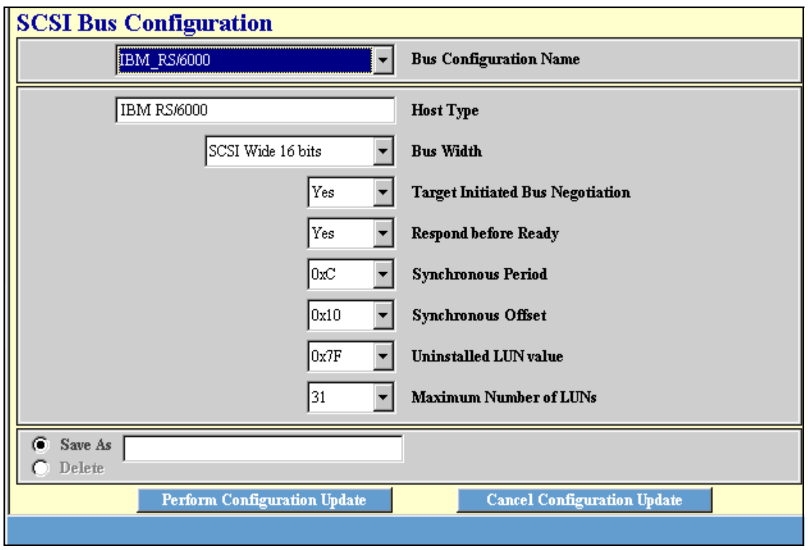

Figure 15-14 ESS Specialist: SCSI bus configuration window

In this example, the IBM_RS/6000 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 *Save As* field. You cannot alter the original bus configuration. Click **Perform Configuration Update** to save the new bus configuration.

Click **OK** to return to the *Configure Host Adapter 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, IBM_RS/6000), should appear under the new configuration, 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.

Adding hosts

After you have selected the correct *SCSI Bus Configuration* as described above, *Available Hosts* lists contains the previously defined hosts that conform to the selected SCSI bus configuration. Another entry, “Unrelated Host or Device” also appears. See Figure 15-13 on page 247. This is normal and should appear here. Select another 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. 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 *"Volumes Attached to SCSI ID"* cannot be changed to a different target ID. See Figure 15-13 on page 247. 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 **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 15-15.

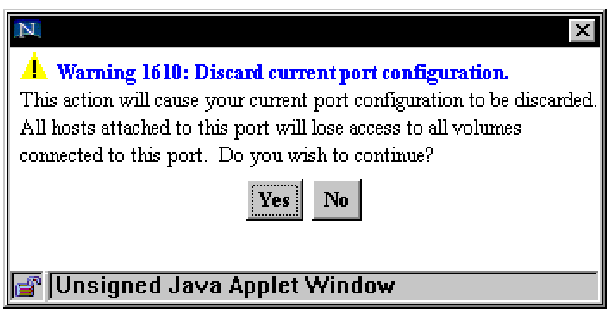

Figure 15-15 ESS Specialist: Warning 1610

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 provides 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 the **Perform Configuration Update** button. The message `Error 1519: Duplicate SCSI ID Assigned` is presented if any ID appears more than once in the list.

Figure 15-16 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 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.

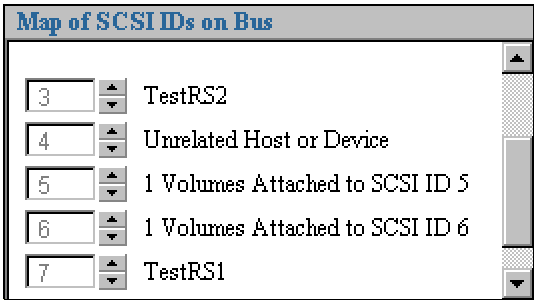

Figure 15-16 ESS Specialist: SCSI ID map

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 that 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 (LSS) 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.See Figure 15-13 on page 247.

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 Host Adapter 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.

Fibre Channel adapter ports

To define the hosts that will attach to a port, on the *Open System View* panel, click **Configure Host Adapter Ports**.

Selecting a Host Adapter port

The first row of the *Configure Host Adapter Ports* panel contains icons for the installed host adapters. The ESCON, FICON and SCSI adapter icons have a different appearance to the Fibre Channel adapter icons. In Figure 15-17, the Fibre adapter is the third one, in bay1.

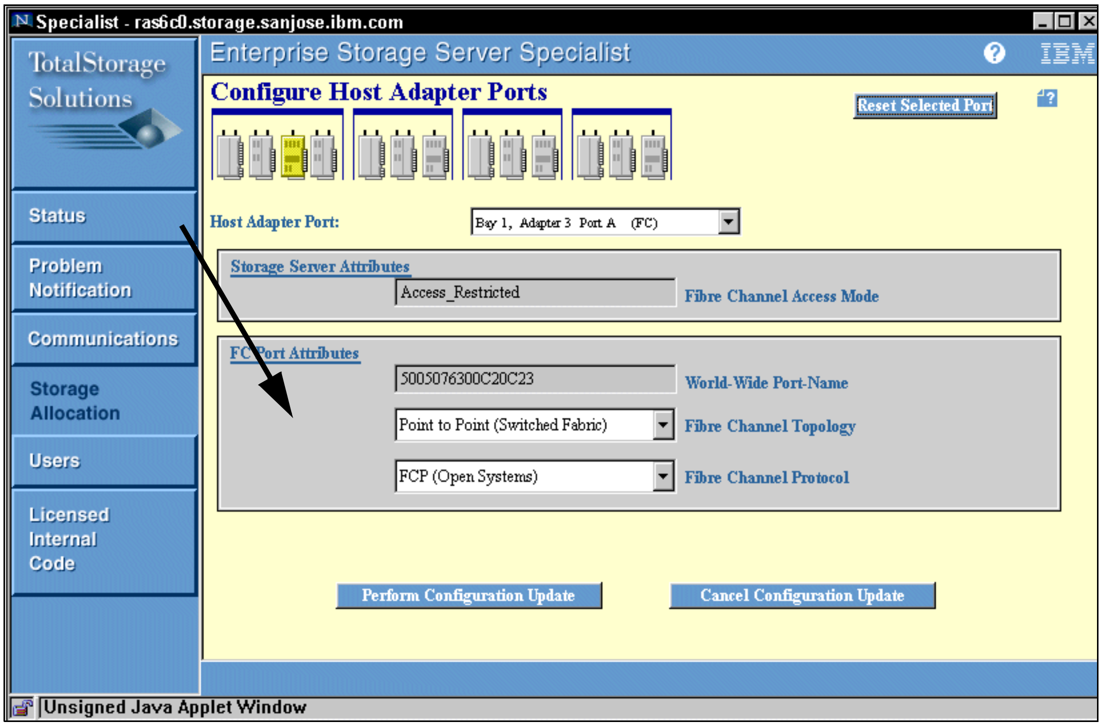

Figure 15-17 ESS Specialist: Configure Fibre port panel

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 *Host Adapter Port*, drop down box.

Figure 15-17 shows the following fields:

- Fibre Channel Access mode. This field is greyed out and not selectable if you chose **Access Restricted** when planning with the configuration worksheets. This function enables a from of zoning and LUN masking by the use of the Nicknames assigned to each host adapter port.
- World Wide Port Name. This field is greyed out and is the WWPN of the ESS I/O host adapter port, not to be confused with the open system host adapters WWPN. You can not change the WWPN of the ESS Host (I/O) Adapter Port.
- Fibre Channel Topology. You may choose **Point to Point (Switched Fabric)** or Arbitrated Loop.
- Fibre Channel Protocol. You can only choose **FCP (Open Systems)**.

15.2.6 Adding fixed block volumes

In order for hosts to access data located in the ESS, the ESS storage must be partitioned into LUNs. Each FB logical volume is seen by the host as a SCSI target/LUN, even though the LUN is physically allocated in SSA disk storage in the ESS. A LUN cannot span ranks, but many LUNs in different ranks can be accessed through a single host port.

Although LUNs 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.

LUNs are selectable from a range of sizes, from 0.1 GB up to 440 GB. For the iSeries and AS/400, only 4.19, 8.59, 17.55, 13.5, 20.56, 35.16, 36 and 70.56 GB are available.

To define LUNs, and associate them with a SCSI or Fibre Channel port, on the *Open System Storage* panel, select **Add Volumes**. The *Add Volumes (1 of 2)* panel is displayed. Refer to Figure 15-18.

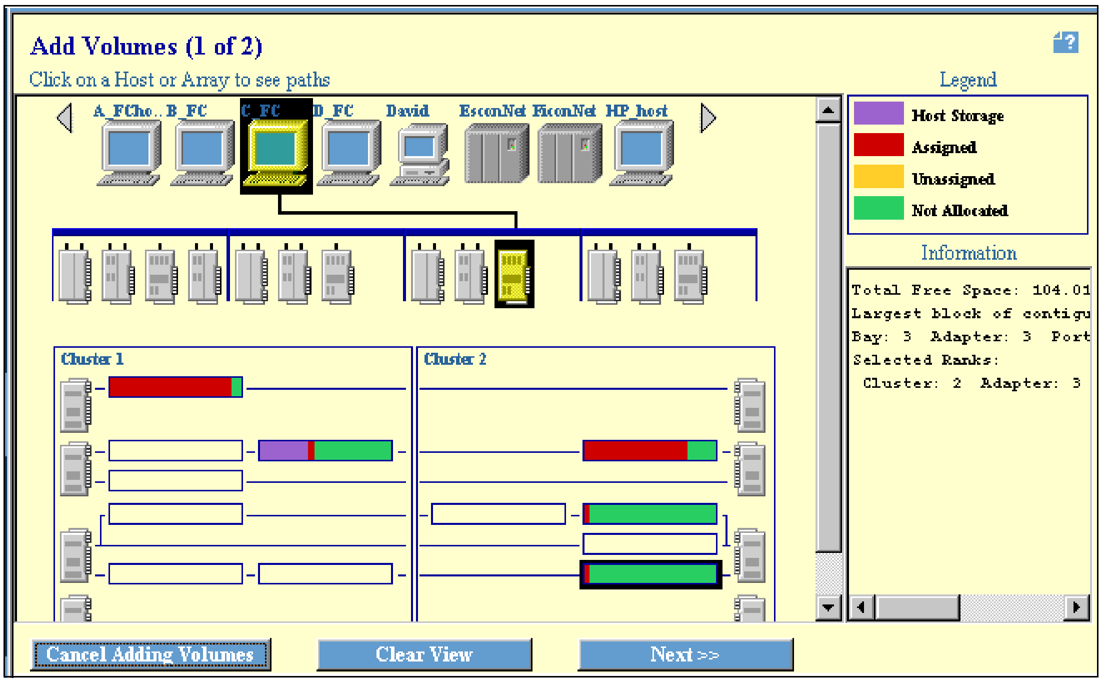

Figure 15-18 ESS Specialist: Add Volumes

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.

