

IBM Enterprise Storage Server™



User's Guide

2105 Models E10, E20, F10, and F20

IBM Enterprise Storage Server™



User's Guide

2105 Models E10, E20, F10, and F20

Note:

Before using this information and the product it supports, read the information under "Safety and environmental notices" on page vii and "Notices" on page 49.

Third Edition (November 2000)

This edition replaces the previous edition of SC26-7295-01.

Publications are not stocked at the address given below. If you want additional IBM publications, ask your IBM representative or write to the IBM branch office serving your locality.

A form for your comments is provided at the back of this document. If the form has been removed, you may address comments to:

International Business Machines Corporation
RCF Processing Department
G26/050
5600 Cottle Road
San Jose, CA 95193-0001
U.S.A.
FAX: **1-800-426-6209**

You can also send your comments electronically to:

starpubs@us.ibm.com

For additional information on IBM storage products, visit the Web site at: <http://www.ibm.com/storage>

Contents

Figures	v	Adapters for UNIX systems	27
Safety and environmental notices	vii	Transferring data for AS/400 systems.	27
Product recycling	vii	Transferring data for UNIX systems	28
Disposing of products.	vii	Data migration affect on data availability	34
About this guide.	ix	Data migration from System/390 operating systems	35
Who should use this guide	ix	Publications about data migration	35
Frequently used terms	ix	Guidelines for migrating data to an ESS.	35
Publications	xii	Data migration and other services	37
The ESS product library	xii	Chapter 5. Problem analysis	39
Ordering publications.	xiii	Analyzing problems that occur during operation	39
Related publications	xiii	Service information messages for System/390 hosts	39
Web sites	xv	Analyzing SNMP trap information	40
Summary of Changes.	xvii	Analyzing e-mail messages	41
Chapter 1. Introduction	1	Informational messages	41
IBM Enterprise Storage Server overview	1	Error messages	41
Host systems supported by the ESS.	3	Status	44
Web interfaces	4	Analyzing pager messages, numeric and alphanumeric.	44
Service interface	9	Obtaining service	47
Fibre-channel connection	10	Technical support	47
Introducing the IBM Subsystem Device Driver.	12	Warranty support or maintenance	47
Chapter 2. Setting up and operating the Enterprise Storage Server	15	Notices	49
Using the control panel	15	Trademarks	50
Turning the power on and off	16	Electronic emission notices	51
Control panel for the expansion enclosure	17	Federal Communications Commission (FCC) statement	51
Attaching the ESS Network to your local area network	17	Industry Canada compliance statement	51
Chapter 3. Using parallel access volumes for System/390 hosts	19	European community compliance statement	51
Overview of parallel access volumes	19	Japanese Voluntary Control Council for Interference (VCCI) class A statement.	52
Parallel access volumes requirements.	19	Korean government Ministry of Communication (MOC) statement	52
Overview of the input/output configuration program	20	Taiwan class A compliance statement.	52
IOCP example	20	IBM agreement for licensed internal code	53
Guidelines for defining the logical control unit	22	Actions you must not take	54
Chapter 4. Migrating data to the IBM Enterprise Storage Server	25	Statement of limited warranty	54
Migrating data from UNIX, Windows NT, Windows 2000 and OS/400	25	Production status	54
Replacing existing storage	26	IBM warranty for machines	54
Reformatting existing disk drive modules	26	Warranty service.	55
Host system views of the IBM Enterprise Storage Server	26	Extent of warranty	55
		Limitation of liability	56
		Glossary	57
		Index	69

Figures

1. Models E10, E20, F10, and F20; front and rear views	2	7. Multipath connections between a host server and ESS LUNs	13
2. Expansion enclosure, front and rear views	2	8. Control panel switches and indicators for Models E10, E20, F10, and F20	16
3. Connecting to ESS Web Copy Services server groups	8	9. Control panel switches and indicators for an expansion enclosure.	17
4. Point-to-point topology	10	10. IOCP example, Part 1 of 2.	21
5. Switched fabric topology	11	11. IOCP example, Part 2 of 2.	22
6. Arbitrated loop topology	12		

Safety and environmental notices

This section contains information about:

- Safety notices that are used in this guide
- Environmental guidelines for this product

To find the translated text for a danger or caution notice:

1. Look for the identification number at the end of each danger or caution notice. In the following examples, look for the numbers **1000** and **1001**.

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury.

1000

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury.

1001

2. Find the number that matches in the *IBM Storage Solution Safety Notices* book, GC26-7229.

Product recycling

This unit contains recyclable materials. Recycle these materials at your local recycling sites. Recycle the materials according to local regulations. In some areas, IBM provides a product take-back program that ensures proper handling of the product. Contact your IBM representative for more information.

Disposing of products

This unit may contain batteries. Remove and discard these batteries, or recycle them, according to local regulations.

About this guide

This publication provides the procedures for setting up and using the IBM® Enterprise Storage Server (ESS). It includes an overview of the ESS, and guidelines for:

- Setting up and operating the ESS
- Use of parallel access volumes for System/390® hosts
- Migration of data to the ESS
- Problem analysis

Who should use this guide

This publication is for the use of customers or their representatives who install, set up, and use the ESS.

DANGER

Because of hazardous voltages, IBM authorizes only experienced service personnel to open the ESS.

1012

Frequently used terms

The following terms appear in this publication to describe the IBM Enterprise Storage Server (ESS) and its configuration. Also see the “Glossary” on page 57 for additional terms and definitions.

Term	Definition
array	An arrangement of related disk drive modules (DDMs) that you have assigned to a group. Also referred to as a rank. For the ESS, an array is a group of up to eight DDMs.
cluster	A partition of a storage server that is capable of performing all functions of a storage server. When a cluster fails in a multiple-cluster storage server, the cluster that remains in the configuration takes over the processes of that cluster.
controller image	A System/390® term for a logical subsystem that you access with an ESCON® I/O interface. One or more controller images exist in each controller. Each image appears to be an independent controller, but all images share a common set of hardware facilities. The ESS can emulate 3990-3, 3990-6, or 3990-3 TPF controller images.
count key data (CKD)	A System/390 term for a logical device that specifies the format of the logical data units on the device. The logical data unit is a track that contains one or more records, consisting of a count field, a key field (optional), and a data field (optional).

device adapter	A physical component of a storage controller that attaches to one or more interfaces. It enables communication between the interfaces and the associated storage devices.
destage	The process of writing modified storage from cache to the disk drive module (DDM).
disk drive module (DDM)	The primary, nonvolatile storage medium that you use for any host data that is stored within a storage server. The number and type of storage devices within a storage server may vary.
Enterprise Systems Connection architecture (ESCON)	
	<ol style="list-style-type: none"> 1. An ESA/390 computer-peripheral interface. The I/O interface uses ESA/390 logical protocols over a serial interface. 2. A set of IBM products and services that provide a dynamically connected environment within an enterprise.
ESCON	See <i>Enterprise Systems Connection architecture</i> .
ESCON channel	A channel that has an ESCON channel-to-controller I/O interface that uses optical cables as a transmission medium.
ESCON host systems	System/390 hosts that you attach to the ESS with an ESCON adapter. Such host systems run on Multiple Virtual Storage (MVS), virtual machine (VM), VSE, or TPF operating systems.
fabric	In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are <i>cascaded</i> .
fibre-channel	<p>Fibre-channel is an architecture that supports full-duplex communication over a serial interface that configures attached units to a communication fabric.</p> <p>The ESS supports fibre-channel communications for open system hosts through its fibre-channel adapters.</p>
fixed-block (FB) devices	An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This is the required format of the logical data units for host systems that you attach with a small computer system interface (SCSI) or fibre-channel interface.
gigabyte (GB)	When referring to hard-disk-drive capacity, GB means 1 000 000 000 bytes; total user-accessible capacity may vary depending on operating environment.

host adapter	A physical component in a storage server that enables the server to attach to one or more host I/O interfaces.
IBM Enterprise Storage Server (ESS)	A member of the Seascape™ product family of storage servers and attached storage devices (disk drive modules). The ESS provides integrated caching and RAID support for the disk drive modules (DDMs). The DDMs are attached through a serial storage architecture (SSA) interface.
logical device	A pointer to a logical volume on a storage server. The pointer enables a host system to access the logical volume. For the ESS, the logical device represents a logical volume.
logical subsystem (LSS)	<p>The LSS is a logical “container” for RAID-5 arrays, disk groups, and the volumes that are located on those arrays or disk groups. One SSA Adapter accesses an LSS primarily. One or more logical subsystems can exist on the ESS. In general, the ESS associates a given set of devices with only one logical subsystem.</p> <p>For System/390 hosts, a logical subsystem represents a controller.</p>
logical unit	The Small Computer System Interface (SCSI) term for a logical disk drive.
logical unit number (LUN)	LUN is a SCSI term for the field in an identifying message that is used to select a logical unit on a given target.
logical volume	<p>The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. For the ESS, you define this unit of storage; the logical volume resides on a RAID-5 array, and is spread over 6 + P or 7 + P drives, where P is parity. A logical volume can also reside on a non-RAID array on one storage device.</p> <p>For CKD, you can define the logical volume size by the device emulation mode (3390 or 3380 track format). For UNIX®, Microsoft® Windows NT®, and Microsoft Windows 2000® operating systems in the fixed-block format, the size is 0.5 GB to 224 GB. For AS/400® operating systems in the fixed-block format, the size is 4.19, 8.59, 17.55, or 36 GB.</p> <p>Note: The AIX® operating system views a logical volume as a logical disk or a hard disk (hdisk).</p>
SCSI	See <i>small computer system interface</i> .
SCSI host systems	Identifies host systems that you attach to the ESS with a SCSI interface. Such host systems run on UNIX, OS/400, Windows NT, Windows 2000, or Novell NetWare operating systems.

SCSI ID	A unique identifier (ID) assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs. A SCSI device is either an initiator or a target.
service personnel	Individuals or a company that you have authorized to service your ESS. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs your ESS.
Small Computer System Interface	An American National Standards Institute (ANSI) standard for a logical interface to computer peripherals and for a computer-peripheral interface. The interface uses a SCSI logical protocol.
staging	To move data from an offline or low-priority device to an online or higher priority device is called staging.
storage server	A unit that manages attached storage devices and provides access to the storage or storage-related functions for one or more attached hosts.
terabyte (TB)	When referring to hard-disk-drive capacity, TB means 1 000 000 000 000 bytes; total user-accessible capacity may vary depending on operating environment.
volume	Refers to a <i>logical</i> volume.

Publications

This section describes the ESS library and publications for related products. It also gives ordering information for these publications.

The ESS product library

The ESS is an IBM Seascope architecture-based product. See the following publications for more information on the ESS:

- *IBM Enterprise Storage Server Introduction and Planning Guide, GC26-7294*
This guide introduces the product and lists the features you can order. It also provides guidelines on planning for the installation and configuration of the ESS.
- *IBM Enterprise Storage Server User's Guide, SC26-7295*
This guide provides instructions for setting up and operating the ESS and for analyzing problems.
- *IBM Enterprise Storage Server Host Systems Attachment Guide, SC26-7296*
This book provides guidelines for attaching the ESS to your host system and for migrating from SCSI to fibre-channel attachment.
- *IBM Enterprise Storage Server SCSI Command Reference, SC26-7297*

This book describes the functions of the ESS. It provides reference information for UNIX and Application System/400® (AS/400®) hosts, such as channel commands, sense bytes, and error recovery procedures.

- *IBM Enterprise Storage Server System/390 Command Reference*, SC26-7298

This book describes the functions of the ESS and provides reference information for System/390 hosts, such as channel commands, sense bytes, and error recovery procedures.

- *IBM Storage Solutions Safety Notices*, GC26-7229

This book provides translations of the danger notices and caution notices that IBM uses in ESS publications.

- *IBM Enterprise Storage Server Web Interface User's Guide*, SC26-7346

This book provides instructions for using the IBM StorWatch™ Enterprise Storage Server Web interfaces, ESS Specialist and ESS Web Copy Services.

- *IBM Enterprise Storage Server Configuration Planner*, SC26-7353

This book provides work sheets for planning the logical configuration of the ESS. This book is not available in hardcopy. This book is only available on the following Web site:

www.storage.ibm.com/hardsoft/products/ess/refinfo.htm

- *IBM Enterprise Storage Server Quick Configuration Guide*, SC26-7354

This booklet provides flow charts for using the StorWatch Enterprise Storage Server Specialist. The flow charts provide a high-level view of the tasks the IBM service support representative performs during initial logical configuration. You can also use the flow charts for tasks that you might perform during modification of the logical configuration.

The hardcopy of this booklet is a 9-inch × 4-inch fanfold.