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.

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.

If you select a non-RAID group, the panel shown in Figure 15-19 is displayed. (Note that AS/400 does not support non-RAID LUNs.)

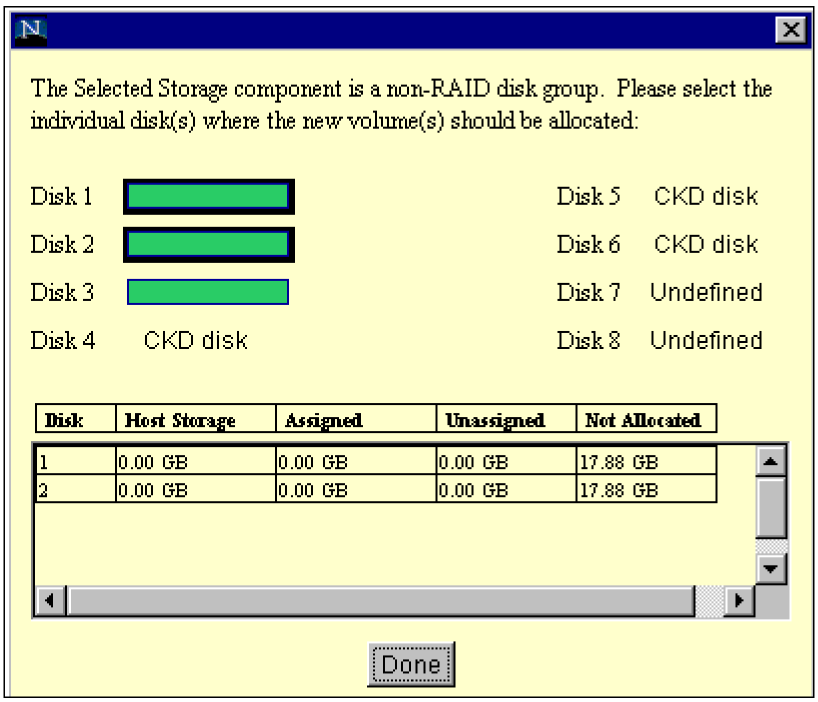

Figure 15-19 ESS Specialist: Non-RAID disk group

The disks colored green are 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 LSSs. Where no expansion frame is installed, there are only two disk groups per loop, and they belong to LSSs in alternate clusters. On the *Storage Allocation -- Graphical View* panel you can click each group and view its status in the Information box. On an SSA loop, even numbered groups belong to one Cluster 1; odd numbered groups belong to Cluster 2.

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 adapter and disk groups is complete, click **Next >>** to proceed to the *Add Volumes (2 of 2)* panel. See Figure 15-20. If you want to cancel instead, then click **Cancel Adding Volumes** to return to the *Open System Storage* panel.

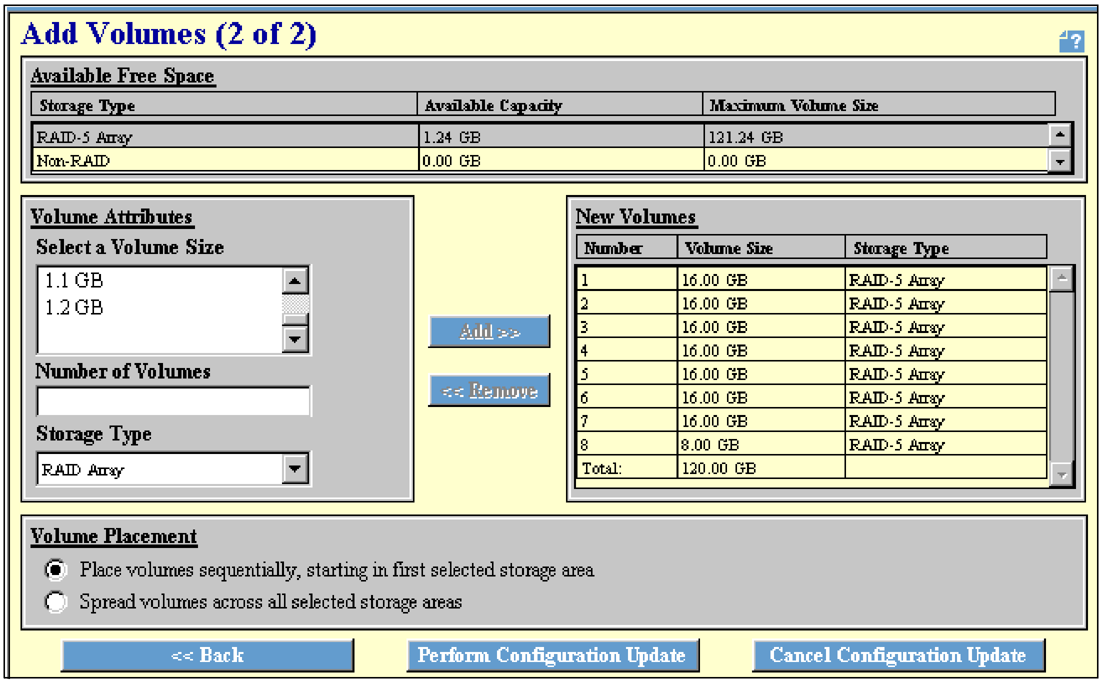

Figure 15-20 ESS Specialist: Add Volumes (2 of 2) panel

This panel is used to allocate volumes in the selected disk groups, and add them to the selected SCSI or Fibre Channel port. We suggest that you add or carve out your LUNs from largest to smallest. This will help you manage other functions like FlashCopy later.

Under *Available Free Space* are three items:

Storage Type RAID 5 or Non-RAID

Available Capacity 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.

Maximum volume size The largest contiguous unallocated space in any FB rank in the selected disk groups.

Add volumes

To add volumes, follow these steps:

- Click the **Storage Type** field at the top of the *Add Volumes (2 of 2)* panel.
- In the **Select a Volume Size** field, use the scroll button to select the required volume capacity you want, in gigabytes (GB).
- In the **Number of Volumes** field, enter the number of volumes of this size you want to create.

- From **Volume Placement**, select to do either of the following:
  - Place volumes sequentially, starting in first selected storage area.
  - Spread Volumes across all selected storage areas.
- Click the **Add >>** button.

Notes:

- 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.
- 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.
- If the selected host is an iSeries, the logical volumes created are iSeries 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 iSeries volumes and open systems volumes are incompatible, you cannot add both types of volumes to a single host port on the *Modify 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 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 Volumes (1 of 2) panel.

When the *New Volumes* to be added list is correct, click **Perform Configuration Update**. See Figure 15-20 on page 254. Clicking on **Cancel Configuration Update** returns you to the Open System Storage panel.

On this panel, when you click 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 **Cancel Adding Volumes** to return to the *Open System Storage* panel.

Note: The *Add 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 *New 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 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 Disk Groups* panel deletes all volumes residing in the group. Any data on the volumes is discarded.

15.2.7 Modify Volume Assignments

The *Modify Volume Assignments* panel is used to associate logical volumes with additional ports, and to remove volumes from ports as needed. The *Modify Volumes Assignments* panel is shown in Figure 15-21.

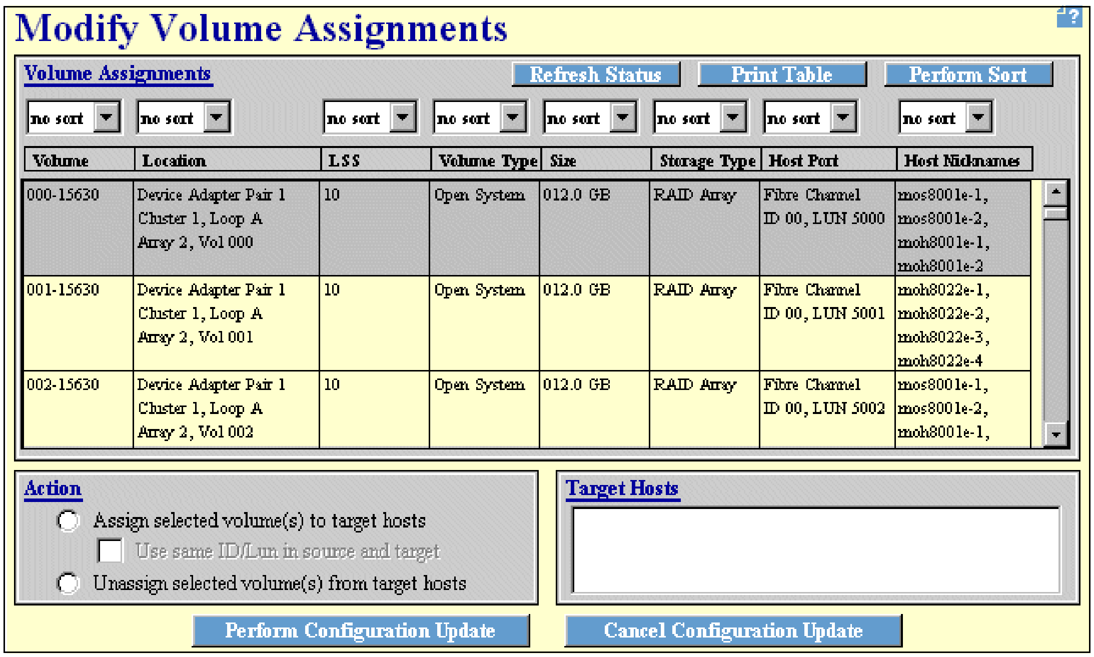

Figure 15-21 ESS Specialist: modify volume assignments panel

The *Modify Volume Assignments* panel is a table of eight 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 Fibre Channel hosts, the table contains an additional Host Nickname entry on the same row, for each extra Fibre Channel host assignment. Refer back to Figure 15-21 to see four entries on the same row. If a volume is assigned to two or more SCSI ports, the table contains an additional row for each extra SCSI port assignment. See Figure 15-22.