- *IBM Enterprise Storage Server Service Guide*, SY27-7605

This service guide is for the service support representatives who are trained to install and repair the IBM 2105 Models E10/E20 and F10/F20. This service guide is broken into three volumes that contain the following chapters:

- Volume 1: Information, Start, and Isolation
- Volume 2: Repair, Install, and Configure
- Volume 3: Locations, Service Terminal, Codes, Power

Ordering publications

All of the publications that are listed in “The ESS product library” on page xii are available on a compact disc that comes with the ESS, unless otherwise noted. You can also order a hardcopy of each publication. For publications on compact disc, order *IBM Enterprise Storage Server Customer Documents*, SK2T-8770.

All the customer documents listed above are available on the following ESS Web site:

www.storage.ibm.com/hardsoft/products/ess/refinfo.htm

Related publications

The following related publications are also available:

- *DFSMS/MVS® Version 1 Release 5 Software Support for the IBM Enterprise Storage Server*, SC26-7318

This book describes the changes you need to make to operating systems for System/390 hosts to support the ESS functions.

- *IBM Storage Area Network Data Gateway Installation and User's Guide, 2108 Model G07, SC26-7304*

This book introduces the IBM Storage Area Network (SAN) Data Gateway and provides the procedures for installing and using it.

- *IBM Subsystem Device Driver*

This book describes how to use the IBM Subsystem Device Driver on open-systems hosts to enhance performance and availability on the ESS. The Subsystem Device Driver creates redundant paths for shared logical unit numbers. The Subsystem Device Driver permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.

This publication is not available in hardcopy. However, it is available on the compact disc that is delivered with your ESS. You can also view and copy this publication from the Web site:

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates/

- *DFSMS Advanced Copy Services, SC35-0355*

This book provides detailed information about Copy Services functions such as Peer-to-Peer Remote Copy (PPRC), Extended Remote Copy (XRC), and FlashCopy.

- *OS/390 MVS System Messages Volume 1 (ABA - ASA), SG24-5465*

This publication helps you understand and use IBM Advanced Copy Services functions on an IBM System/390. It describes two dynamic-copy functions and several point-in-time copy functions. These functions provide backup and recovery of data should a disaster occur to your data center. The dynamic-copy functions are Peer-to-Peer Remote Copy and Extended Remote Copy, and are known collectively as remote copy. FlashCopy, SNAPSHOT[®], and Concurrent Copy are the point-in-time copy functions.

- *IBM Enterprise Storage Server, SG24-5465*

This book, from the IBM International Technical Support Organization, introduces the IBM Enterprise Storage Server (ESS) and provides an understanding of its benefits. It also describes in detail the architecture, hardware and functions of the ESS.

- *Implementing the IBM Enterprise Storage Server, SG24-5420*

This book, from the IBM International Technical Support Organization, can help you install, tailor, and configure the IBM Enterprise Storage Server (ESS) in your environment.

- *IBM Enterprise Storage Server Performance Monitoring and Management, SG24-5656.*

This book, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your Enterprise Storage Server (ESS) to ensure optimum performance.

- *IBM System/360[™] and System/370[™] I/O Interface Channel to Control Unit Original Equipment Manufactures' Information, GA22-6974*

This publication provides a description of the interface between System/360, System/370, and System/390 channels and controllers that are designed to operate with this interface. This publication covers the parallel I/O interface. It does not cover the interface between the controller and the I/O device, nor does it cover the ESCON-I/O interface for the System/390.

This publication is for designers of programs and equipment that are associated with parallel I/O interfaces and for service personnel that maintain the equipment. Anyone concerned with the functional details of this interface will also find it useful.

- *Fibre Transport Services (FTS) Direct Attach, Physical and Configuration Planning*, GA22-7234

This publication provides information about fibre-optic and ESCON-trunking systems.

- *DFSMS/MVS Version 1 Remote Copy Guide and Reference*, SC35-0169

This publication provides guidelines for using remote copy functions with System/390 hosts.

- *Implementing ESS Copy Services on S/390*, SG24-6480

This publication tells you how to install, customize and configure Copy Services on IBM Enterprise Storage Server on a S/390 host system. Copy Services functions include, Peer-to-Peer Remote Copy, Extended Remote Copy, FlashCopy, and Concurrent Copy. This publication describes the functions, pre-requisites and co-requisites. The publication also describes how to implement each of the functions into your environment.

- *Implementing ESS Copy Services on UNIX and Windows NT*, SG24-5757

This publications tells you how to install, customize, and configure Copy Services on an IBM Enterprise Storage Server on the UNIX and Windows NT host system. Copy Services functions include, Peer-to-Peer Remote Copy, FlashCopy, Extended Remote Copy, and Concurrent Copy. Extended Remote Copy and Concurrent Copy are not available for UNIX and Windows NT host systems. They are only available on System/390. This publication describes the functions, pre-requisites and co-requisites. The publications also shows you how to implement each of the functions into your environment. This publications also shows how to implement these solutions in an HACMP cluster.

- *NUMA-Q ESS Integration Release Notes*, Part Number 1003-80094-01

This publication contains information about special procedures and limitations involved in running ESS with Copy Services on a NUMA-Q host system. It is not available through IBM ordering systems. Contact your IBM NUMA-Q sales representative to obtain this document

Web sites

For general information about IBM storage servers, see the following Web site:

- www.storage.ibm.com/hardsoft/products/ess/ess.htm

To view and print the IBM Enterprise Storage Server publications, see the following Web site:

- www.storage.ibm.com/hardsoft/products/ess/refinfo.htm

To get current information about the host system models, operating systems, and adapters that the ESS supports, see the following Web site:

- www.storage.ibm.com/hardsoft/products/ess/supserver.htm

For information about the IBM Subsystem Device Driver, see the following Web site:

www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates/

To attach a SAN or host system that uses an industry-standard, fibre-channel arbitrated loop (FC-AL) through the IBM 2108 SAN Data Gateway Model G07, see the following Web site:

www.storage.ibm.com/hardsoft/products/sangateway/sangateway.htm

For information about the latest updates to Copy Services components XRC, PPRC, SnapShot, Concurrent Copy, and FlashCopy for S/390, see the following Web site:

www.storage.ibm.com/software/sms/sdm/sdmtech.htm

Summary of Changes

Revision bars in the left margin indicate technical changes to the IBM Enterprise Storage Server (ESS). Minor editorial changes are not marked. This edition includes the following major changes:

- Added support for the IBM Versatile Storage Server (VSS)
- Added support for ESS Copy Services:
 - FlashCopy for open system hosts
 - Peer-to-peer remote copy (PPRC) for System/390 and open system hosts
 - Alphanumeric pager message sample

Chapter 1. Introduction

This book provides instructions for using the IBM Enterprise Storage Server (ESS). It provides an overview of the ESS, and guidelines for:

- Setting up and operating the ESS
- Use of parallel access volumes for System/390 hosts
- Migration of data to the ESS
- Problem analysis

An IBM service support representative (SSR) configures your ESS during installation. See *IBM Enterprise Storage Server Web Interface User's Guide* for information about accessing the IBM StorWatch Enterprise Storage Server Specialist (ESS Specialist) interface. You use the ESS Specialist to modify the configuration on your ESS.

See the *IBM Enterprise Storage Server Introduction and Planning Guide* for more detailed information about ESS models and on feature codes.

IBM Enterprise Storage Server overview

This section provides an overview of the IBM Enterprise Storage Server (ESS) and describes its interfaces and components.

The ESS is a member of the Seascape family of storage servers. The ESS provides integrated caching and support for redundant arrays of independent disks (RAID) for the disk drive modules (DDMs). The DDMs are attached through a serial storage architecture (SSA) interface.

The ESS uses peripheral component interconnect (PCI) architecture that is designed to provide high availability and to avoid single points of failure.

The ESS Models E10 and E20 provide 6 GB of cache. The ESS Models F10 and F20 provide 8 GB of cache; an optional 16 GB of cache; and faster PCI architecture.

The ESS provides the following features:

- RAID 5 or non-RAID disk groups
- Fast SSA drives
- Fast reduced instruction-set computer (RISC) processors
- A fault-tolerant system
- Sharing of storage between open-systems hosts
- OS/390[®] parallel input/output (I/O) interface
- Instant copy solutions
- Disaster recovery solutions

Models E20 and F20, with expansion enclosures, can provide up to 11 terabytes (TB) of storage capacity, with a choice of 9.1, 18.2, or 36.4 GB DDMs. When IBM refers to DDM capacity, TB means 1 000 000 000 000 bytes. For processor storage, real and virtual, a TB equals 2⁴⁰ or 1 099 511 627 766 bytes. Total user-accessible capacity may vary depending on the operating environment.

See Figure 1 and Figure 2 for illustrations of the ESS models and the expansion enclosure. ESS Models E10 and F10 do not support an expansion enclosure.




Figure 1. Models E10, E20, F10, and F20; front and rear views

Each ESS expansion enclosure has dual power cables and redundant power. The redundant power system allows both the storage controller and DDMs to continue normal operation when one of the power cables is inactive. Redundancy also ensures continuous data availability.




Figure 2. Expansion enclosure, front and rear views

ESS Models E10 and F10 support a maximum of 64 DDMs. The minimum configuration for a Model E10 or F10 base enclosure is 16 DDMs. ESS Models E20 and F20 support a maximum of 128 DDMs. With the expansion enclosure, Models E20 and F20 support a maximum of 384 DDMs.

Host systems supported by the ESS

The ESS provides heterogeneous data storage that you can share between workloads for open-systems hosts. The following sections contain information about:

- Open-systems hosts with small computer system interface (SCSI) adapters
- Open-systems hosts with fibre-channel adapters

Note: Because of possible interactions between host bus adapters in switched configurations, IBM recommends that you establish zones that contain a single port attached to a host bus adapter with the desired number of ports attached to ESS. You can establish the zones by using either the hard (by port) or soft (by WWN) zoning method. You can configure ports that are attached to the ESS in more than one zone, which allows multiple host bus adapters to share access to the ESS fibre-channel ports. Shared access to an ESS fibre-channel port might be from host platforms that use a combination of supported host bus adapter types and operating systems.

SCSI host systems

Each SCSI adapter has two ports. With SCSI adapters you can connect up to 128 host systems when you serially interconnect them. The ESS supports the following host systems for SCSI attachment:

- Hewlett Packard with the HP-UX operating system
- IBM RISC System/6000[®] and RISC System/6000 SP with the IBM AIX operating system
- IBM AS/400 with the IBM Operating System/400 (OS/400) Version 3 or Version 4
- Personal computer servers with Microsoft Windows NT 4.0 operating systems
- Personal computer servers with Microsoft Windows 2000 operating systems
- Sun with the Solaris operating system
- Data General with the DG UX operating system
- Personal computer servers with Novell Netware

See the following ESS Web site for details about types, models, adapters, and the operating systems that the ESS supports for SCSI-attached host systems:

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

Fibre-channel host systems

Each ESS fibre-channel adapter has one port. With fibre-channel adapters you can connect up to 127 host systems. See “Fibre-channel connection” on page 10 for information on the fibre-channel protocols that the ESS supports. The ESS supports the following host systems for fibre-channel attachment:

- IBM RISC System/6000 and RISC System/6000 SP with the IBM AIX operating system
- Personal computer servers with Microsoft Windows NT 4.0 operating systems
- Personal computer servers with Windows 2000 operating systems
- IBM NUMA-Q with DYNIX/ptx
- Hewlett Packard with the HP-UX operating system
- Personal computer servers with Novell NetWare
- Sun with the Solaris operating system

See the following ESS Web site for details about types, models, adapters, and the operating systems that the ESS supports for fibre-channel attached host systems:

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

System/390 host systems

With ESCON adapters, the ESS supports up to 32 connections for System/390 host systems, each with up to 64 logical paths. IBM attaches an ESS to a System/390 server for fibre-channel connection (FICON™) through the bridge card on the IBM 9032 Model 005 ESCON Director.

The ESS supports the following operating systems:

- OS/390
- Multiple Virtual Storage/Enterprise Storage Architecture (MVS/ESA™)
- Virtual Machine/Enterprise Storage Architecture (VM/ESA®)
- Virtual Storage Extended/Enterprise Storage Architecture (VSE/ESA™)
- Transaction Processing Facility (TPF)

The ESS also supports the following operating systems for the new IBM eserver zSeries 900 (z900 server):

- z/OS Version 1 Release 1
- OS/390 Version 2 Release 6 or higher
- VM
 - z/VM Version 3 Release 1
 - VM/ESA Version 2 Release 2
 - VM/ESA Version 2 Release 3
 - VM/ESA Version 2 Release 4
- VSE/ESA Version 2 Release 3 and higher
- Transaction Processing Facility (TPF) Version 4 Release 1

See the preventive service planning (PSP) bucket for operating system support and planning information that includes APARs and programming temporary fixes (PTFs).

For additional information about the changes you need to make to System/390 hosts to support ESS functions, see *DFSMS/MVS Version 1 Release 5 Software Support for the IBM Enterprise Storage Server*.