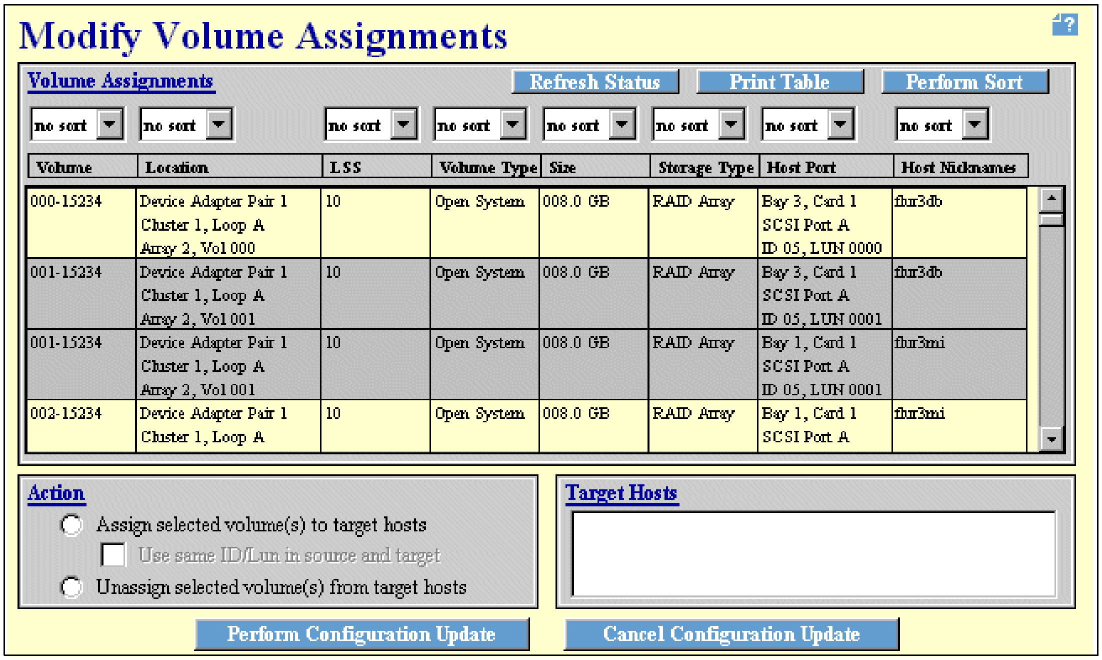

Figure 15-22 ESS Specialist: Dual row SCSI port assignments

Volume table

These are the Volume table columns:

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. For example, if the serial number of the ESS is FCA35, then the very first LUN created in LSS 10 will be 000-FCA35. The first three digits (000) are the LUN hex ID.

Location Displays the location of the logical volume expressed as DA pair/cluster/loop/rank/volume number. Rank is either array or disk group, as applicable.

LSS Displays the Logical Subsystem that the volume is in.

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. For FC only one row is displayed for Fibre Channel.

Hosts Nicknames Displays a list of host system host nicknames that are associated with the volume and port number combination.

Use the scroll bar to move the table past the viewing window as required.

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 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.

Volume selection

Click 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 use Shift-left-click on the row at the other end. To deselect highlighted rows, use Ctrl-left-click on them. To deselect all but one highlighted row, simply click on it.

Select action

When configuring SCSI, you assign volumes to the port. When configuring Fibre Channel, you assign LUNs to a host not a port. Perform one of the following actions for a selected volume(s) by clicking one of the action buttons at the bottom of the panel:

Assign selected volume(s) to target hosts:

Adds selected volumes to the port/host you select from the pull-down in the associated port selector box.

Use same ID/LUN in source and target:

Uses the same LUN ID when selecting the same LUN to assign to an additional port/host.

Unassign selected volume(s) from target hosts:

Removes selected volume(s) from a port/host.

Click the **Perform Configuration Update** button. You are returned to the same panel. To return to the Storage Allocation panel, click **Cancel Configuration Update**, or on **Storage Allocation**.

> **Note:** To add a volume, you must select a port in the selector box. See Figure 15-17 on page 251.

You cannot assign 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.

- Using the *Assign* and *Unassign* functions, you can:
  – Assign a volume to a second or subsequent host port/host. 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/host.
  – Move a volume from one port to another.
  – Leave the volume unassigned by removing its last port/host assignment.
- 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/host.
- A volume, once created, cannot be individually deleted from the disk group. 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 Figure 15-6 on page 242. You can then redefine the group and add new volumes.
- 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.

15.3 Sample procedures

In this section we present sample configuration procedures. We recommend that you do them in the order listed, in order to obtain the best results.

15.3.1 Formatting RAID-5 disk groups

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation**.
3. Click **Open System Storage**.
4. Click **Configure Disk Group**.
5. Select disk group to configure (area will darken).
6. Select *Storage Type*, **RAID 5**.
7. Select *Track Format*, **Fixed Block**.
8. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.2 Creating a SCSI logical host

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation**.
3. Click **Open System Storage**.

4. Click **Modify Host Systems**.
5. Enter Appropriate value into Host Nickname Field.
6. Select **Host Type**, RS6000 is default.
7. Select **Host Attachment Type**, SCSI is default.
8. Enter a valid host name in Host Name/IP Address field.
9. Click **Add**.
10. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.3 Creating a Fibre Channel logical host

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation**.
3. Click **Open System Storage**.
4. Click **Modify Host Systems**.
5. Enter appropriate value into Host Nickname Field.
6. Select **Host Type**, RS6000 is default.
7. Select **Fibre Channel Attached** Type.
8. Enter a Valid host name in Host Name/IP Address field.
9. Enter the **World-Wide Port-Name**.
10. Select **All installed ports**.
11. Click **Add**.
12. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.4 Assigning SCSI Host Adapters to a logical host

Perform these steps:

1. We recommend that you reserve SCSI ID 7 by setting an unknown device to the ID.
2. Make sure all adapters that are assigned to a single host have the same SCSI ID, (preferably SCSI ID 6).
3. Use SCSI ID 6 for the primary HACMP node and SCSI ID 7 for the failover HACMP node.
4. Log into the ESS Specialist application.
5. Click **Storage Allocation** button in the Navigation Frame on the left side of the browser window. This will bring up the graphical storage allocation screen.
6. Click **Open System Storage** button at the bottom of this window. This will bring up the *Open System Storage* screen.
7. Click **Configure Host Adapter Ports** button.
8. Select the adapter (for example: Bay1 Adatper1 PortA), from the graphical bay representation or, from the *Host Adapter Port* selector bar.
9. Select the host type from the *SCSI Bus Configuration* selector bar.
10. Select the host name from the *Available Hosts* selector field.

11. Click **Add**.
12. The Host Name now appears in the Map of SCSI IDs on Bus. Select desired SCSI ID from Selector Bar.
13. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.5 Assigning Switched Fibre Channel host adapters to a logical host

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation** button in the Navigation Frame on the left side of the browser window. This will bring up the graphical storage allocation screen.
3. Click **Open System Storage** button at the bottom of this window. This will bring up the *Open System Storage* screen.
4. Click **Configure Host Adapter Ports** button.
5. Select an adapter (for example: Bay1 Adatper1 PortA), from the graphical bay representation or, from the *Host Adapter Port* selector bar.
6. Verify that *Storage Server Attributes* is **Access_Restricted**. This attribute is previously set by the IBM SSR upon installation in a switch environment.
7. Select from the Fibre Channel Topology selector bar, **Fibre Channel Point to Point**.
8. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.6 Assigning Direct Fibre Channel host adapters to a logical host

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation** button in the Navigation Frame on the left side of the browser window. This will bring up the graphical storage allocation screen.
3. Click **Open System Storage** button at the bottom of this window. This will bring up the Open System Storage screen.
4. Click **Configure Host Adapter Ports** button.
5. Select an adapter (for example: Bay1 Adatper1 PortA), from the graphical bay representation or, from the **Host Adapter Port** selector bar.
6. Verify that Storage Server Attributes is **Access_Restricted**. This attribute is previously set by the IBM SSR upon installation in a Switch environment.
7. Select from the Fibre Channel Topology selector bar, **Fibre Channel Arbitrated Loop**.
8. Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.

15.3.7 Carving LUNs from RAID5 Disk Groups

All carving should take place using a temporary DUMMY Host and DUMMY Host Adapter.

This procedure requires that the DUMMY Host has already been created by you, and that the DUMMY Host Adapter has already been assigned.

Note: If there are no un-used/spare host adapters then you can temporarily use a REAL Host and one of the REAL Host Adapters already assigned to it. We recommend that you select a REAL Host not running HA or SDD.

Procedure

Perform these steps:

1. Log into the ESS Specialist application.
2. Carve disk groups in the order they were installed. Carve the entire disk group before moving to the next.
3. Click **Storage Allocation**.
4. Click **Open System Storage**.
5. Click **Add Volumes**.
6. From the *Add Volumes (1 of 2)* panel. Select the host system to be used as the DUMMY Host.
7. Select the host adapter to be used as the DUMMY Host Adapter.
8. Select the next disk group to be carved.
9. Click **Next**.

Note: Always start with the largest LUN size required for this disk group and work down; for example: 16, 8, 1.

10. Click **Add Volumes (2 of 2)**:
    – First LUN loop pass: Select the largest and primary LUN size for this disk group from the *Volume Attributes* Panel.
    – Second LUN loop pass: Select the next largest LUN size for this disk group from the *Volume Attributes* panel.

Note: Use caution when selecting sizes (observe decimal placement in selections!)

11. Select the LUN count for this LUN Size into the *Number of Volumes* entry field.
12. Click **Place Volumes Sequentially**
13. Click **Add**

Note: Your entries will transfer to the *New Volumes* field, the *Available Space* field will update and, your LUN size selections will change in the *Volume Attributes* field.

14. Repeat the previous steps until the *Available Space* field indicates you have selected all the space available.
15. After you have selected ALL available space in this disk group, then do the following:
    – Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.
    – Repeat all the previous steps until all un-allocated disk groups have been carved.
    – Notify to the IBM SSR to have them verify that ALL disk groups have been successfully carved.

16. Un-assign all new LUNs from the DUMMY host and host adapter.

15.3.8 Assigning LUNs to Fibre Channel host adapters

Perform these steps:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation**.
3. Click **Open Systems Storage**.
4. Click **Modify Volume Assignments**
5. Sort by volume. (This is the default selection.)
6. Select the first Target LUN in *Modify Volume Assignments* panel. The selection will darken.
7. Select the second and subsequent target LUNs in *Modify Volume Assignments* panel while holding down the Shift key; each selection will darken.
8. Click **Assign Selected Volume(s)** to the target hosts after you have selected all LUNs for current operation.
9. Select the first target host/adapter combination in *Target Host* field, the selection will darken.
10. Select the second and subsequent target host/adapter combinations while holding down the Shift key; each selection will darken.
11. After you have selected ALL host/adapter combinations for this operation, then do this:
    - Click **Perform Configuration Update** to commit.

      or
    - Click **Cancel Configuration Update** to cancel.

15.3.9 Assigning LUNs to SCSI Host Adapters

Perform these steps:

1. Select next LUN group. See important notes previously, for criteria of a LUN group. Extract the following information from the LUN from the tabular view of the *Modify Volume Assignments* panel.
   - Target LUN(s) IDs
   - Target Host Name(s)
   - Target Host Adapter(s)
2. Log into the ESS Specialist application
3. Click **Storage Allocation**.
4. Click **Open Systems Storage**.
5. Click **Modify Volume Assignments**
6. Sort by Volume.
7. Click Volume Spinner Bar; Select: **First**.
8. Click **Perform Sort**
9. Select the first target LUN in *Modify Volume Assignments* panel. The selection will darken.
10. Select the second and subsequent target LUNs in *Modify Volume Assignments* panel while holding down the Shift key, each selection will darken.

11. Click **Assign Selected Volume(s)** to the target hosts after you have selected all LUNs for current operation.
12. Select the first target host/adapter combination in *Target Host* panel, the selection will darken.
13. Select the second and subsequent target host/adapter combinations while holding down the Shift key; each selection will darken.
14. After you have selected ALL host/adapter combinations for this operation then:
    - Click **Perform Configuration Update** to commit or click **Cancel Configuration Update** to cancel.
    - For each LUN, verify that the SCSI ID is identical for all of its adapter assignments; see the procedure in the next section.

15.3.10 Verifying a LUN and SCSI IDs

This is the procedure to verify that the LUNs have identical SCSI IDs for each adapter assignment:

1. Log into the ESS Specialist application.
2. Click **Storage Allocation**.
3. Click **Tabular View**.
4. Sort by Volume and Host.
5. Click Volume Spinner Bar, select: **First**.
6. Click Host Spinner Bar, select **Second**.
7. Click **Perform Sort**.
8. Scroll through *Storage Allocation - Tabular View* panel to the desired LUN.
9. You should see a LUN entry for each host adapter, assigned to it.
10. Verify that the volume has the same SCSI ID for each adapter to which it is assigned.

16

Data migration in the open systems environment

This chapter gives you the information necessary to plan the methods and tools your installation will be using when doing the migration of the existing data into the IBM TotalStorage Enterprise Storage Server.

The information presented in this chapter is basically oriented to the open systems users (UNIX; iSeries; and Intel-based servers environments) that must move data from non-ESS storage subsystems into ESS storage-subsystems.

To complement the recommendations and information discussed in this chapter, you may also refer to Chapter 4 (Migrating data to the ESS) of the *IBM TotalStorage Enterprise Storage Server User's Guide 2105 Models E10, E20, F10 and F20,* SC26-7296.

16.1 Migrating in the open system environment

Due to the broad differences in open systems environments, we suggest that you use the general methodology we are presenting here and then determine the specific courses of action based on your host platform.

You can complement the recommendations and information discussed in this chapter referring to chapter 4 (Migrating data to the ESS) of the *IBM TotalStorage Enterprise Storage Server User's Guide 2105 Models E10, E20, F10 and F20,* SC26-7296.

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 spare disks or tape devices for temporary storage
- The format of the data itself

You should select the method that is the best compromise between efficiency and the least impact on the users of the system.

16.1.2 Replacing existing storage

In general 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.

7133 serial disks

The ESS supports existing 7133-020 or 7133-D40 drawers. Whenever these 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.

For configuration guidelines, when attaching your existing 7133 Serial Disk to the IBM TotalStorage Enterprise Storage Server you can refer to the IBM redbook *IBM TotalStorage Enterprise Storage Server,* SG24-5465.

Non-IBM disks

It is possible to free up redundant cabling to external disk storage units and use that port to connect the replacement ESS in parallel. This way you can mirror LUNs or disks on the existing unit to the replacement unit, by means of software techniques. Once the volumes are mirrored, you can then break the mirror from the existing unit, and remove the old unit.

Table 16-1 shows an example of the overall procedure on and AIX platform, migrating from a Symmetrix unit to an ESS unit. You may be able to apply the concept and strategy to other open system platforms.

Table 16-1 Example migration procedure

Process	Commands	Explanation
Phase 1, Identify hdisks, gatekeepers and channels		
Identify the gate keepers on the Symmetrix, look for the disks that are 2880 in size. make note of the rhdisk#	`/usr/lpp/Symmetrix/bin/inq`	where # is the number of the rhdisk
Identify the new adapter connections to obtain the SCSI adapter number	`lsdev -Cc adapter | grep scsi`	If you are using Fibre Channel then replace `scsi` with `fcs`
Identify the hdisks on each of two channels, make note of which hdisk# resides on which adapter#	`lsdev -Cc disk | grep 10-08`	For example: the scsi # previous step is equal to 10-08
Phase 2, Copy lv data from one hdiskpower on the Symmetrix to the hdisk on ESS		
Identify which Volume groups to mirror	`lsvg -o`	Shows all volumes varied on
Check the current information	`lsvg -l vg_name` `lsvg -p vg_name`	This identifies all the logical volumes in the volume group and what disk they are on
Add disk(s) to volume group. Spread the disks evenly across the vgs for best performance	`extendvg vg_name hdisk#`	Where `vg_name` is the volume group name and # is the hdisk number on the ESS unit.
Verify that new hdisk has been added to the volume group	`lsvg -p vg_name`	Where vg_name is the volume group name
Mirror data onto new hdisk	`mklvcopy lv_name 2 hdisk`	Where `lv_name` is the logical volume name and # is the new hdisk number you have added
Sync the volume groups	`syncvg -l lv_name`	Where lv_name is the logical volume name
Verify that logical volume has copied on hdisk# and is still on hdiskpower#	`lspv -l hdisk#` `lspv -l hdiskpower#`	Where `hdisk#` is the new hdisk and `hdiskpower#` is the old disk you are migrating from
Verify that the sync isn't showing `stale, it` should show as `sync'd`	`lsvg -l vg_name`	If the lv still shows stale then you need to resync it before proceeding
Phase 3, remove the disk from the Symmetrix		
Once sync'd, remove the mirrored copy from old hdiskpower#	`rmlvcopy lv_name 1 hdiskpower#`	

Process	Commands	Explanation
Make sure data is off old hdiskpower	`lspv -l hdiskpower#`	
Make sure data is on new hdisk	`lspv -l hdisk#`	
Move the old hdiskpower out of the volume group	`reducevg vg_name hdiskpower`	
Remove the old device definition from the host	`lrmdev -dl hdiskpower#`	Where x is the number of disk you want to delete
Remove the cabling and the old unit from the host.		

16.2 Data migration for iSeries systems

The ESS can attach to the iSeries by either SCSI (V3R1 or later) or Fibre Channel (V5R1 only). The steps to add the external storage to the iSeries are the same regardless of the adapter. With a SCSI adapter, the LUNs will report into the iSeries as device type 9337. With a Fibre Channel adapter the LUNs report in as 2105 device types.

DASD, direct access storage devices (more commonly known as disk drives) are either internal to the iSeries or attached externally. In the iSeries, disk drives are grouped together into auxiliary storage pools (ASPs). The iSeries 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 iSeries.

If you are using the ESS to migrate into it, data that is currently on the internal drives of the iSeries then you must consider that the Load Source Unit (LSU) is a special disk in the iSeries. This is the disk that is used to IPL the system (among other things), it is similar to a boot drive. All other “user data” can be located on external DASD units, but the LSU must be always an internal drive. This is because the system can not be IPL'd from an I/O adapter (IOA) supporting external drives.

If you have an external disk subsystem to the iSeries, that you are replacing with a ESS, you need to migrate the data to the ESS. You can use your existing host utilities for this data migration.To compliment this information, you can refer to chapter 4 of the *Users Guide* found on the Web at:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

You can select from several methods to migrate data to the ESS:

- You can use logical ADD and REMOVE functions.
- You can use save and restore methods with tape devices.
- You can also use these methods if you remove an existing disk subsystem before you install the ESS.

> **Note:** Whether the ESS LUNs report to the iSeries system as Unprotected or Protected (Redundant Array of Independent Disks), any OS/400 reference to starting or stopping device parity on these disk units does not apply. With the appropriate V4R5 PTFs or with V5R1, ESS LUNs created for iSeries as Unprotected can be used to mirror either the LSU or other internal or external unprotected units.

Since the ESS always reports to the iSeries as a Redundant Array of Independent Disks (RAID) (protected 9337 or 2105), any reference to starting or stopping device parity does not apply.