For information about the latest updates to Copy Services components XRC, PPRC, SnapShot, Concurrent Copy, and FlashCopy for S/390, see the following Web site:

www.storage.ibm.com/software/sms/sdm/sdmtech.htm

Web interfaces

This section describes the Web security, the ESSNet, and the Web interfaces for ESS. The Web interfaces include:

- StorWatch Enterprise Storage Server Specialist (ESS Specialist)
- StorWatch Enterprise Storage Server Web Copy Services (ESS Web Copy Services) that supports Web Interfaces to manager peer-to-peer remote copy or FlashCopy.
- StorWatch Expert, an optional software product

See your IBM Field office for information about how to order StorWatch Expert. See the *IBM Enterprise Storage Server Web Interface User's Guide* which provides detailed descriptions of the ESS Specialist and the ESS Web Copy Services interfaces and how to use them.

Web connection security

You connect to the ESS through the ESSNet.

All data that is sent between the ESS and the Web browser through the ESSNet is encrypted to avoid unauthorized modification of configuration commands. Passwords and authorization levels protect access to the interface.

You control user access by assigning levels of access and passwords.

IBM Enterprise Storage Server Network

The IBM Enterprise Storage Server Network (ESSNet) is a private network resident on an IBM workstation. It is a required feature. IBM installs the ESSNet when they install your IBM Enterprise Storage Server. The ESSNet hardware includes:

- The IBM workstation (a PC) and monitor
- An external Ethernet hub that provides cable connections from the ESSNet to the IBM Enterprise Storage Server. The hub speed is 10 megabits per second (Mbps) or 10/100 megabits per second, depending on your system.

Note: You can attach your Ethernet LAN to the external hub. You provide any hardware that is needed for this connection.

- A modem and a modem expander that allow your ESS to communicate between you and IBM

Note: The feature code for Remote Service Support includes the modem and modem expander.

The ESSNet includes the following software on the workstation:

- Windows NT 4.0 operating system
- Browser software (Microsoft Internet Explorer and Netscape Navigator) that allows access to the ESS Specialist
- The ESSNet application for installation and configuration. The ESSNet workstations includes an HTML page that provides links to all ESS connected to the ESSNet for Specialist access.

The ESSNet provides the following features:

- Support for multiple ESSs. A hub with 16 ports will support 7 ESSs.
- Connectivity between your ESS and the ESS Specialist Web interface. The ESSNet provides browser software at the appropriate level for this connection.
- Improvements in Web performance compared to previous Web browser connections.
- Faster network connectivity and elimination of network setup problems when using previous Web connections.
- Ethernet connectivity through an Ethernet hub to the ESSNet.
- An independent platform that facilitates installation and configuration of the ESS.
- Software for maintenance and configuration.

IBM installs the ESSNet when they install your first IBM Enterprise Storage Server.

Accessing ESS Specialist and ESS Web Copy Services

You access the IBM StorWatch Enterprise Storage Server Specialist (ESS Specialist) from the ESSNet. The ESSNet includes browser software for this access. You access ESS Web Copy Services from the ESS Specialist Navigation frame.

You can access ESS Specialist from your own workstation if you connect the ESSNet into your intranet by using one of the ports on the ESSNet hub. If you use your own workstation IBM recommends at least 128MB of memory in the workstation. You can access the ESS Specialist by using a browser that is fully enabled for Java(TM) 1.1.x. The ESS Specialist supports the following browsers:

- Netscape Navigator Version 4.04 or greater, which supports Windows NT and Windows 95.

Note: Version 4.04 requires Java 1.1 fixpack; later versions such as 4.5 may not require this fixpack. For better web performance, IBM recommends using Netscape Navigator Version 4.5.1 or greater.

- Microsoft Internet Explorer (MSIE) Version 4.0 or greater, which supports Windows NT and Windows 95.

Note: Version 4.0 requires Service Pack 1.

The ESSNet PC is pre-loaded with current release levels of Microsoft Internet Explorer and Netscape Navigator.

If you do not choose to attach ESSNet into your intranet, you can only access ESS Specialist from a browser on the ESSNet workstation.

You must provide another network connection to the remote ESS through ESSNet hubs for Copy Services support.

IBM updates ESS Specialist and ESS Copy Services through engineering changes (ECs).

ESS Specialist

The ESS includes the ESS Specialist. The ESS Specialist is a Web-based interface that allows you to configure the ESS.

From the Web interface, you can perform the following tasks:

- Monitor error logs
- View and change the configuration:
 - Add or delete SCSI and fibre-channel host systems
 - Configure SCSI host ports and fibre-channel host ports on the ESS
 - Define controller images for System/390
 - Define fixed-block (FB) and count-key-data (CKD) disk groups
 - Add FB and CKD volumes
 - Assign volumes to be accessible to more than one host system
 - Change volume assignments
- Change and view communication resource settings, such as electronic mail (e-mail) addresses and telephone numbers
- Authorize user access

With the ESS Specialist, you can view the following information:

- The external connection between a host system and an ESS port

- The internal connection of SCSI ports and fibre-channel ports to cluster 1 or cluster 2
- The allocation of storage space to FB and CKD volumes

IBM updates the ESS Specialist through licensed internal code (LIC) updates.

ESS Copy Services

ESS Copy Services collects information from your IBM storage servers onto a single Copy Services server. The ESS supports Copy Services through both a Web-based interface and a command-line interface.

This section describes the two interfaces for ESS Copy Services, shows an example of the connections between two Copy Services server groups, and gives you an overview of the functions you can perform using the ESS Web Copy Services panels.

Web-based interface: The ESS Web Copy Services interface initiates and manages Peer-to-Peer Remote Copy (PPRC) and FlashCopy functions for open systems and IBM System/390 hosts.

You access ESS Web Copy Services from the ESS Specialist Navigation frame.

Command-line interface: If you choose to install the command-line interface (CLI) application on your host system, you can write customized scripts using CLI commands. You can execute the CLI commands from a host system. You can use the CLI commands to invoke Copy Services tasks in an ESS Web Copy Services server.

The command-line functions (TSO or API) on IBM System/390 hosts initiate and manage Extended Remote Copy (XRC), Concurrent Copy, FlashCopy and PPRC. An API has been developed which uses the TSO command line function. For more information, see *OS/390 DFSMSdfp Advanced Services*.

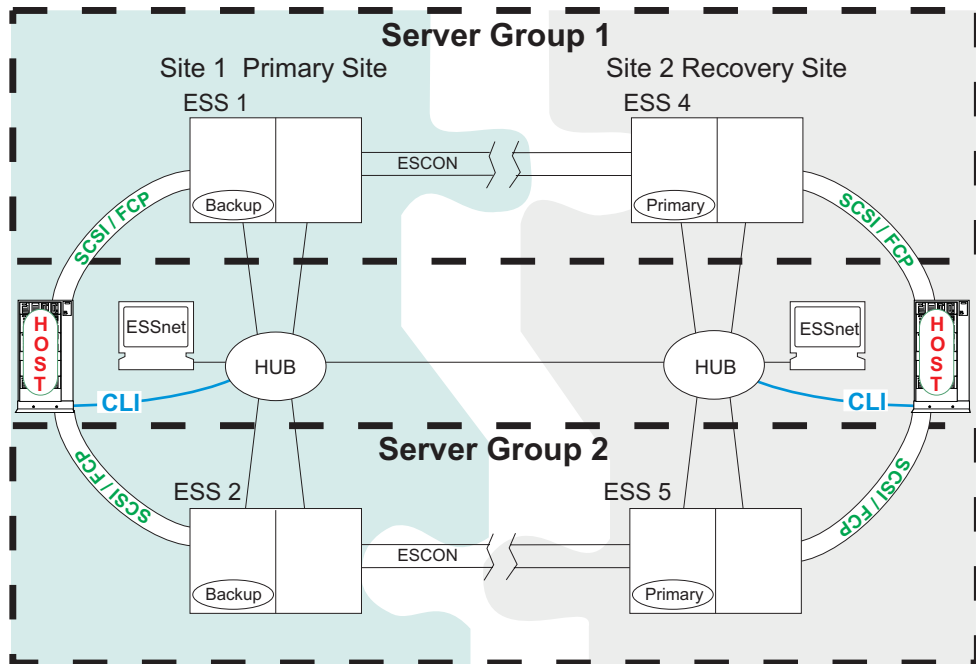
The CLI on the open-systems hosts initiates and manages PPRC and FlashCopy on ESS.

See *IBM Enterprise Storage Server Host Systems Attachment Guide* for information about using the Copy Services command-line interface from your open systems hosts to invoke Copy Services functions. See the web site at

<http://www.ibm.com/storage/hardsoft/products/ess/refinfo.htm>

for information on which open systems host platforms are supported for CLI. The command-line interface (CLI) is updated when a new LIC is provided.

Copy Services server groups: Figure 3 on page 8 shows the connections between two ESS Web Copy Services server groups. A Copy Services server group is a collection of ESS clusters that participate in Copy Services functions that include the volumes within the clusters.



CS000005

Figure 3. Connecting to ESS Web Copy Services server groups

Copy Services panels: Use the ESS Web Copy Services panels to view and define:

- **Volumes**

The volumes panel allows you to view volumes and define them as source or target volumes for PPRC or FlashCopy.

- **Storage Server**

The Storage Server panel allows you to work with logical servers as complete entities. You can build tasks to place all the volumes of a logical server within a peer-to-peer relationship with all the volumes of another logical server.

You can also build a task to remove similar groups of volumes from an existing peer-to-peer relationship.

You can determine the function of Consistency Groups and freeze and create commands.

- **Paths**

The Paths panel displays the current status of paths between an ESS and the other ESSs or switches to which it is connected. The paths panel also shows ESCON connected hosts and channel extenders.

You can use the paths panel to add or remove copy services paths or to create consistency group relationships for the LSS pairs.

- **Tasks**

The Tasks panel allows you to manage tasks you have defined. You can run, remove, save, group, or ungroup tasks. The host system command line interface function uses these predefined tasks.

- **Configuration**

The Configuration panel allows the customer to view the copy services event and timing logs, set password protection, refresh the panels and mail logs to a specified userid.

ESS Expert

The IBM StorWatch Enterprise Storage Server Expert (ESS Expert) is an optional software feature of the IBM StorWatch Expert. The StorWatch Expert can be purchased to use with the ESS. The StorWatch Expert Web interface provides storage-resource management for the IBM storage servers.

- Asset management
StorWatch Expert collects and displays asset management data on all of your ESSs.
- Capacity management
StorWatch Expert collects and displays capacity management data on all of your ESSs.
- Performance management
StorWatch Expert collects and displays performance management data on all of your ESSs. For example:
 - Number of I/O requests
 - Number of bytes transferred
 - Read and write response time
 - Cache usage statistics

The IBM StorWatch Expert allows you to schedule the information collection. With this information, you can make informed decisions about volume placement and capacity planning as well as isolate I/O performance bottlenecks.

Service interface

The ESS provides service interface ports for external connection of a service terminal.

IBM or your service provider can perform service on the ESS by using an IBM mobile service terminal (MoST) or equivalent.

Remote services support

The ESS service interface also provides remote service support with call-home capability for directed maintenance for service personnel.

You provide an analog telephone line to enable this support. The service interface provides an RS232 connection through a modem switch and modem to the analog telephone line.

You need to order a modem and modem switch. The first ESS you order requires this equipment. The modem and modem switch support up to seven ESSs. The cable length from the ESS to the modem switch should be a maximum of 15 m (50 ft.).

The ESS provides the following service functions:

- Continuous self-monitoring initiates a call (call home) to service personnel if a failure has occurred. Because service personnel who respond to the call know about the failing component, the service personnel can reduce the repair time.
- Error and problem logs are available that service personnel can access remotely to analyze potential failures.
- Remote support allows the ESS to correct many types of problems. When the ESS reports a problem, service personnel can often correct the problem from a remote location.

Configuration

IBM logically configures your ESS during installation using the work sheets that you completed from the *IBM Enterprise Storage Server Configuration Planner*. After IBM installs the ESS, you can perform additional configuration by using the ESS Web interfaces. You can also change the remote service functions.

Fibre-channel connection

The ESS provides fibre-channel connection to open-systems hosts that it supports. Fibre-channel interconnection architecture provides a variety of communication protocols on the ESS. The units that are interconnected are referred to as *nodes*. Each node has one or more ports.

An ESS is a node in a fibre-channel network. Each port on an ESS fibre-channel host adapter is a fibre-channel port. A host is also a node in a fibre-channel network. Each port attaches to a serial-transmission medium that provides duplex communication with the node at the other end of the medium.

With fibre-channel technology, the distance between an ESS and attached host systems increases from the 25 m (75 ft) maximum imposed by SCSI connections up to 500 m (1500 ft). The ESS can be up to 11 km (6.8 mi) from your host when you connect to it with an IBM SAN fibre-channel switch. The ESS can also be the same distance from your host when you connect to it with an IBM fibre-channel storage hub.

ESS architecture supports three basic interconnection topologies:

- Point-to-point
- Switched fabric
- Arbitrated loop

Point-to-point topology

The point-to-point topology, also known as direct connect, allows you to interconnect ports directly. Figure 4 shows an illustration of a point-to-point topology.



S008944L

Legend

- 1** is the host system.
- 2** is the ESS.


Figure 4. Point-to-point topology

Switched fabric topology

The switched fabric topology provides the underlying structure that allows you to interconnect multiple nodes. You can use a fabric that provides the necessary switching functions to support communication between multiple nodes. You can implement a fabric by using available vendor products that the ESS supports. For a list of these vendor products, see the ESS Web site at:

- www.storage.ibm.com/hardsoft/products/ess/supserver.htm

The ESS supports the switched-fabric topology with point-to-point protocol. The ESS fibre-channel adapter should be configured to operate in point-to-point mode when you connect it to point-to-point topologies or fabric topologies. See Figure 5.



Legend


- 1** is the host system.
- 2** is the ESS.
- 3** is a switch.

Figure 5. Switched fabric topology

Arbitrated loop

Fibre-channel arbitrated loop (FC-AL) is a ring topology that allows you to interconnect a set of nodes. The maximum number of ports you can have for a fibre-channel arbitrated loop is 127. See Figure 6 on page 12.

The ESS supports FC-AL as a private loop. It does not support the fabric-switching functions in FC-AL.



S008943N

Legend

1 is the host system.

2 is the ESS.

Figure 6. Arbitrated loop topology

Introducing the IBM Subsystem Device Driver

The IBM Subsystem Device Driver resides in the host server with the native disk device driver for the ESS. It uses redundant connections between the host server and disk storage in an ESS to provide enhanced performance and data availability. Figure 7 on page 13 is an example of the type of configuration that IBM Subsystem Device Driver supports. These connections comprise many different components through which data flows during input and output processes. Redundancy and the ability to switch between these components provides many different paths for the data to travel. In the event of failure in one input-output path, it automatically switches to another input-output path. This automatic switching in the event of failure is called *failover*.

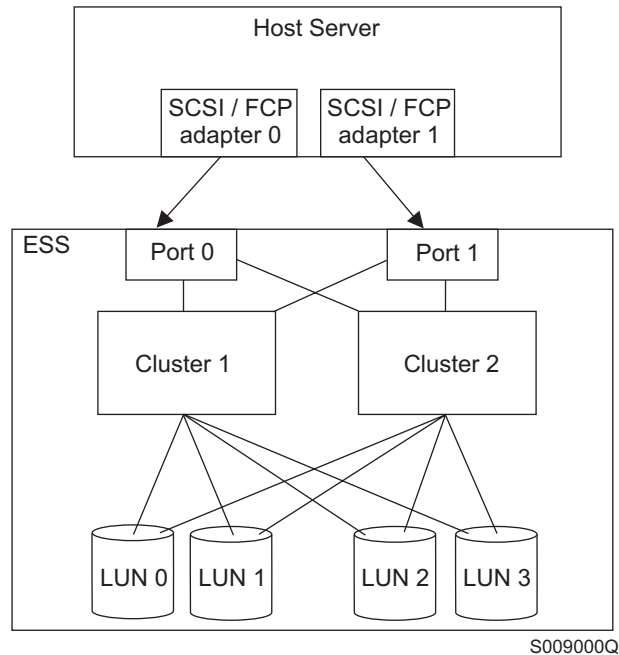


Figure 7. Multipath connections between a host server and ESS LUNs

The IBM Subsystem Device Driver provides the following functions:

- Enhanced data availability
- Automatic path failover and recovery to an alternate path
- Dynamic load balancing of multiple paths
- Path selection policies for the AIX operating system
- Concurrent download of licensed internal code

In most cases, host servers are configured with multiple host adapters with SCSI or fibre-channel connections to an ESS that, in turn, provides internal component redundancy. With dual clusters and multiple host interface adapters, the ESS provides more flexibility in the number of input-output (I/O) paths that are available.

When there is a failure, the IBM Subsystem Device Driver reroutes I/O operations from the failed path to the remaining paths. This function eliminates the following connections as single points of failure: a bus adapter on the host server, an external SCSI cable, a fiber-connection cable, or a host interface adapter on the ESS.

In addition, multipath load balancing of data flow prevents a single path from becoming overloaded with I/O operations.

The IBM Subsystem Device Driver path recovery algorithms now test *Dead* paths that are not manually offline every 30 minutes.

Concurrent download of licensed internal code is the capability to download and install licensed internal code on an ESS while applications continue to run. During the time when new licensed internal code is being installed in an ESS, the upper interface adapters inside the ESS may not respond to host I/O requests for approximately 30 seconds. The IBM Subsystem Device Drivers makes this transparent to the host through its path selection and retry algorithms.

Note: For updated and additional information not included in this publication, see the README file on the IBM Subsystem Device Driver compact disk recordable (CDR) or visit the Subsystem Device Driver Web site at:
www.ibm.com/storage/support/techsup/swtechsup.nsf/support/sddupdates/

Chapter 2. Setting up and operating the Enterprise Storage Server

This chapter tells you how to set up and operate the IBM Enterprise Storage Server. It includes information about the following tasks:

- Using the operator panel to power the ESS on and off
- Attaching your ESS and the ESS Network to your local area network

Using the control panel

This section describes the control panel switches and indicators, and provides the procedures for turning the power for the ESS on and off. It also shows the control panel for the expansion enclosures.

Figure 8 on page 16 shows the IBM Enterprise Storage Server control panel for the ESS Models. The control panel has the following switches and indicators:

- The Local Power switch, which controls direct current (dc) power in the ESS
- The Power Complete light-emitting diodes (LEDs) for Line Cord 1 and Line Cord 2:
 - LEDs are on solid when the power sequence is complete with no errors
 - LEDs flash slowly to indicate a power fault
 - LEDs flash rapidly to indicate that the on- or off-sequence is in progress
 - LEDs are off when the power is off

Note: A power sequence on the first enclosure causes the green power LEDs to flash fast, that is, 350 milliseconds (ms) on and 350 ms off. For the other enclosure, the power sequence causes the green power LED to turn on for 400 ms, then off for 400 ms.

- Messages LEDs for Cluster 1 and Cluster 2:
 - LEDs are on when the licensed internal code (LIC) or the cluster detect a hardware fault. The LEDs also indicate where the service action should begin.
 - LEDs flash rapidly to indicate that a cluster power sequence is in progress.
- The Unit Emergency power switch, which you use to turn the power off to the ESS in an emergency. Do *not* use it for normal power-off procedures.




Figure 8. Control panel switches and indicators for Models E10, E20, F10, and F20

Turning the power on and off

This section provides the procedures for turning the power to the ESS Models E10, E20, F10, and F20 on and off.

Turning on the power

Note: The Unit Emergency switch and the remote switch must be in the *on* position.

To turn on the power to the ESS:

1. Set the Local Power switch to the *on* position. See Figure 8 for switch locations and indicator locations:
 - The Power LEDs for Line Cord 1 and Line Cord 2 flash alternately for 3 – 10 seconds when the power switch is pressed.
 - If the power LED flashes slowly, you need to have service personnel check for a power fault.
 - If the power LEDs are flashing rapidly, the power-on sequence is in progress.

Turning off the power

To turn off the power to the ESS:

1. Set the Local Power switch to the *off* position.

Note: IBM recommends that you also put the remote switch in the *off* position.

2. When the Power Complete LEDs are off for Line Cord 1 and Line Cord 2, the power is *off*. (This can take a minute or longer.)

The ESS should power off within 3 - 5 minutes after you set the Local Power switch to the *off* position. The line cord LEDs will be off when the system power off has completed. If the line cord LEDs do not go off, call your service representative for repair.

Control panel for the expansion enclosure

Figure 9 shows the IBM Enterprise Storage Server control panel for the expansion enclosure.

Note: The Local Power switch on the Models E10 and E20 control direct current (dc) power to the expansion enclosures.

The control panel has the following switches and indicators:

- The Power Complete LEDs for Line Cord 1 and Line Cord 2:
 - LEDs are on solid when the power sequence is complete with no errors
 - LEDs flash slowly to indicate a power fault
 - LEDs flash rapidly to indicate that the on- or off-sequence is in progress
 - LEDs are off when the power is off
- The Unit Emergency power switch, which you use to turn the power off to the ESS in an emergency. Do *not* use it for normal power-off procedures.

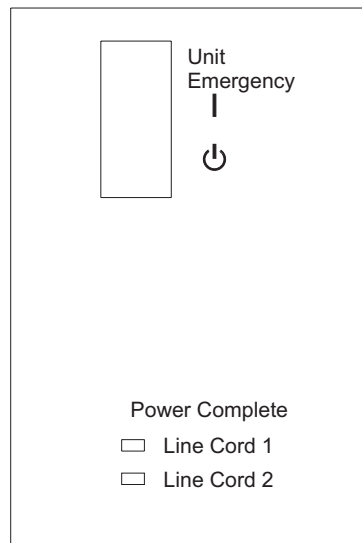


Figure 9. Control panel switches and indicators for an expansion enclosure

Attaching the ESS Network to your local area network

While IBM does not require you to have a local area network (LAN) attached to the IBM Enterprise Storage Server, IBM recommends that you do so. A LAN allows you to receive problem notification messages. IBM or a service provider connects your LAN to the ESSNet. (IBM attaches the ESS to the ESSNet during installation.)

Note: You provide any cables or adapters that you might need to attach your LAN to the ESSNet.

LAN attachment to the ESSNet allows you to manage user access and to monitor errors on the ESS.

Note: Problems on the LAN will not affect the ability of the ESS to read and write data. However, some service functions will be affected, and the ESS could generate and send messages as electronic mail (e-mail) to indicate a problem on the LAN. It could also generate and send messages as a Simple Network

Management Protocol (SNMP) trap. See “Analyzing e-mail messages” on page 41 for more information about e-mail messages.

Follow these guidelines to prepare your LAN for attachment to the ESSNet:

1. In the domain name server, assign a Transmission Control Protocol/Internet Protocol (TCP/IP) address and a hostname to each of the clusters in the ESS.
2. If you allow e-mail, enable support on the mail server for the TCP/IP addresses that you assigned to the ESS.
3. You need one 10BASE-T cable, already connected to a 10BASE-T LAN¹ that will reach to the ESSNet hub.

If you are attaching the ESSNet to a non-10BASE-T LAN, you need a cable and an adapter, as appropriate for the LAN type.

4. IBM or your service provider will make the connections to the ESSNet.
5. Attach the other end of the cables to your LAN hardware.
6. The service personnel can check the LAN connections and run notification tests.

The notification tests verify the addresses that you entered in the ESS Specialist Web pages. These tests also verify that the trap destination receives the Simple Network Management Protocol (SNMP) traps.

1. These guidelines assume that the LAN is an Ethernet.

Chapter 3. Using parallel access volumes for System/390 hosts

This chapter provides the following guidelines for using parallel access volumes (PAVs):

- An overview of PAVs
- PAV requirements
- An overview of the input/output configuration program (IOCP)
- Guidelines for defining the logical control units
- An example of an input/output (I/O) configuration program (IOCP)

To use the ESS advanced feature for System/390 PAVs, you need to define the ESS in your I/O configuration using IOCP.

Overview of parallel access volumes

The ESS supports concurrent or parallel data transfer operations to or from the same volume from the same system or system image for System/390 hosts.

A System/390 with PAV software support allows multiple users and jobs to simultaneously access a logical volume. Read operations can be satisfied simultaneously as well as write operations to different domains.²

The parallel I/O capability in the ESS system allows higher I/O rates to a logical volume. As a result, the ESS has increased overall throughput and has reduced response time.

PAVs can improve the performance of large volumes. You will get better performance with one base and two aliases³ on a 3390 Model 9 than from three 3390 Model 3 volumes with no PAV support. This would also reduce storage management costs that are associated with maintaining large numbers of volumes.

Parallel access volumes requirements

The following is a list of the requirements for using PAV:

- The configuration you create on the ESS (2105) must match the configuration you create using IOCP.
- The control unit (CU) image numbers that exist on the ESS depend on what logical subsystems (LSSs) you configured:
 - The CU image addresses are 0 - 15. For IOCP the CU addresses are hexadecimal, 0 - F.
 - LSS addresses are 0 - 15.

The configuration might only have a subset of these addresses, and the subset might not be contiguous.

2. The domain of an I/O operation is the specified extents to which the I/O operation applies.

3. An alias is an alternate path to the base device. For example, a 3380 or a 3390 with one alias has only one device to write to, but two paths for getting there.