For information on how to use these procedures, see the AS/400 Backup and Recovery Guides for the release of the operating system you have on your host system. Also refer to the IBM redbook *iSeries in Storage Area Network A Guide to Implementing FC Disk and Tape with iSeries,* SG24-6220.

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 devices.

16.3.1 Migration methods

In this section we discuss various migration methods.

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.

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.

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 disks are being removed before you install the ESS, the only way to move the data is with an intermediate 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.

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.

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.

You may not be able to use the volume management methods of migrating data in the following cases:

- 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.

16.3.2 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.

Data migration using AIX LVM mirroring

This is the procedure for using AIX LVM mirroring for data migration:

1. First determine which logical volumes need to be migrated.

   A sample scenario could be one in which 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 two steps describe how to identify those logical volumes:

   a. 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.

   b. 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.

2. Mirror the logical volumes.

   Mirroring does degrade system performance slightly. However, most users 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 user.

   The following steps describe how to mirror logical volumes.

   a. 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.

   b. 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.

3. 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.

   The following process should be used to split each logical volume:

a. 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.

b. 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

   In this command:

   -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

4. 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

   a. 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.

   b. Remount the filesystem by using the **mount** command or using **smit**

5. 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.

6. 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.

   a. Remove each logical volume copy.

   b. 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 and 2000

You can migrate data to an ESS by using standard Windows NT and 2000 commands. No special tools or methods are required for transferring data to ESS. Windows NT and 2000 recognizes the logical volumes in an ESS as normal physical SCSI or FCP disks. For details on migration methods and commands, refer to chapter 4 of the *Users Guide* found on the Web at:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

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 also find many products on the market from different vendors who provide the storage management software for migration and copying. Widely accepted as such a product is IBM TotalStorage Tivoli Storage Manager (TSM).

16.5 Data migration across ESSs

The information in the previous sections is oriented to the data migration scenarios where data from non-ESS storage subsystems is moved (or copied) to ESS storage subsystems.

If your data migration has to do with moving data from one ESS to another ESS you still can use the techniques described so far. But for this ESS to ESS migration scenarios you also have the additional alternative of using the ESS Copy Services functions. More precisely, you can take advantage of the optimized ESTABLISH and RESYNCH operations over PPRC pairs of logical volumes at long distances using the INRANGE Fibre Channel Directors. See “INRANGE” on page 22.

16.6 IBM migration services

This is the easiest way to migrate data, because IBM will assist you throughout the complete migration process. In several countries IBM offers a migration service. Check with your IBM sales representative about migration services for your specific environment and needs.

17

Managing and monitoring the ESS

In this chapter we examine some of the tools that are available for managing and monitoring the ESS in the open systems environment, on the host operating system level. Each operating system has some form of storage management capability. In addition to this, you can use the IBM TotalStorage family of products.

17.1 IBM StorWatch Enterprise Storage Server Expert

The IBM StorWatch Enterprise Storage Server Expert (ESS Expert) is a software tool designed to help storage administrators manage and monitor IBM storage servers. It is an optional product that the storage administrator has for managing and monitoring the performance, the capacity, and the components of the storage environment.

This tool is described in Chapter 6, "IBM StorWatch Enterprise Storage Server Expert" on page 87 of this redbook. To know how to use the ESS Expert you can consult the IBM redbook *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide,* SG24-6102.

17.2 General considerations

When planning for an efficient administration of your storage environment, and thinking in terms of management and monitoring tools and procedures you will need, you should ask yourself the following questions:

- What tools are available to help you in the administration of your system?
- Do you have access to any Web site tools and documentation?
- Do you have access to software technical support?

We suggest that after you read this chapter, you will be able to address these considerations. The following sections in this chapter should help you start to formulate ideas to resolve these proposed questions.

17.3 Managing and monitoring tools

Many open system platforms have software tools designed to help system and storage administrators to manage and monitor disk storage devices. It is beyond the scope of this text to go into detail about the various tools and commands available on each open system platform. We will only examine a few.

17.3.1 AIX

System Management Interface Tool (SMIT)

Typing in the commands **`smit`** on a Graphical User Interface (GUI) or **`smitty`** on a non GUI, present excellent tools you can use to control adapters, disks, and arrays. Once you have launched the file, you can press F1 for help using **`smit`,** see Figure 17-1. For additional on-line help using **`smit`**, refer to chapter 7 of the Redbook *IBM Certification Study Guide pSeries AIX System Support*, SG24-6199.

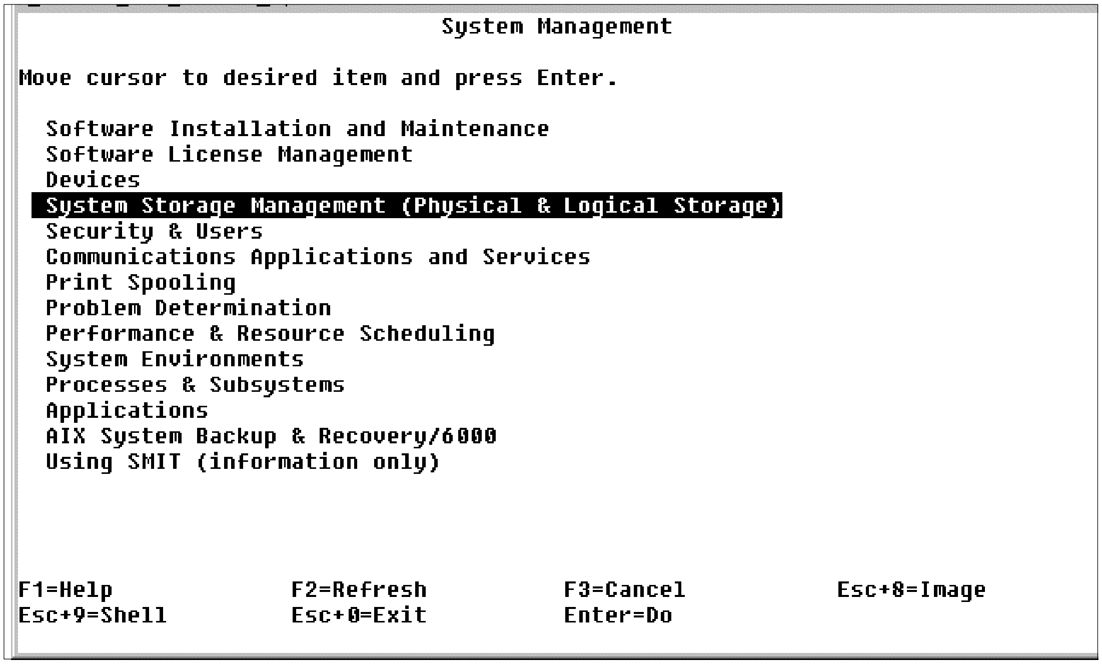

Figure 17-1 AIX smitty panel

Being familiar with the Logical Volume Manager (LVM) will be of great benefit to you. Among, many things LVM allows you to mirror, stripe, partition, create, grow, and delete volume groups and file systems on the disks.

With the **diag** command you can invoke a standard AIX diagnostics menu. This menu lists these aids that allow you to manage external disk operations.

A variety of scripts exist that have been developed for disk management.

17.3.2 HP-UX and Sun Solaris

HP and Sun can utilize SAM and Veritas Volume Manager. Solitus Disk Suite is only available for Solaris. These tools are and additional purchase and are used separately from the operating system management functions. These tool allows disk and data management functions. Figure 17-2 shows an example of the VERITAS Volume Manager GUI tool, that allows the administrator to manage the volumes and disk groups. This tool enables you to create and delete filesystems. Set sizes and permissions. Manage and monitor the ESS disks presented to the host on the host operating system level.

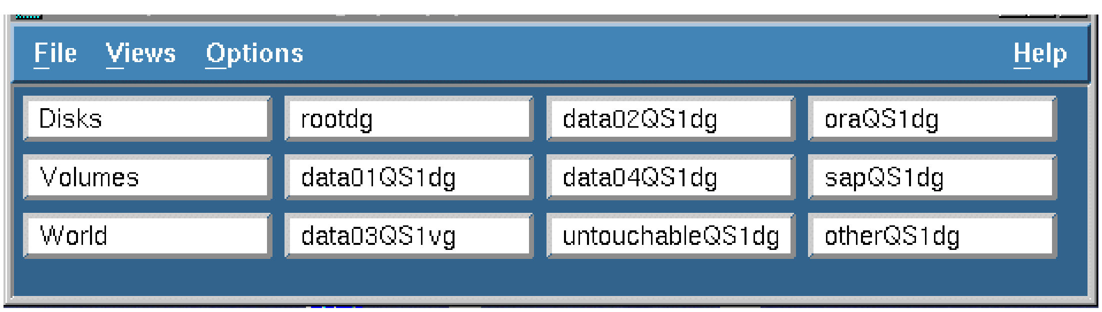

Figure 17-2 Veritas Volume Manager GUI

17.3.3 OS/400

There are several tools and commands on iSeries you can use for managing and monitoring disks. The Dedicated Service Tools (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 (see Figure 17-3).

```
                    Use Dedicated Service Tools (DST)

Select one of the following:

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

   Selection

      4

   F3=Exit     F12=Cancel
```

Figure 17-3 DST main menu

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 these options:

- Display disk configuration
- Work with disk configuration
- Work with disk unit recovery

Operations Navigator

The Operations Navigator provides a Windows interface to common iSeries management functions. Operations Navigator is supported on Windows 95, Windows 98, Windows NT 4.0.and Windows 2000.

Operations Navigator can be used to configure and manage ASPs comprised of both internal and external (ESS) disk units. However, as of V5R1, Operations Navigator provides a graphical display of only its internal disk units, and not of ESS hardware LUNs.

Other 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

With the Operations Navigator installed and configured, you can refer to the Operations Navigator online help for specific task-based information. Also you can visit the iSeries Info Center Web site by clicking this option at:

http://publib.boulder.ibm.com/pubs/html/as400/v5r1/ic2924/index.htm

At the Info Center you will find many examples on how to exploit the functions that Operations navigator provides for managing and monitoring.

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.

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.

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.

17.3.4 Windows 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 do the following:

- 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.

For additional information, please refer to your Windows Disk Administrators users guide.

17.3.5 Linux

With the Linux source code being freely available, several companies have developed different distributions of Linux. A distribution is a complete system, the key component of which is the Linux kernel. Other utilities, services, and various applications can be included as well, depending on the distribution and the intended use. You may want to check the tools available with the distribution you are using.

Bonnie

Bonnie is a performance measurement tool written by Tim Bray. Bonnie performs a series of tests on a file of selectable size (100MB default). Read and write access of different kinds are performed during the verifications. For each, the processed bytes per second and the CPU usage are measured. If you are interested in Bonnie, you will find more details on:

http://www.textuality.com/bonnie/advice.html

17.3.6 NUMA-Q

ConfigApp, Advanced Detection Availability Manager (ADAM) and VERITAS Volume Manager provide disk and data management tools.

Visit the following Web sites for more details on NUMA-Q:

http://techdocs.beaverton.ibm.com/

http://techdocs.beaverton.ibm.com/docs/ncrnab00/ch_5.htm#SE178140

The Advanced Detection Availability Manager (ADAM) subsystem provides NUMACenter systems management. ADAM is composed of the ADAM Server, the ADAM hardware, the ADAM software stack, and the ADAM Console. ADAM provides device access, management of servers and SAN elements, single point of access, resource monitoring, reconfiguration of components regardless of vendor, error and event consolidation, auto discovery of devices, and phone-out capability. For additional information on installation and use of ADAM, visit the Web at:

http://techdocs.beaverton.ibm.com/docs/ncagaa07/ncagaa07.htm

18

ESS Copy Services for open systems

PPRC and FlashCopy are ESS functions used for disaster recovery, data backup, creation of test data, and for data migration. In this chapter we examine some considerations that will help you when planning for the implementation of these functions in your open systems environment. For detailed information and recommendations on how to implement and use this powerful ESS functions you must refer to the IBM redbook *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757.

18.1 ESS Copy Services: considerations

ESS 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 or Fibre Channel 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 or Fibre Channel bus target/LUN group (LSS). CLI is provided to allow these functions to execute on supported host systems at the command line level. For a current list of CLI supported hosts please refer to Table 13-1 on page 207.

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

- The FlashCopy license feature ordered must be equal to or greater than the total capacity of the ESS
- The PPRC license feature must be equal to or greater than the total capacity of the ESS. The PPRC feature must also be purchased and installed on both the primary and secondary ESS

You should consider the following when planning:

- Only one FlashCopy at a time can be active on a volume, however, you can perform a PPRC concurrently with FlashCopy on the same volume. The designated FlashCopy target volume cannot be a primary volume in a PPRC volume pair.
- The primary and secondary volumes must reside within an ESS, for both PPRC and FlashCopy. You cannot FlashCopy from an ESS to different external device type.
- You need to manage PPRC using the ESS Copy Services Web interface, 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.
- The source and target logical volumes must be the same, or larger in size.
- If you want to use PPRC or FlashCopy to do a hardware copy of the disks attached to the iSeries, you must mirror the LSU from the internal drive into the ESS to ensure the whole single level storage is copied.
- There is a maximum of 2000 volumes allowed for a ESS Copy Services server. This number includes all the primary and secondary PPRC volumes plus all the source and target FlashCopy volumes.

18.2 Preliminary setup

In order to use ESS Copy Services you must configure one ESS to be the primary ESS Copy Services server. All information related to the copy services is stored in this ESS, such as volumes and their state, ESCON connectivity between ESS, and much more.

On each ESS that is configured to use ESS Copy Services there is a client running. Whenever the copy services configuration changes this client notifies the ESS Copy Services server of the changes.

Optionally, there could be one ESS defined as the backup ESS Copy Services server. In case the primary ESS Copy Services server is lost, the backup server could be used for controlling the ESS Copy Services. Once the primary ESS is up again, the backup server will notify all clients, so the clients can be switched back to the primary ESS Copy Services server.

We recommend that you have the primary and backup ESS Copy Services servers on different sites if your ESS Storage Network spans multiple locations. In case the entire side with the primary ESS Copy Services server is going down for whatever reason, the backup ESS Copy Services server can keep copy services alive. You may also consider to run the primary ESS Copy Services server on the remote site ESS cluster.

The information on the primary and backup ESS Copy Services server has to be specified on each cluster of all the Enterprise Storage Servers that are going to be used for ESS Copy Services. We recommend that you specify a backup for ESS Copy Services whenever possible. This task is done by an IBM SSR with the Master Console.

18.3 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 operates only over ESCON connections, and therefore, even if it is being used in an open systems environment, ESCON adapters and cabling are required.

PPRC is supported at distances over 103 km but requires IBM approval, which can be requested by submitting a Request for Price Quotation (RPQ).

ESTABLISH/RESYNC with INRANGE allows for new implementations:

Enhancements to the INRANGE 9801 SNS exploit the PPRC Establish Copy design, allowing data to stream across the network without unnecessary intermediate messages (protocols) for data transmission over long distances. These enhancement enable PPRC solutions with the INRANGE 9801 for data movement and migration; remote backup (in combination with FlashCopy); and long distance disaster recovery in two ways: sending log files, or implementing RESYNCH/SUSPEND loops (if the primary site can tolerate application quiesce).

18.3.1 Rules for configuring PPRC links

In this section we provide rules for configuring PPRC links.

Establishing ESCON connections between two ESS subsystems

These are some requirements to observe:

- A primary LSS can be connected via ESCON links to up to four secondary LSSs.
- A secondary LSS can be connected to any number of primary LSSs, limited by the number of ESCON links available.
- PPRC links are unidirectional, because the ESCON port at the primary ESS is reconfigured to act like an ESCON channel in a host S/390 processor. The primary ESCON port is dedicated to PPRC.
- The ESCON protocol has been streamlined with less handshaking and larger frames transmitted between ESS.

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.

18.3.2 Planning for PPRC

PPRC is possible only between Enterprise Storage Servers. Other disk storage units that support PPRC can also communicate to the same type of unit only.

You need to have the PPRC feature purchased and PPRC-capable microcode activated on all ESS units that will be used for PPRC.

PPRC operates at a volume level from one LSS to another LSS.That means you need to have the target volumes available on the secondary ESS, and you need to identify the LSSs where the primary and secondary volumes are located.

ESCON connections have to be configured between the units. There can be up to eight ESCON links between the subsystems. A primary ESS can communicate with up to four secondary ESSs. A secondary ESS can be connected to any number of ESS primary subsystems.

You will need to purchase ESCON cables and possibly some other equipment, depending on the distance between the primary and the secondary ESS. However, each of the ESS Copy Services servers can only control two ESSs. Therefore, to have one primary ESS communicate with more than one secondary ESS, this will mean that multiple ESS Copy Services servers will have to be configured.

The ESS units involved in PPRC must be connected with their standard Ethernet and TCP/IP to the units that are the primary and backup ESS Copy Services servers. All ESS cluster hostnames, including its own, must be added to the cluster hostname list, the /etc/hosts file, during installation. This is configured by the IBM SSR during the installation of the PPRC feature.

A browser for the Web interface has to be installed on the machines that will be used to control PPRC with the ESS Copy Services Web Interface.

If you plan to use CLI, install the Java Developers Kit (JDK) JDK version 1.1.8 on the machines that will run the CLI commands. *Host Systems Attachment Guide,* SC26-7296 for the latest recommended JDK revision.

You need to check that the CRIT=YES - Light or Heavy options have been set correctly in the ESS VPD by your hardware specialist.

You 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 ESS Copy Services primary.

When planning your secondary ESS volume layout for PPRC, optimize your disk capacity. It is important to realize that the capacity needed on the secondary ESS for disaster recovery may not have to be initially as large as the primary ESS. A disaster recovery plan (DRP) requires significant investment financially in technology, people, and process. Every company will be different, but the I/T components of a disaster recovery plan are essentially driven by the applications and data you require for business continuity, should a disaster occur. Some applications and data will be more critical than others. An organization will typically require its core business systems to be available in a short time, whereas less critical systems quite possibly could be restored over a number of days.

Bearing in mind that the disk space you need for PPRC secondary volumes is real disk space, size your secondary ESS based on your critical business requirements, possibly with some headroom for applications of intermediate importance. Create PPRC pairs for the critical data so that is copied in real time. Then, if a disaster happens, you will have the core systems available on the secondary copies. After the initial recovery priorities have been handled, you can add more disk ranks for the applications of lower importance and restore them from tape.

18.3.3 PPRC Read from Secondary option

This function allows a host to read from a Secondary PPRC volume, once the PPRC pair has reached full duplex state. This function is enabled on open system volumes only, and the ESS LIC level required is 1.5. This option can be enabled when you select the copy options (see Figure 18-1).

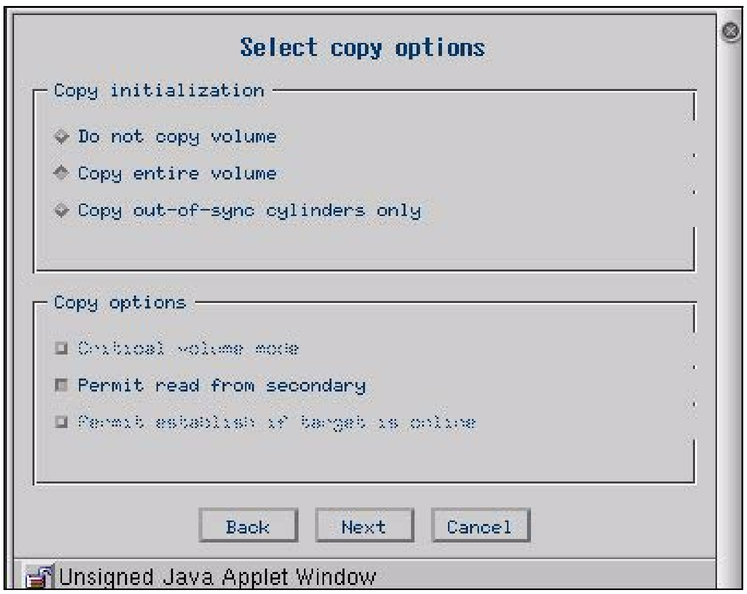

Figure 18-1 ESS Copy Services: Read from Secondary allowed

The following rules apply:

- Read from Secondary will not be allowed unless designated in the Establish PPRC pair command.
- Read from Secondary will not be allowed until the pair has transitioned to full duplex state. The pair must have reached full duplex state in order for Read from Secondary to be successful. (if the pair suspends after reaching full duplex state, Read from Secondary will still be allowed).
- Read from Secondary permission may be “added” or “removed” by issuing another Establish Pair command with the desired Read from Secondary setting.
- Read from Secondary is only allowed on open volumes. If this request is made on a S/390 volume, the indication will be stored, but S/390 read commands will still be rejected.