You need to know what CU images you have created before you can update the I/O configuration. Your host system uses the updated I/O information in order to know what CU images are on the ESS.

Overview of the input/output configuration program

The output generated by the input/output configuration program (IOCP) contains the I/O definition for the channel subsystem of a System/390 host. It is not used as input to the I/O definition for OS/390, VM, or VSE. Therefore, you do not need to distinguish between bases and aliases in the input to IOCP. However, doing so provides an accurate representation of the I/O configuration. See “IOCP example”.

The OS/390 operating system requires that the input/output definition file (IODF) that is created by the hardware configuration definition (HCD), specifies base devices and alias devices. OS/390 dynamically discovers which base and alias devices are associated with each other.

The VM and VSE operating systems dynamically discover which devices are bases and aliases, and which are associated with each other. VM supports PAV for guest usage.

See these publications for detailed information:

- *IBM ES/9000 Input/Output Configuration Program User's Guide*, GC38-0401
This publication provides information about using the input/output configuration program.
- *IBM OS/390 Hardware Configuration Definition User's Guide*, SC28-1848
This publication provides detailed information about the input/output definition file (IODF). It also provides details about configuring PAVs. OS/390 uses the IODF.

IOCP example

The following IOCP example defines Enterprise Systems Connection (ESCON) and fibre-connection (FICON) channel path IDs that you have attached to an ESCON director. This example uses the ESS machine number, 2105.


```

*****
*
*   DEFINE ESCON CHPIDS
*
*****
      CHPID PATH=(70),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(71),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(72),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(73),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(74),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(75),PARTITION=(HOST1),SWITCH=50,TYPE=CNC
      CHPID PATH=(FC),PARTITION=(HOST1),SWITCH=50,TYPE=FCV
      CHPID PATH=(FD),PARTITION=(HOST1),SWITCH=50,TYPE=FCV
*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 0
*
*****
      CNTLUNIT CUNUMBR=7000,PATH=(70,71,72,73),UNITADD=((00,256)),
      LINK=(24,2D,34,3D),CUADD=0,UNIT=2105
*****
*   DEFINE 3390-9 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 0
*
      16 BASE ADDRESS, 3 ALIASES PER BASE
*****
      IODEVICE ADDRESS=(9000,016),CUNUMBR=(7000),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(90D0,048),CUNUMBR=(7000),STADET=Y,UNIT=3390A
*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 1
*
*****
      CNTLUNIT CUNUMBR=7001,PATH=(70,71,72,73),UNITADD=((00,128)),
      LINK=(24,2D,34,3D),CUADD=1,UNIT=2105
*****
*   DEFINE 3390-9 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 1
*
      16 BASE ADDRESS, 3 ALIASES PER BASE
*****
      IODEVICE ADDRESS=(9100,016),CUNUMBR=(7001),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(9150,048),CUNUMBR=(7001),STADET=Y,UNIT=3390A
*****

```

Figure 10. IOCP example, Part 1 of 2

```

*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 2   *
*                                             *
*****
      CNTLUNIT CUNUMBR=7002,PATH=(70,71,72,73),UNITADD=((00,256)), *
      LINK=(24,2D,34,3D),CUADD=2,UNIT=2105
*****
*   DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 2 *
*   64 BASE ADDRESS, 3 ALIASES PER BASE   *
*                                             *
*****
      IODEVICE ADDRESS=(9200,064),CUNUMBR=(7002),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(9240,192),CUNUMBR=(7002),STADET=Y,UNIT=3390A
*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 3   *
*                                             *
*****
      CNTLUNIT CUNUMBR=7003,PATH=(70,71,72,73),UNITADD=((00,256)), *
      LINK=(24,2D,34,3D),CUADD=3,UNIT=2105
*****
*   DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 3 *
*   64 BASE ADDRESS, 3 ALIASES PER BASE   *
*                                             *
*****
      IODEVICE ADDRESS=(9300,064),CUNUMBR=(7003),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(9340,192),CUNUMBR=(7003),STADET=Y,UNIT=3390A
*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 4   *
*                                             *
*****
      CNTLUNIT CUNUMBR=7004,PATH=(74,75,FC,FD),UNITADD=((00,256)), *
      LINK=(24,2D,34,3D),CUADD=4,UNIT=2105
*****
*   DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 4 *
*   32 BASE ADDRESS, 7 ALIASES PER BASE   *
*                                             *
*****
      IODEVICE ADDRESS=(9400,032),CUNUMBR=(7004),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(9420,224),CUNUMBR=(7004),STADET=Y,UNIT=3390A
*****
*   DEFINE 2105-E20 LOGICAL CONTROL UNIT 5   *
*                                             *
*****
      CNTLUNIT CUNUMBR=7005,PATH=(74,75,FC,FD),UNITADD=((00,256)), *
      LINK=(24,2D,34,3D),CUADD=5,UNIT=2105
*****
*   DEFINE 3390-3 BASE AND ALIASES ADDRESS ON LOGICAL CONTROL UNIT 5 *
*   32 BASE ADDRESS, 3 ALIASES PER BASE   *
*                                             *
*****
      IODEVICE ADDRESS=(9500,032),CUNUMBR=(7005),STADET=Y,UNIT=3390B
      IODEVICE ADDRESS=(95A0,096),CUNUMBR=(7005),STADET=Y,UNIT=3390A

```

Figure 11. IOCP example, Part 2 of 2

Guidelines for defining the logical control unit

The following is a list of the information you need in order to define the logical control unit (LCU):

- Which devices (arrays) are 3380 format, and which are 3390 format
- Which devices have PAVs, and which do not
- How many devices you define to a given LCU

Notes:

1. An LCU is equivalent to a logical subsystem (LSS). Like an LSS, an LCU can have a maximum of 256 logical devices or volumes.
2. Every device has a single unit address (in range 00 - FF).
3. An LCU can manage a maximum of 256 devices and identifies a device based on the unit address.

Update the I/O configuration with the appropriate devices, using the following guidelines:

- There are three device types possible for 3380: 3380, 3380B, and 3380A.
- There are three device types possible for 3390: 3390, 3390B, and 3390A.
- Any non-PAV 3380 is a 3380, regardless of model.
- Any non-PAV 3390 is a 3390, regardless of model.
- Any 3380 PAV base is a 3380B.
- Any 3390 PAV base is a 3390B.
- Any alias of a 3380B, regardless of model, is a 3380A
- Any alias of a 3390B, regardless of model, is a 3390A.
- Only define the actual base and alias devices to IOCP or HCD to avoid unnecessary definition of processor subchannels.
- Only define the unit address range to IOCP or HCD as necessary for the installed devices in an LCU (64, 128, or 256). This allows you to put more LCUs and devices on a channel path.

IBM recommends that you document the configuration you define for the ESS. The ESS configuration and the IOCP input must match, with the exception that IOCP does not care about the various 3380 and 3390 models.

Note: The CU device is a 2105 for any ESS device type. In HCD you can use 3990 as a control unit type for 3380 and 3390 devices. However, you must use 2105 as the control unit type for 3380B, 3380A, 3390B, and 3390A devices.

Chapter 4. Migrating data to the IBM Enterprise Storage Server

This chapter provides the following guidelines for migrating data to the ESS:

- Data migration for UNIX, Windows NT, Windows 2000 and OS/400 operating systems
- Data migration for System/390 operating systems
- IBM data migration and other services

Migrating data from UNIX, Windows NT, Windows 2000 and OS/400

This section provides guidelines for migrating data in UNIX, Windows NT, Windows 2000, and OS/400 operating system environments.

The following commands below are guidelines for migrating data to the ESS.

- Volume management software
Volume management software provides specific tools for wholesale movement of data. "AIX logical volume manager" on page 29 describes some methods and provides examples.
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 with **cpio -p**
The **cpio** command is a standard UNIX command for archiving and copying data. The *-p* (pass) option allows the operating system to copy data between file systems without creating an intermediate archive.
The **cpio** command reads a list of files to copy from its standard input. You typically use it with the UNIX **find** command to copy complete directory structures.
- AIX **backup** and **restore** commands
You commonly use these commands on AIX systems to archive and restore data. However, they do not support a direct disk-to-disk copy. They require an intermediate device such as a tape drive or spare disk drive to store the archive that is created by the **backup** command.
- The **dump** and **restore** commands
These commands are similar in function to the AIX **backup** and **restore** commands; you mostly find them on other forms of UNIX. They, too, require an intermediate device.
- Other commands
Other UNIX commands provide archival functions, such as the **tar** command, that you can use to transport data. Again, these commands require that you create an intermediate archive, usually on a tape drive or a spare disk drive.

Your system administrator should select the method that is the best compromise between efficiency and impact on the users of the system:

- Management software provides simple robust methods that you can generally use during production without disturbing users.

- The AIX logical volume manager (LVM) provides methods that you can use at any time without disrupting user access to the data. You might notice a small performance degradation, but this is better than having to shut down databases or require users to log off the system.
- Methods that use 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.

Replacing existing storage

When the ESS replaces existing storage, you must partition it so that its virtual disks are similar in configuration to the drives that it is replacing. New configurations should be large enough to accommodate the existing data.

Reformatting existing disk drive modules

To provide integrity checking of data and to allow different operating systems to use the ESS, IBM formats the disks to a 524-byte sector. The 524-byte sector replaces the traditional 512-byte sector of fixed-block architecture DDMs. The 524-byte sector format requires that you reformat all disks that you are migrating to the ESS. Because reformatting will erase any data that is already contained on the drives, you must back up the data to some form of temporary storage.

You can use these drives for either direct migration of data or as temporary storage while the existing drives are moved. If this is not possible, use a removable media device, such as a tape drive, to temporarily store the data while you reformat the drives.

Most fixed-block disk architectures use a fixed-byte sector of 512 bytes. This includes most UNIX systems, including AIX. When used in an ESS subsystem, the format of a DDM has a fixed-byte sector of 524 bytes. The 524-byte sector format enables the ESS subsystem to connect to a wide range of host systems and share data between them.

The IBM AS/400 host systems use eight bytes at the start of the sector. UNIX host systems do not use these eight bytes when you attach them to the ESS. The data portion of the sector remains at 512 bytes for all host systems. A 2-byte sequence number and a 2-byte LRC increase the size of the sector to 524 bytes.

The sequence number is a modulo 64K value of the logical block address (LBA) of this sector. It is an extra method of ensuring that you are accessing the correct block.

The LRC, generated by the SCSI host adapter, is calculated on the 520 data bytes and header bytes. The ESS uses LRC for error checking as the data progresses from the host, through the ESS controller, into the device adapter, to the array. See “Adapters for UNIX systems” on page 27 for a detailed description of data flow through the subsystem. The ESS also uses LRC as an error checking mechanism as it reads data from the array and passes data to the host adapter. The ESS never transfers the sequence number or the LRC to the host system.

Host system views of the IBM Enterprise Storage Server

The ESS supports connection to open system hosts through SCSI adapters and fibre-channel adapters. A variety of UNIX-based systems and AS/400 systems support the ESS as peripheral devices. The following sections discuss how the host systems view the ESS.

Adapters for UNIX systems

For UNIX-based systems, the ESS emulates multiple SCSI DDMs. The host system accesses the virtual drives of the ESS as if they were generic SCSI DDMs. The AIX operating system contains entries in its object distribution manager database to identify the ESS. However, the AIX operating system accesses the ESS through its generic SCSI DDM. The ESS appears to be a standard physical volume or *hdisk* to AIX. The ESS appears similarly to Solaris and HP-UX systems.

When you use Ultra-or wide-SCSI adapters in your host systems, a total of 16 SCSI IDs per interface are available on the ESS. The host system SCSI IDs are known as initiators; the ESS SCSI IDs are the targets.

If only one host system connects to an ESS SCSI port, the ESS can assign up to 15 unique target IDs. If the maximum of four host systems connect to the ESS, the ESS can only assign up to 12 unique SCSI target IDs.

You can configure an ESS to appear as 64 LUNs per SCSI interface. LUN sizes can be defined from 0 to 245.5 GB in size and in 100 MB increments with the 7/28/00 ESS code release. If the ESS has 36 GB DDM, then the additional LUN sizes of 112, 128, 160, 192, and 224 GB are also available. If you do not specify a value, an 8 GB LUN is the default.

Transferring data for AS/400 systems

The ESS emulates a 9337 subsystem when you attach it to an AS/400 host. Because an AS/400 host requires 9337 subsystems to have groups of four to eight drives, you must also configure the ESS to appear as four to eight drives. The 4.5 GB drives emulate the 9337-580, and the 9.1 GB drives emulate the 9337-590.

The AS/400 expects to see a separate device address for each drive in the subsystem. The ESS meets this requirement by reporting unique addresses for each virtual drive defined to the AS/400 host.

Because the AS/400 host does not support 7133 drawers, you will have no existing 7133 drawers to reformat. If you have an existing disk subsystem that you are replacing with an ESS, you need to migrate the data to the ESS. Use your existing host utilities for the data migration.