> **Note:** The value of Read from Secondary is that it allows a server attached to a PPRC secondary device to discover the physical configuration on the PPRC secondary volumes. This will reduce the amount of time to switch over to a recover site after a disaster while PPRC is active. You must then terminate PPRC and launch the application from the secondary site.

18.3.4 How to invoke PPRC

There are two different methods of using the ESS Copy Services in the open systems environment:

- Using the ESS Copy Services Web interface
- Using the Java-based Command Line Interface (CLI)

To compliment the information presented in this section, please refer to the following publications: *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface Reference,* SC26-7434 and *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757.

ESS Copy Services Web interface

The ESS Copy Services function that runs in the ESS provides a Web browser interface that can be used to control the PPRC functions in the open systems environments. You will be able to establish PPRC pairs in three different ways:

- From the *Volumes* panel — based on volumes
- From the Storage Servers panel — based on entire logical subsystems
- From the Task panel — once a task for PPRC is created

The ESS Copy Services panel shows volumes (LUNs) by their ESS internal serial numbers. You will first have to find out what is the serial number of the volumes you intend to copy (source and target) from your operating system, in order to identify the volume on the ESS Copy Services *Volume* panels.

To invoke PPRC, you need to perform the following tasks from the ESS 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 an LSS.
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-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 systems 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 LSSs, volumes or paths.

Detailed information regarding each panel described above can be found in the *Web Interface Users Guide*, SC26-7346.

Command Line Interface (CLI)
The CLI interface allows administrators to execute Java-based ESS Copy Services commands from a command line. This command line interface is currently available for the following operating systems: AIX, Sun Solaris, HP-UX, Windows NT and 2000, Intel-based Linux (See Section 13.5, "Command Line Interface (CLI)" on page 209 for information on operating system support).

18.4 FlashCopy

FlashCopy makes a single point in time (T0) copy of a LUN. The target or copy LUN is almost immediately available after the command has been processed. For detailed information on how FlashCopy works and for considerations when planning for FlashCopy implementations we recommend you refer to the publication *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757.

18.4.1 Invoking FlashCopy

There are two different methods of using the ESS Copy Services in the open systems environment:

- **Using the ESS Copy Services Web interface:** The ESS Copy Services function provides a Web browser interface that can be used to control the FlashCopy functions in the open systems environment. For this situation you will be establishing a FlashCopy pair using the *Volumes* panel, or using the *Task* panel that the Web browser presents you. Detailed information regarding the use of each panel can be found in the *Web Interface Users Guide*, SC26-7346.
- **Using the Java-based Command Line Interface (CLI):** The CLI interface allows administrators to execute Java-based ESS Copy Services commands from a command line. For this situation you will be starting a pre-defined task from the command line or from your own customized scripts. This command line interface is currently available for the following operating systems: AIX, Sun Solaris, HP-UX, Windows NT and 2000, Intel-based Linux (See Section 13.5, "Command Line Interface (CLI)" on page 209 for information on operating system support).

To compliment the information presented in this section, please refer to the following publications: *IBM TotalStorage Enterprise Storage Server Copy Services Command-Line Interface Reference,* SC26-7434 and *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757.

18.4.2 Invoking FlashCopy scenario for AIX

The first step requires you to identify the internal serial number of the logical source and target disks on the host. For example, the following command is issued on AIX:

```
lscfg -vl hdisk#| grep Serial
```

Here is an example of the returned output:

```
Serial Number...............30115234
```

The first 3 digits represent the LUN's hex ID. The last 5 digits represent the serial number of the ESS. The first 3 digits are what you are looking for to identify the Source LUN serial number.

We suggest that you create a LUN mapping table that represents the layout of the ESS, such as in Figure 15-2 on page 237. This will help you strategically see at a glance the LSSs and individual LUNs.

To invoke FlashCopy, use the VOLUMES panel to find and select the source and target volumes. Left-click on the LUN you have identified as the source (the volume being copied from). Right-click on the LUN that you have identified as the target (the volume being copied to). Right-click one more time to execute the wizard window.

Once presented with the wizard window, select **Establish FlashCopy pair** or **Withdraw FlashCopy pair**. Figure 18-2 represents establishing a FlashCopy pair. Click **Next** to go to the next panel. Figure 18-3 shows the wizard window with No Copy Background selected. This means that the data will only be accessible to the target LUN as long as the pair remains established. In the first wizard window, notice on the ESS Master Console, that if the button appears to be pushed in, then it is selected. See Figure 18-2.

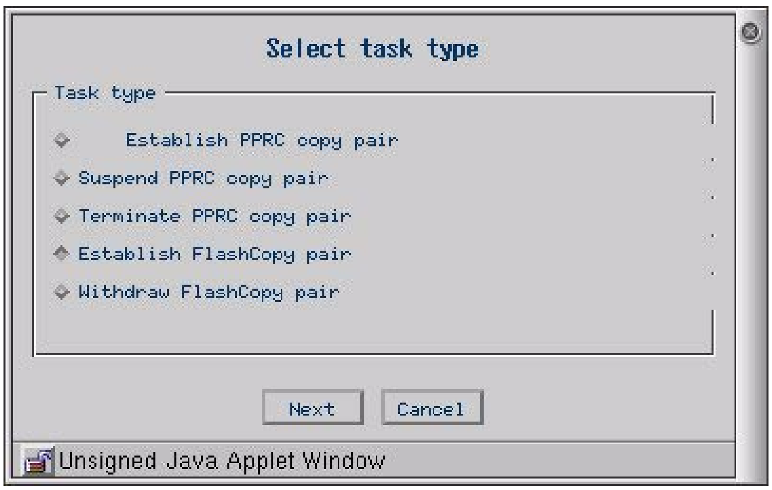

Figure 18-2 ESS Copy Services: First wizard window

From the ESS Copy Services panel, choose:

Volumes -> LSS -> source LUN (with a left-click) **-> target LUN** (with a right click) **->** (right-click again) **-> Establish FlashCopy pair -> Next ->** choose copy option (see Figure 18-3) **-> Next ->** name the task (see Figure 18-4), and either **Run** or **Save** the task.

Figure 18-3 Establishing a pair with No Copy background

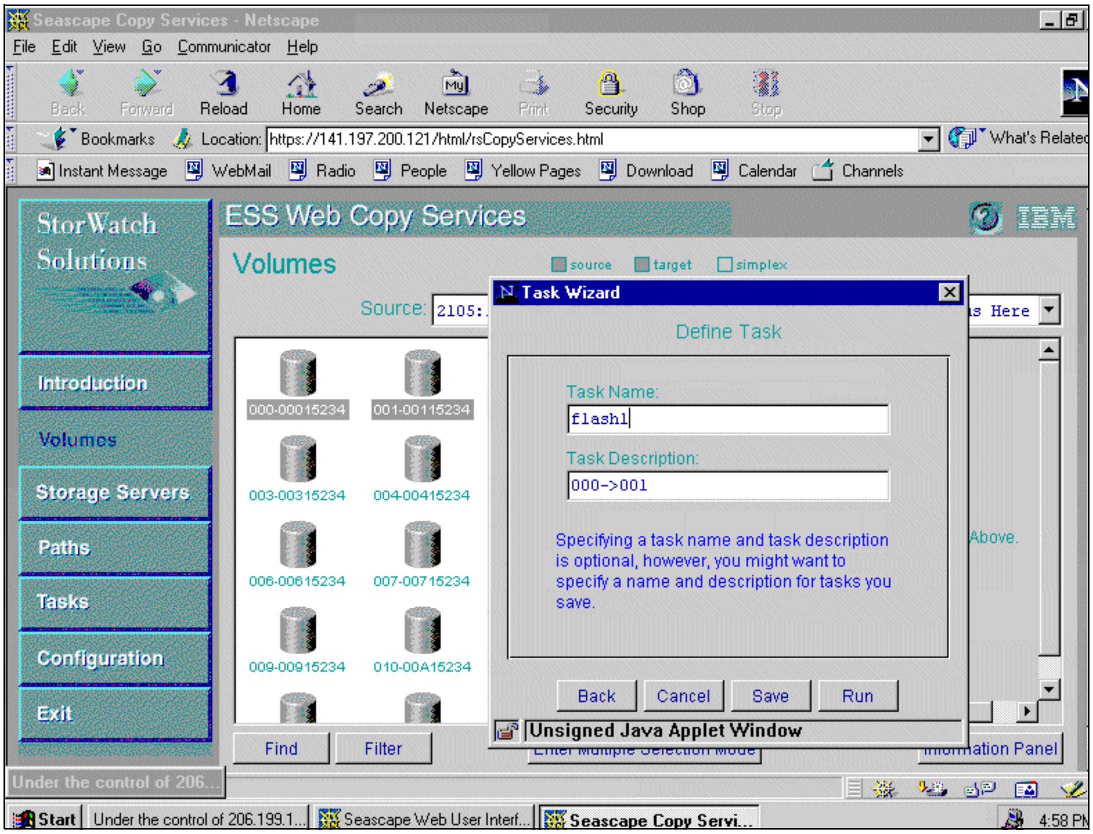

Figure 18-4 Wizard window with example task name and description

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. Figure 18-4 shows an example task name and description.

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 Copy Services Web interface panels, 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 Copy Services Web interface panels. If NOCOPY was not specified, the relationship ends automatically once the background source copy has been completed.