You can select from several methods to migrate data to the ESS:

- You can use the logical add and remove functions.
- You can use save methods 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: The AS/400 allows software mirroring on a unprotected 9337 model and prevents software mirroring on a protected 9337 model. From a ESS logical configuration view point all AS/400 volumes are RAID-5 and are protected. When you create the AS/400 volume from Web Copy Services Specialist you can create it as a protected or unprotected volume. This allows you to see the protected and unprotected models. Unprotected models can be software mirrored.

For more information about these procedures, see the AS/400 backup and recovery guides for the release of the operating system that you have on your host system. IBM also offers data migration services.

Transferring data for UNIX systems

You can migrate data to an ESS by using standard host utilities; the host sees the ESS as one or more generic DDMs. Select the method that best meets your criteria for:

- The amount of data that you want to migrate
- The amount of time available
- The availability of tape devices or spare disks for temporary storage
- The format of the data itself

The following sections describe the data migration methods.

Using logical volume management software

AIX and HP-UX 10.xx ship with logical volume management (LVM) software as part of the base operating system. LVM provides complete control over all disks and file systems that exist on an AIX system. HP-UX has similar volume management software. “AIX logical volume manager” on page 29 describes AIX methods and examples in more detail.

Sun Microsystems has a basic volume management product called Solstice, which is available for the Solaris systems. You can also purchase the Veritas Volume Manager (VxVM) and Veritas File System (VxFS) as optional products for Solaris.

All of these volume managers provide tools and utilities to move data:

- At a logical volume level
- At a physical volume (DDM) level
- By selective use of the volume manager’s mirroring features

Using direct copy

If the data that you are migrating resides as individual files on UNIX file systems and no volume management software is available, use a utility. A utility that supports a direct copy feature, such as the **cpio** command with the *-p* (pass) option, provides the next easiest method of moving data.

The **cpio** command is available on all of the UNIX operating systems that support the ESS. It is easy to use; **cpio** reads a list of files to copy from its standard input.

The easiest way to produce the list of files is to use the UNIX **find** command. Pipe its standard output to the standard input of the **cpio** command. The following output shows a typical example of using the **cpio** command to move data.

```
# mount /dev/lv00 /mnt
# cd /data
# find . -print | cpio -pdmuv /mnt
.
.
.
# umount /mnt
# umount /data
# mount /dev/lv00 /data
```

For the example:

1. Assume that you have made a file system on the `/dev/lv00` logical volume. AIX LVM uses this file system to view part or all of the virtual disks that the ESS has made available to the system.
2. Mount the logical volume on a temporary mount point, in this case `/mnt`.

3. Then change directories to the directory at the top of the file system you want to move (`cd/data`).
4. Using the **find** command produces a list of file names, which a pipe (`|`) passes to the **cpio** command.
5. Once finished, unmount both file systems, and mount the new file system over the original mount point directory.

Backup and restore

In some cases, the only method available to transfer data is to back it up to tape, then restore it to the new disk. This method is obviously slower as tape devices are essentially slow devices. However, if you are removing disks before you install the ESS, the only way to move the data is with a tape device.

There are a number of different utilities found on UNIX systems with which to archive data. The **cpio** command that is mentioned above can also create and read archives on tape devices. Using the `-o` (output) option creates archives, and the `-i` (input) option reads and extracts data from archives.

AIX provides a pair of commands, **backup** and **restore**. The **backup** command has two different methods of creating archives, either by file name or by the i-node. Use the **restore** command for reading any archive that is created by the **backup** command. Solaris and HP-UX provide the **dump** and **restore** commands, which provide backup and restoration by the i-node.

The **tar** command is available on all UNIX systems that support the ESS, and is another form of creating and extracting tape archives. Use the `c` option to create archives. Use the `x` option to extract files from archives.

You might not be able to use the volume management methods to migrate the following types of data:

- Databases that use raw file systems
- Logical volumes or methods other than a UNIX file system

You probably will be unable to use the volume management methods if the database uses volume serial numbers in its licensing code or validity checking. If databases use licensing methods or validity checking, you might 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 can take the form of a standard database backup and restore if it does not have any specific tools for movement of data.

AIX logical volume manager

The AIX LVM provides useful tools and utilities for migrating 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. This section describes some of the methods that you can use to migrate data and shows some example commands.

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.

Copying a complete logical volume: The AIX LVM provides the **cplv** command for copying logical volumes within volume groups or to different volume groups.

You can use this command to create a new logical volume while running the command, or you can overwrite an existing logical volume. The following shows two examples of how to use the **cplv** command:

```
# cplv -v datavg -y newlv oldlv
# cplv -e existinglv oldlv
```

In the first example, the **cplv** command copies data from the existing logical volume `oldlv`. It creates a new logical volume that it calls `newlv` (`-y`) in the volume group `datavg` (`-v`). If you omit the `-v` option, the volume group to which the existing logical volume belongs receives the new logical volume. When the **cplv** command creates a new logical volume, it creates the new volume with exactly the same characteristics as the existing logical volume.

In the second example, the **cplv** command copies data from existing logical volume `oldlv` to existing logical volume `existinglv` (`-e`). When you use the `-e` option, it overwrites the existing target-logical volume with the data from the source-logical volume. When you use the `-e` option, it maintains the characteristics of the existing target-logical volume. Take care when you use this option.

The **cplv** command is a good method for copying or migrating a single logical volume. Sometimes, however, you might need to migrate all the data from a physical volume.

Using the migratepv command: The AIX LVM provides the **migratepv** command to migrate a complete physical volume of data from one volume to another. It also provides options to migrate only portions (at the logical volume level) of data from one physical volume to another. For example:

```
# migratepv hdisk1 hdisk2
# migratepv -l data1lv hdisk3 hdisk9
```

In the example, all data migrates from `hdisk1` to `hdisk2`. The **migratepv** command updates all LVM references. From the time that the command completes, the LVM no longer refers to `hdisk1` to access data that was previously stored there. As the data is physically moved, the target physical volume must have enough spare physical volumes to accommodate the source physical volumes data. After this command completes, you can remove the source-physical volume from the volume group.

To avoid disturbing users, use the **migratepv** command while the system is active.

The **migratepv** command works by:

1. Creating a mirror of the logical volumes that you are moving
2. Synchronizing the logical volumes
3. Removing the original logical volume

You can use the **migratepv** command to move data from one physical volume to another physical volume, within the same volume group. The usage parameters for this command are:

```
migratepv [-i] [-l LVname] SourcePV DestinationPV...
```

Note: You can specify more than one destination physical volume.

First identify the source disk from which you want to migrate the data. Then identify the target disk to which you want to migrate the data. You can only migrate to disks that are already in the rootvg volume group. To get a list of disks that are already in the rootvg volume group, run the **lsvg -p rootvg** command. For example:

```
# lsvg -p rootvg
rootvg:
PV_NAME  PV STATE  TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk0   active    515        116       57..00..00..00..59
hdisk1   active    515        515       00..00..00..00..00
```

Now, determine the space that is currently in use on the disk you want to migrate. This will be the total PPs value minus the free PPs value for the desired disk. In the preceding example, select hdisk0, which is using (515 - 116) PPs or 399 physical partitions.

Next, find a disk or disks that have the available space. In this case, hdisk1 has 515 free physical partitions, which is more than the required space of 399 physical partitions.

The only situation that is specific to rootvg is if the desired disk contains the boot image. Generally, this is the logical volume called hd5. This logical volume has a type of "boot", as you see from a partial list of **lsvg -l rootvg**:

```
LV NAME   TYPE  Ps  PPs  PVs  LV STATE   MOUNT POINT
hd5       boot  1   1   1   closed/syncd  N/A
```

To determine if the boot image is on the disk you want to migrate, run the **lslv -l** command. For example:

```
# lslv -l hd5
hd5:N/A
PV          COPIES          IN BAND          DISTRIBUTION
hdisk0     001:000:000    100%            001:000:000:000:000
```

In this case, the boot image resides on the disk you want to migrate. You need to first move the boot image with the **migratepv -l** command.

Using mirroring to migrate data: One function of AIX LVM is to provide RAID level 1, data mirroring, in software. You can use mirroring to create copies or move data. This is similar to using the **cplv** and **migratepv** commands.

The following sections describe two ways of using mirroring to duplicate and migrate data.

Using mirroring to migrate data, example 1: The **mklvcopy** command creates a second copy of the data within the logical volume. The **mklvcopy** command is all that you need to start mirroring data in a logical volume. You can use it on an

active system without disturbing users. You might notice some performance degradation while the new data copy synchronizes with the existing copy.

Use the **splitlvcopy** command to break a mirrored logical volume into separate logical volumes that use copies of the data. This command creates new logical volumes; it leaves one intact with its original name. You can use the **splitlvcopy** command on an active logical volume. However, be aware that it can cause corruption if the ESS is actively updating data while the split operation is running.

You can use the **rmlvcopy** command to delete a copy of a mirrored logical volume. Removing the copy frees the space for use by other logical volumes in the volume group. You can also move data from one physical volume to another by creating a copy of the data with **mklvcopy**. Once the synchronization of the new copy has completed, you can remove the old copy. Use this method in order to avoid disrupting operations.

The following is an example of the **mklvcopy** command. It shows how to create a mirror copy in a logical volume:

```
# mklvcopy -e m -s y -k data1v 2 hdisk3 hdisk7
.
.
.
# splitlvcopy -y splitlv data1v 1
```

The options specify:

(-e m) To use minimum inter-disk allocation policy

(-s y) To strictly allocate mirror copies on separate physical volumes

(-k) To synchronize new copies immediately

data1v The name of the logical volume where you want to start a mirroring operation

2 The number of copies that you want to make of the data (this can be a maximum of 3)

hdisk3 **and** *hdisk7*

The physical volumes on which the logical volume will reside

hdisk3 The physical volume that is already holding the data of *data1v*

hdisk7 The physical volume that will hold the mirror copy and where you want to move the data

The **splitlvcopy** command creates a new logical volume, *splitlv* with the *-y* option. The source-logical volume is *data1v*. The number 1 specifies how many copies are to remain in the logical volume.

The source-logical volume, *data1v*, will still exist as it did before, residing on *hdisk3*. The new *splitlv* will reside on *hdisk7*.

Using mirroring to migrate data, example 2: The second example is similar to example 1. It uses the **mklvcopy** command to start mirroring the *infx1v* logical volume.

```
# mklvcopy -e m -s y -k infxlv 2 hdisk4 hdisk10
.
.
# rmlvcopy infxlv 1 hdisk4
```

In this example, `hdisk4` is where the data for `infxlv` already exists; `hdisk10` is where the mirror copy resides. It is also where we want to ultimately move the data.

The **rmlvcopy** command specifies to remove a copy from `infxlv`. The options specify to leave 1 remaining copy of the data and remove the existing copy from `hdisk4`. The remaining data is on `hdisk10`. Effectively, the data has moved from `hdisk4` to `hdisk10`.

Using the mirrorvg command: With AIX 4.2.1 you can use **mirrorvg** and **unmirrorvg** to simplify mirroring or unmirroring the entire contents of a volume group. The commands detect if the entity to be mirrored or unmirrored is `rootvg`.

Note: The commands give slightly different completion messages that are based on the type of volume group.

The **mirrorvg** command mirrors all of the `rootvg` to only one disk - `hdiskx`. For example:

```
mirrorvg rootvg hdiskx
```

Note: Use the **mklvcopy** command if you want to mirror logical volumes to different `hdisks`.

If you mirrored the `rootvg`, you need to break the mirror prior to doing any updates. By using the **ls -l rootvg** command, you can determine what portions of the `rootvg` are mirrored. On an AIX 4.2.1 or higher system, you can unmirror the volume group with the command:

```
unmirrorvg rootvg hdiskx
```

where `hdiskx` is the disk to remove.

This breaks the copies. It is possible that the primary copy was moved during operation of the system, so be sure you know which copy you want to remove. To remove the boot image from `hdisk1` if you are unmirroring `therootvg`, enter:

```
mkboot -c -d /dev/hdisk1
```

To replace a failed disk in a mirrored volume group, enter:

```
unmirrorvg workvg hdisk7
reducevg workvg hdisk7
rmdev -l hdisk7 -d
```

Replace the drive. Then enter:

```
cfgmgr
extendvg workvg hdisk7
mirrorvg workvg
```

cfgmgr adds the new disk to the system in the first available device name.

For more detailed information about data migration for AIX, and current information about PTF and APARS, call 1-800-426-4FAX.

Data migration affect on data availability

Most methods of data migration affect everyday operation of a computer system. When data is moved, it must be in a 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.

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 you must consider when selecting a method to use for migrating data.