We will now summarize these operations.

To invoke a no-copy background FlashCopy pair

1. Log into ESS Copy Services.
2. From the ESS Copy Services panel, select the **Volumes** panel.
3. Click the desired LSS (from the source **Select Focus here**, drop-down box).
4. Left-click on the source LUN.
5. Right-click on the target LUN.
6. Right-click again.
7. The wizard window will now appear. Select **Establish FlashCopy pair**.
8. Click **Next**.
9. Select, **Do not perform background copy if checked**.
10. Click **Next**.
11. The wizard window will now appear. Enter in the task name and task description.
12. Click **Save,** to run it at another time, or from the command line on the target, or click **Run**, to run it immediately.

To invoke a withdraw

For every FlashCopy pair that you establish with the No Copy background option, you must also create a task to withdraw that pair.

1. Log into ESS Copy Services.
2. From the ESS Copy Services panel, select the **Volumes** panel.
3. Click the desired LSS.
4. Left-click on the source LUN.
5. Right-click on the target LUN.
6. Right-click again.
7. The wizard window will now appear. Choose **Withdraw the Established Pair**.
8. Click **Next**.
9. Enter the task name and task description.
10. Click **Save,** to run it at another time, or from the command line on the target, or click **Run**, to run it immediately.

For detailed instructions on setting up FlashCopy pairs and withdrawing them, refer to the IBM publication *IBM TotalStorage Enterprise Storge Server Web Interface User's Guide,* SC26-7346.

A

ESS configuration planning process

This appendix describes the procedures and support available for the IBM Field Technical Support Specialist (FTSS) and IBM Business Partners (BP). This supplement is designed to help FTSSs and BPs to support the customers in the installation planning of the ESS. We cover the following topics:

- Considerations for planning the storage capacity for each host
- Where you should start:
  - Gathering the information to submit the ESS Configuration worksheets
  - Gathering information to complete the Communication Resources worksheets
- Where you can find the reference material
- Submitting the ESS Configuration worksheet

Considerations for planning

You should ask yourself the following questions when planning the total storage capacity of the ESS:

- How many TB of usable disk space do you have? How many disk groups do you have?
- What is the total capacity in GB of each disk group. Remember that with the disk drive intermix, your disk groups will vary in size. You can have 8 packs of 9.1, 18.2, 36.4, and 72.8 GB sizes. Also remember that you can upgrade the installed eight-packs to the 36.4 GB and 72.8 GB higher capacity disk drives.
- What size LUNs should you carve out of each disk group to meet the needs of the applications data requirements?
- How much disk space will you need for Flash Copy? A good rule of thumb is double the amount of disk space that you will be using to flash and designating it for FlashCopy. Remember that the target LUN in the ESS must be in the same LSS and be of equal or greater size than the source LUN that you will be flashing.
- How should you lay out LUNs for best performance and FlashCopy.
- What type of hosts are you connecting to the ESS and how many do you have?
- How many host adapter ports will you be connecting to the ESS or SAN Fabric switch?
- Will you be using SDD?
- Will you be using PPRC or XRC copy services?
- Should you install CLI?
- What is the order number, feature code, and serial number of the ESS?

> **Note:** If your ESS is coming with a mix of drive sizes, then the feature code will not be valid. You instead need the order number readily accessible. The minimum quantity of Eight Packs per order is two.

Where should you start?

After considering the questions asked previously, you should start gathering information about each topic.

You should contact your IBM Field Technical Support Specialist (FTSS) or IBM Business Partner (BP).

Gathering information to submit the ESS Configuration worksheets

You should download the ESS Configuration worksheets from the web at the following site:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

Click on the *ESS Configuration Planner* and look under Appendix A, B, C and D.

Obtain the order number, feature code and Serial number of the ESS from the Sales Representative, Project Manager, or Solutions Architect. It would be helpful for you to:

1. Involve the Project office in San Jose to help with the overall ESS Plan.
2. Gather the names of the resources you need to help you configure the ESS.
3. Identify the IBM SSR
4. Identify the storage specialist.
5. Identify the various team members, such as Disk Admin/Sys Admin, SAP/Basis, DBA members and Network team.
6. Receive the file system requirements.

Completing the Communication Resources worksheets

You should download the Communication Resource worksheets from the Web at the following site:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

Click on the *Introduction and Planning Guide* and look under Appendix A.

You should decide on connecting through a private or public network or LAN.

You will probably need to gather the information from all of the team members. You will need IP addresses, phone numbers, etc. Refer to the worksheets for the information needed.

After completing the Communications worksheets you must E-mail or send them to the IBM SSR that will be installing the ESS.

Where you can find the reference material

You can find information on all of the above questions and topics by visiting the Web site:

http://ssddom02.storage.ibm.com/disk/ess/documentation.html

Topics and subjects include the following:

- ESS overview
- ESS spec Sheet
- Host systems Attachment Guide
- IBM TotalStorage Redbooks
- Introduction and Planning Guide
- SCSI Command Reference Guide
- System/390 and open systems
- Users Guide
- Web Interface User's Guide
- ESS Copy Services Command-Line Interface User's Guide
- ESS Configuration Planner (Includes a chart with various feature codes and capacities).
- Sample storage configuration worksheets.

Submitting the ESS Configuration worksheets

This is the process for the Americas and it may differ from geography to geography.

Configuration Worksheet Assistance details

The Configuration HelpLine is available to complete the Logical Configuration for the ESS install. Follow these steps to contact the Configuration HelpLine:

1. Call (800) 237-5511.
2. Select Option #0 (there is a long menu, so try it as soon as you get to this spot).
3. The Agent will ask for the following information:
   a. The customer number or order number and serial number of the ESS.
   b. The caller's name and telephone number (also be sure to give them the email address of the SSR who is to receive the completed configuration.)
   c. The system or component type: Say “Shark” or “ESS” (these are the “trigger words”).
   d. The severity of the problem / question: Reply “Sev 1” if you need immediate response, otherwise please use “Sev 2”.
   e. The topic or nature of the question: Reply “Logical Configuration Help”.
4. The Agent will set up a PMR number and Branch number, assigned to the question (Important! These may be used later to check on the status of the question).
5. The Agent will forward the request for a call-back to the Toronto Configuration Help Desk.
6. If you can't access the 800 number, you can use the following tieline or outside line numbers: 8/367-5079, 770-858-5079.
7. After receiving the PMR number, the Advanced Logical Configuration Team member must fill out the Web Configuration Worksheets at:

   `http://ssddom01.storage.ibm.com/ess/sirtstor.nsf/Welcome/?OpenForm`

- You must have a PMR opened with the IBM Support Center in order to proceed with submitting the worksheets to Toronto.
- The worksheets are then converted to a different format for the IBM SSR to use for the ESS installation.
- The completed worksheets are either returned to you or the IBM SSR to do the logical configuration.
- The completed communication resource worksheets are also used by the IBM SSR.
- You will need to provide the SSR and Systems Administrator with cabling diagrams for the host adapter attachments to the ESS.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see "How to get IBM Redbooks" on page 298.

- *IBM TotalStorage Enterprise Storage Server,* SG24-5465
- *Implementing Linux with IBM Disk Storage,* SG24-6261
- *IBM StorWatch Enterprise Storage Server Expert Hands-on Usage Guide*, SG24-6102
- *Implementing Fibre Channel Attachment on the ESS,* SG24-6113
- *ESS Solutions for Open Systems: Compaq Alpha Server, HP, and Sun,* SG24-6119
- *iSeries in Storage Area Network A Guide to Implementing FC Disk and Tape with iSeries,* SG24-6220
- *Implementing ESS Copy Services on UNIX and Windows NT/2000,* SG24-5757
- *Implementing ESS Copy Services on S/390,* SG24-5680
- *Planning for IBM Remote Copy,* SG24-2595

Referenced Web sites

These Web sites are also relevant as further information sources:

- ESS Web site

  `http://www.storage.ibm.com/hardsoft/products/ess/ess.htm`
- ESS manuals

  `http://ssddom02.storage.ibm.com/disk/ess/documentation.html`
- ESS supported servers

  `http://www.storage.ibm.com/hardsoft/products/ess/pdf/1012-01.pdf`
- SDD support:

  `http://ssddom01.storage.ibm.com/techsup/swtechsup.nsf/support/sddupdates`
- Host System Attachment Guide

  `http://www.storage.ibm.com/hardsoft/products/ess/pubs/f2ahs04.pdf`

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Special notices

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 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.

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 other countries licensed exclusively through The Open Group.

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

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

Index

F

H

I

T

U

V

W

X

Z

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set

IBM TotalStorage Enterprise Storage Server: Implementing the ESS in Your Environment

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set

IBM TotalStorage Enterprise Storage Server

Implementing the ESS in Your Environment

Learn how to install, tailor, and configure the ESS

Plan for migrating data and changing configurations

Know the tools for monitoring and managing the ESS

Since the initial availability of the IBM TotalStorage Enterprise Storage Server (ESS) there have been many modifications and additional functionality. The original ESS Implementation Guide was written in November 1999. The 2002 version is an updated guide that includes all of the changes since then.

This IBM Redbook is a guide for the installation, implementation, and administration activities of the ESS in both the S/390 and open systems environments. It will help you plan and accomplish the installation, tailoring, and configuration of the ESS in your environment. It explains how you can use the functions available for the ESS to efficiently manage your disk storage data as well as the ESS, once it is operative.

We cover the latest announcements on the ESS: disk capacity intermix; 72.8 GB capacity disk drive; flexible configurations; Control Unit Initiated Reconfiguration (CUIR) support; large volume support (LVS); read from secondary; ESS Master Console; Subsystem Device Driver (SDD) and the Command Line Interface (CLI) support for additional operating systems; INRANGE Channel Extender support; TPF support for PPRC and FlashCopy.

We also provide information on the new models F10 and F20; FICON native host attachment; new Fibre Channel/FICON host adapters (short wave and long wave); Linux support for Intel-based servers and zSeries servers; iSeries and AS/400 support for copy services; and new cache options.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-5420-01 ISBN 0738424390