Volume management methods for data migration

Use the **cplv** command to make entire copies of logical volumes. It can operate without disrupting access to data that is contained within the logical volume that it is copying. If the data is actively being updated, however, the resulting copy might not look anything like the logical volume it started out as. Some updates make it to the copy; others might not. Inconsistency is a drawback of using the **cplv** command on active logical volumes. You should close any logical volumes before you copy them using the **cplv** command. To close a logical volume, either unmount the file system that is built on it, or shutdown any database which has an open raw logical volume. After running the **cplv** command, update the system configuration file, `/etc/filesystems`, to include the relevant configuration data on the new logical volume and file system. Run the file-system integrity checker **fsck** to ensure data consistency within the new logical volume.

The **migratepv** command is one of the better methods for moving data without disrupting users. You can run the **migratepv** command on an active system. It first creates a mirror of each logical volume contained within the physical volume, and then it synchronizes both copies. Once the synchronization is complete, it removes the original copy and leaves the new copy active and available to users. Users might notice some performance degradation due to the automatic nature of the creation and the synchronization of the mirror. This process writes each physical partition and locks it from access by any other process. This can slow down access to the data, but it ensures data integrity.

The **mklvcopy**, **splitlvcopy** method is ideal for creating a copy of logical volumes. You can run with both copies, one in a production environment, and one in a test environment. The **mklvcopy** command ensures data integrity by creating a mirror copy of the data and synchronizing it automatically. You should not run the **splitlvcopy** command on an active logical volume, nor you should you run the **cplv** command on an active logical volume. If processes are updating the data while the split is taking place, the consistency of the data on both copies cannot be guaranteed. After you run the **splitlvcopy** command, update the configuration file of the file system, `/etc/filesystems`. Include the relevant configuration data

regarding the new logical volume and file system. Then run the file system integrity checker **fsck** to ensure data consistency within the new logical volume.

The **mklvcopy**, **rmlvcopy** method is similar to the method that the **migratepv** command uses. It creates and removes mirrors automatically, ensuring data integrity. This method is ideal for migrating active logical volumes when a slight performance loss is acceptable.

Direct copy method for data migration

This method suffers from the same drawback as when using the **cplv** and **splitlvcopy** commands on logical volumes. You can use it on active file systems, but data consistency between the original and new copies cannot be guaranteed. Use the **cpio -p** command only on file systems that are in a quiescent state to guarantee consistency between the copy and the original. This in turn mandates that some disruption of service to users must occur.

Backup and restore methods for data migration

All backup and restore methods require that the file systems or logical volumes that you are backing up be in a quiescent state. Typically, you run backups after hours when there is minimal use of the system and you can shutdown databases.

Some databases require that you export the data to move it between file systems or logical volumes. You must then import the data onto the new file system or logical volume. You might have to perform some reconfiguration of the database so that it will point to the new data locations.

Data migration from System/390 operating systems

This section includes the following information for System/390 operating systems:

- A list of publications that provide information about data migration
- Guidelines for migrating data to the ESS
- Information about IBM services for data migration

Publications about data migration

The following publications can guide you through data migration:

- *3390 Installation and Migration for DFSMS*, GG66-3155
- *DASD Migration in DFSMS*, GG24-3515
- *IBM 3390 Direct Access Storage Migration Guide*, GG24-3373
- *MVS Guide for IBM 3390 Direct Access Storage*, GC28-1105
- *Using IBM 3390 Direct Access Storage in a VM Environment*, GC26-4575

Guidelines for migrating data to an ESS

The following sections provide guidelines for migrating data for MVS, VM, and VSE operating systems.

Migrating data in an MVS environment

Use DFDSS to migrate data to the ESS for an MVS environment. While DFDSS is by far the simplest way to move most data, you can also use the following functions in MVS environments:

- IDCAMS EXPORT/IMPORT (VSAM)
- IDCAMS REPRO (VSAM, SAM, BDAM)
- IEBCOPY (partitioned data set (PDS), including load module libraries, and PDSE)

- ICEGENER (SAM), part of DFSORT™
- IEBGENER (SAM)
- Specialized data base utilities for Customer Information Control System (CICS), DB2®[®], or Internet Management Specification (IMS)

Another tool that offers non-disruptive data migration in the MVS environment is the Amdahl Transparent Data Migration Facility (TDMF). This tool is available from IBM as a data migration service.

Migrating data using Extended Remote Copy (XRC)

You can use XRC if the data is on disk drive modules (DDM) if you have XRC on the ESS, and it is supported in an S/390 environment only. XRC is also useful in the occurrence of disaster recovery. You can use XRC in the following environments:

- Remote mirroring of volumes
- An asynchronous copy process
- Supporting large distances

Migrating data in a VM environment

You can simplify data migration for a VM environment by using DFSMS/VM®[®], which is a no-charge feature of VM/ESA. DFSMS/VM lets you manage and control the use of VM disk space.

DFSMS/VM includes a data mover, an automated move process, and an interactive user interface for migrating data. A description of these features follows:

- DASD Dump Restore (DDR) is a service utility shipped with VM. You can use it 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 that you can use to move minidisks from any device types supported by VM, to any other type.
- COPYFILE is a CMS command that you can use to copy files or minidisks between devices with the same or different track modes.
- PTAPE is a Control Program (CP) command that you can use to dump spool files to tape and to load files from tape to disk.

Migrating data in a VSE environment

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 following backup and restore dialogs:

- Export and import VSAM files
- Backup and restore VSAM files
- Backup and restore ICCF libraries
- Backup and restore the system history file
- Backup and restore the system residence library
- Create a loadable tape with the system residence library and system, history file ready to restore

You can also use the following commands and functions to migrate data in a VSE environment:

- VSE **fastcopy** command to move volumes and files between devices with identical track formats
- VSE **ditto** command to copy files

- VSE **power** command to transfer the SPOOL queue from one device to another
- VSE/VSAM to move any VSAM dataset by using either REPRO or EXPORT/IMPORT functions

Other vendors also provide utilities for moving data from one device to another.

Data migration and other services

IBM offers the following services for large systems:

- Customer migration services
- Data storage environment analysis (DSEA)
- Data migration tools and services for MVS

In the U.S., you can arrange to have IBM provide these services. Call 1-800-IBM-4YOU and request information about storage service offerings.

Chapter 5. Problem analysis

This chapter provides guidelines for:

- Analyzing problems that occur during operation
- Analyzing service information messages for System/390 hosts
- Analyzing Simple Network Management Protocol (SNMP) trap information
- Analyzing electronic mail (e-mail) messages
- Analyzing pager messages
- Obtaining service

Analyzing problems that occur during operation

If the ESS encounters an error that requires action, it illuminates one of the message lights on the expansion enclosure. See “Using the control panel” on page 15 for a description of the control panel.

The ESS also sends an alert as an e-mail or an SNMP trap. The method depends on how you define alert handling during configuration of TCP/IP on the ESS Specialist Web interface. The ESS also places an automatic call for service if you enable this function.

If your warranty still covers the ESS, or if you have a service maintenance agreement, contact your authorized service representative. You will need to describe the error message, error code, or problem that you observed. Error messages help identify what service action the ESS requires, and help service personnel provide quick and effective service to you.

Note: Your warranty agreement or service maintenance agreement might be affected if you perform your own maintenance.

Service information messages for System/390 hosts

The ESS generates service information messages (SIMs) for all attached System/390 hosts and sends them to your host console. However, SIM generation is not the primary notification for service. The primary notification is to your ESS Specialist Problem Log. IBM uses the problem number in the problem log to track service. The ESS assigns a problem number and a matching SIM ID to a problem. Check the problem log for the problem number that matches the ID of the SIM that you want to track.

You determine how often the ESS sends the SIMs and the SIM severity type. You provided the SIM settings to IBM when you completed the Communication Resources work sheet before IBM installed your ESS. See *IBM Enterprise Storage Server Introduction and Planning Guide* for more information about the Communication Resources work sheet.

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

- Direct access storage device (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.

You select the number of times that the ESS sends messages to your console. The ESS can send a SIM 1 to 9 times, or not at all (0). You also select the severity level of the SIMs. The ESS does not send any SIMs that are below the severity level that you select. Severity levels are:

- | | |
|-------------------|---|
| 1 acute | An irrecoverable error with possible loss of data. This severity level only applies to DASD SIMs. |
| 2 serious | An irrecoverable error or a data check with loss of access to data. |
| 3 moderate | A system path or paths are not operational and performance might be degraded. This severity level does not apply to media SIMs. |
| 4 service | A recoverable error, equipment checks, or data checks. You can defer repair. |

If you did not fill in the SIM information in the Communication Resources work sheet, give your IBM service representative the following information. IBM will use the information to configure the ESS for SIM generation:

1. The SIM severity reporting level for DASD SIMs.
1 to 4
2. The number of times the ESS is to send a DASD SIM.
0 to 9
3. The SIM severity reporting level for media (data check) SIMs.
2 or 4
4. The number of times the ESS is to send a media SIM.
0 to 9
5. The SIM severity reporting level for a storage server SIM.
2, 3, or 4
6. The number of times the ESS is to send a storage server SIM.
0 to 9

Analyzing SNMP trap information

The ESS sends Simple Network Management Protocol (SNMP) traps to the addresses specified by you in the Communication Resources work sheet. You filled out this work sheet during the planning process.

The ESS generates SNMP traps and supports a read-only management information base (MIB) to allow monitoring by the network.

The ESS generates both generic and product-specific SNMP traps. The ESS supports the usual generic traps, such as cold start and warm start. Product-specific traps provide information about problems detected by the ESS that require action by you or the service provider.

Each product specific trap provides the following information:

- Problem ID: this is the identifier assigned by the ESS for the particular problem that caused the trap.
- Presenting time: includes the time and date of the most recent problem report.

- Exception symptom code (ESC): the Exception symptom value for the problem id.
- Description: the error condition description for the problem id.
- User Action: recommended actions for the user.

There are two types of Product-specific trap:

1. Problems detected by the ESS that require action by you or the service provider.
2. With Copy Services, state changed events detected by the ESS that might require action by you.

Note: Problem id is a default value for all Copy Services state changed events, and it is not in the Management Information Base (MIB).

Product-specific traps contain the identifier assigned by the ESS for the particular 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 find additional details about the problem by using a MIB browser. Perform the following steps to view problem details:

1. Point the MIB browser to the product that created the trap.
2. Look within the MIB *ibm2100.mib* objects, *ssProblem* or *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, then look at other MIB variables within *ssProblem*.

Analyzing e-mail messages

The ESS sends e-mail messages to the addresses that you specified in the Communication Resources work sheet. You filled out this work sheet during the planning process.

The ESS generates e-mail messages in the following categories:

- Information
- Errors

Informational messages

Examples of *informational messages* are notifications that:

- IBM has installed a new level of licensed internal code (LIC)
- The service provider has run the customer-notification diagnostic test

This test verifies that e-mail messages are being received by those who should receive them. You need to attach your LAN to the ESSNet to receive e-mail messages from the ESS.

Error messages

The ESS sends *error messages* when it detects a situation that requires customer action.

The error messages typically contain the following fields:

Product manufacturer ID and date

The ID of the ESS and the date it was manufactured.

ESS location

The ESS location is entered by the installer during the initial installation of the product.

Product machine type and model number

Assigned by IBM. For example a 2105 Model F20

Product-serial number

Assigned by IBM.

Customer voice phone number

The phone number for customer voice contact that you enter in the TCP/IP Web page of the ESS Specialist Web interface.

LIC level of local storage server

The level of the licensed internal code of your primary ESS.

LIC level of remote storage server

The level of the licensed internal code of your secondary, or backup ESS.

Report time and date stamp

The time that this report was generated.

Problem ID

The problem ID assigned to the this problem by the storage facility. This problem ID can be used by the service provider to access detailed problem information.

ESC

A detailed error code used by the service provider.

SRN

A detailed error code used by the service provider.

Problem status

See "Status" on page 44 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 that reported the failure (1 or 2).

Failing resource

This coded resource name is used during the repair process by a service provider.

Failure occurred

Date and time when the failure first occurred.

Last occurrence

Date and time of the last noticed occurrence.

Failure count

The number of times this failure occurred.

Presentation interval

The time between successive e-mail copies of this problem.

Remaining presentations

The number of additional times this e-mail 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 following fields are the most useful to you in identifying DDM failures:

- Description
- Failing cluster
- Reporting cluster
- Failing resource

- Last occurrence

Status

A problem can exist in various states. Each state represents the action that IBM has taken or will take in the resolution of a problem. The problem states are defined as follows:

State Description

Pending

The initial problem state. The ESS reports the problem SNMP by either trap or e-mail. You need to attach your LAN to the ESSNet in order to receive e-mail messages or SNMP traps from the ESS.

Received

The ESS uses this state to indicate that a host has received the problem notification, either by 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 is complete. 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.

Analyzing pager messages, numeric and alphanumeric

If you enabled your pager to receive problem information from the ESS, you will receive notifications to numeric and alphanumeric pagers. Numeric pagers receive the information in the following format:

The numeric pager test message is 01234-56789

The normal pager message is 111-aaaaaa-222-bbbbbbb-333-c-ddd, where:

aaaaaa are the numeric characters from the product-serial number

bbbbbbb is the customer telephone number

c is the cluster number (1 or 2)

ddd is the problem ID

Alphanumeric pagers receive the information from the ESS as text. An example of a typical alphanumeric pager message follows:

```
COMPANY RAS6 SERIAL MMFC3A5 PHONE 6-7025 CLUSTER 1 PID 10 SRN 1244 RESOURCE
ssa01
```

Following is the same Problem as sent to Retain:

```
# Begin of call home record.....= BEGIN OF FILE (2676)
```



```

# Product Manufacturer ID and Date.....= MMMJJYYYY
# Business/Company Name.....= RAS6
# Product Machine type and model number.= 2105F20 # Product serial
number.....= MMFC3A5
# Remote service call back number.....= 6-6404
# Call back expander port ID.....= 7
# Call back user ID.....= support
# Call back password.....= ras6
# Outside line prefix.....=
# Customer voice phone number.....= 6-7025
# Customer offshift voice phone number..=
# LIC level of local complex.....= SC00724
# LIC level of remote complex.....= SC00724
# Reporting Cluster Number.....= 1
# Record type.....= 1
# Report time/date stamp.....= 08/15/00 12:09:38 PDT
# Host system type(s).....= AS/400_B, IBM_RS/6000, escon
# Level_CPSS information.....= shark/R6.CPSS/engr/build/gurs0802
# Possible FRUs to replace:
#
# Engineering      FRU                Likely FRU Location
# FRU Name         Name                to fix Description
#-----
rsDDM0101         SSA Disk Drive Modul  33%  R1-U1-W4-D1
rsDDM0102         SSA Disk Drive Modul  33%  R1-U1-W4-D2
rs8pkplnr01      Frame (8-Pk)         33%  R1-U1-W4-P1
#-----
# Additional Information
#

```

```

# Engineering FRU Name ... = rsDDM0101
# Part Number ..... = 09L4295
# Engineering FRU Name ... = rsDDM0102
# Part Number ..... = 09L4295
#
# Primary Failing Unit ... = 2105-E20      13-FCA14
# Reporting Unit ..... = 2105-E20      13-FCA14
#
# Problem ID ..... = 10
# ESC ..... = 1244
# SRN ..... = 1244
# Problem Status ..... = PENDING
# Description ..... = FRU FAILURE
# Additional Message ..... = NONE
# Failing Cluster ..... = 1
# Reporting Cluster ..... = 1
# Failing Resource ..... = ssa01
# First Occurrence ..... = Thu Sep  9 08:46:27 1999
# Last Occurrence ..... = Thu Sep  9 08:46:27 1999
# Failure Count ..... = 1
# Re-presentation Interval = 28800
# Remaining Presentations. = 3
# Isolation Procedure .... = 3010
# Failure Actions ..... = D521 "SERVICE ACTION REQUIRED" # Probable
Cause ..... = D12C "FRU FAILURE"
# Failure Cause ..... = D12C "FRU FAILURE"
# User Actions ..... = FFFF "NONE"
# End of record flag ..... = END OF FILE

```

You enable your pager for this service when you fill out the Communication Resources work sheet. See *IBM Enterprise Storage Server Introduction and Planning Guide* for more information about the work sheet.

Obtaining service

The ESS is capable of remote error notification and remote support for those machines that are under warranty or a maintenance agreement. An IBM service support representative (SSR) configures your ESS for remote service during installation. For those machines not under warranty or a maintenance agreement, call:

1-800-IBM-SERV

For more information about remote service, see “Remote services support” on page 9.

Technical support

If you have a problem configuring or using the ESS, call your technical support provider or contact IBM.

Warranty support or maintenance

If you have a problem that requires a hardware service technician to repair the ESS, contact IBM at 1-800-IBM-SERV (426-7378).

Have the following information available when you call:

IBM Customer Number _____ (9 digit number)

IBM Product Number _____

Model _____

IBM Product Serial Number _____

Date of Purchase _____

Your Phone Number _____

Your Fax Number _____

Write your service telephone numbers here for later reference if you do not allow remote support.

_____.

_____.

_____.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publications. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Trademarks

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

AIX
AS/400
DB2
DFSMS/MVS
DFSMS/VM
DFSORT
Enterprise Storage Server
ES/390
ESCON
FICON
FlashCopy
IBM
MVS/ESA
Netfinity
NUMA-Q
Operating System/400
OS/390
OS/400
RISC System/6000
RS/6000
S/370
S/390
Seascape
StorWatch
System/360
System/390
System/370
Versatile Storage Server
VM/ESA
VM/VSE

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

Pentium is a trademark of Intel Corporation in the United States, other countries, or both.

Microsoft Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean government Ministry of Communication (MOC) statement

Please note that this device has been approved for business purpose with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for a nonbusiness purpose one.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

IBM agreement for licensed internal code

Read Before Using
IMPORTANT

YOU ACCEPT THE TERMS OF THIS IBM LICENSE AGREEMENT FOR MACHINE CODE BY YOUR USE OF THE HARDWARE PRODUCT OR MACHINE CODE. PLEASE READ THE AGREEMENT CONTAINED IN THIS BOOK BEFORE USING THE HARDWARE PRODUCT. SEE "IBM agreement for licensed internal code".

You accept the terms of this Agreement⁴ by your initial use of a machine that contains IBM Licensed Internal Code (called "Code"). These terms apply to Code used by certain machines IBM or your reseller specifies (called "Specific Machines"). International Business Machines Corporation or one of its subsidiaries ("IBM") owns copyrights in Code or has the right to license Code. IBM or a third party owns all copies of Code, including all copies made from them.

If you are the rightful possessor of a Specific Machine, IBM grants you a license to use the Code (or any replacement IBM provides) on, or in conjunction with, only the Specific Machine for which the Code is provided. IBM licenses the Code to only one rightful possessor at a time.

Under each license, IBM authorizes you to do only the following:

1. execute the Code to enable the Specific Machine to function according to its Official Published Specifications (called "Specifications");
2. make a backup or archival copy of the Code (unless IBM makes one available for your use), provided you reproduce the copyright notice and any other legend of ownership on the copy. You may use the copy only to replace the original, when necessary; and
3. execute and display the Code as necessary to maintain the Specific Machine.

You agree to acquire any replacement for, or additional copy of, Code directly from IBM in accordance with IBM's standard policies and practices. You also agree to use that Code under these terms.

You may transfer possession of the Code to another party only with the transfer of the Specific Machine. If you do so, you must 1) destroy all your copies of the Code that were not provided by IBM, 2) either give the other party all your IBM-provided copies of the Code or destroy them, and 3) notify the other party of

4. Form Z125-4144

these terms. IBM licenses the other party when it accepts these terms. These terms apply to all Code you acquire from any source.

Your license terminates when you no longer rightfully possess the Specific Machine.

Actions you must not take

You agree to use the Code only as authorized above. You must not do, for example, any of the following:

1. Otherwise copy, display, transfer, adapt, modify, or distribute the Code (electronically or otherwise), except as IBM may authorize in the Specific Machine's Specifications or in writing to you;
2. Reverse assemble, reverse compile, or otherwise translate the Code unless expressly permitted by applicable law without the possibility of contractual waiver;
3. Sublicense or assign the license for the Code; or
4. Lease the Code or any copy of it.

Statement of limited warranty

International Business Machines Corporation
Armonk, New York, 10504

The warranties provided by IBM in this Statement of Limited Warranty⁵ apply only to Machines you originally purchase for your use, and not for resale, from IBM or your reseller. The term "Machine" means an IBM machine, its features, conversions, upgrades, elements, or accessories, or any combination of them.

Unless IBM specifies otherwise, the following warranties apply only in the country where you acquire the Machine. If you have any questions, contact IBM or your reseller.

Machine: IBM Enterprise Storage Server 2105 Model E10 and E20

Warranty Period: Three Years.*

**Contact your place of purchase for warranty service information.*

Production status

Each Machine is manufactured from new parts, or new and used parts. In some cases, the Machine may not be new and may have been previously installed. Regardless of the Machine's production status, IBM's warranty terms apply.

IBM warranty for machines

IBM warrants that each Machine 1) is free from defects in materials and workmanship and 2) conforms to IBM's Official Published Specifications. The warranty period for a Machine is a specified, fixed period commencing on its Date of Installation. The date on your receipt is the Date of Installation, unless IBM or your reseller informs you otherwise.

During the warranty period IBM or your reseller, if authorized by IBM, will provide warranty service under the type of service designated for the Machine and will manage and install engineering changes that apply to the Machine.

5. Form Z125-4753

For IBM or your reseller to provide warranty service for a feature, conversion, or upgrade, IBM or your reseller may require that the Machine on which it is installed be 1) for certain Machines, the designated, serial-numbered Machine and 2) at an engineering-change level compatible with the feature, conversion, or upgrade. Many of these transactions involve the removal of parts and their return to IBM. You represent that all removed parts are genuine and unaltered. A part that replaces a removed part will assume the warranty service status of the replaced part.

If a Machine does not function as warranted during the warranty period, IBM or your reseller will repair it or replace it with one that is at least functionally equivalent, without charge. The replacement may not be new, but will be in good working order. If IBM or your reseller is unable to repair or replace the Machine, you may return it to your place of purchase and your money will be refunded.

If you transfer a Machine to another user, warranty service is available to that user for the remainder of the warranty period. You should give your proof of purchase and this Statement to that user. However, for Machines which have a life-time warranty, this warranty is not transferable.

Warranty service

To obtain warranty service for the Machine, you should contact your reseller or call IBM. In the United States, call IBM at **1-800-IBM-SERV (426-7378)**. In Canada, call IBM at **1-800-465-6666**. You might be required to present proof of purchase.

IBM or your reseller will provide certain types of repair and exchange service, either at your location or at IBM's or your reseller's service center, to restore a Machine to good working order.

When a type of service involves the exchange of a Machine or part, the item IBM or your reseller replaces becomes its property and the replacement becomes yours. You represent that all removed items are genuine and unaltered. The replacement may not be new, but will be in good working order and at least functionally equivalent to the item replaced. The replacement assumes the warranty service status of the replaced item. Before IBM or your reseller exchanges a Machine or part, you agree to remove all features, parts, options, alterations, and attachments not under warranty service. You also agree to ensure that the Machine is free of any legal obligations or restrictions that prevent its exchange.

You agree to:

1. Obtain authorization from the owner to have IBM or your reseller service a Machine that you do not own; and
2. Where applicable, before service is provided:
 - a. Follow the problem determination, problem analysis, and service request procedures that IBM or your reseller provide,
 - b. Secure all programs, data, and funds contained in a Machine, and
 - c. Inform IBM or your reseller of changes in a Machine's location.

IBM is responsible for loss of, or damage to, your Machine while it is 1) in IBM's possession or 2) in transit in those cases where IBM is responsible for the transportation charges.

Extent of warranty

IBM does not warrant uninterrupted or error-free operation of a Machine.

The warranties might be voided by misuse, accident, modification, unsuitable physical or operating environment, improper maintenance by you, removal or alteration of Machine or parts identification labels, or failure caused by a product for which IBM is not responsible.

THESE WARRANTIES REPLACE ALL OTHER WARRANTIES OR CONDITIONS, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THESE WARRANTIES GIVE YOU SPECIFIC LEGAL RIGHTS AND YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF EXPRESS OR IMPLIED WARRANTIES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU. IN THAT EVENT SUCH WARRANTIES ARE LIMITED IN DURATION TO THE WARRANTY PERIOD. NO WARRANTIES APPLY AFTER THAT PERIOD.

Limitation of liability

Circumstances might arise where, because of a default on IBM's part or other liability, you are entitled to recover damages from IBM. In each such instance, regardless of the basis on which you are entitled to claim damages from IBM (including fundamental breach, negligence, misrepresentation, or other contract or tort claim), IBM is liable only for:

1. Damages for bodily injury (including death) and damage to real property and tangible personal property; and
2. The amount of any other actual direct damages or loss, up to the greater of U.S. \$100,000 or the charges (if recurring, 12 months' charges apply) for the Machine that is the subject of the claim.

UNDER NO CIRCUMSTANCES IS IBM LIABLE FOR ANY OF THE FOLLOWING: 1) THIRD-PARTY CLAIMS AGAINST YOU FOR LOSSES OR DAMAGES (OTHER THAN THOSE UNDER THE FIRST ITEM LISTED ABOVE); 2) LOSS OF, OR DAMAGE TO, YOUR RECORDS OR DATA; OR 3) SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), EVEN IF IBM OR YOUR RESELLER IS INFORMED OF THEIR POSSIBILITY. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

Glossary

This glossary includes terms for the IBM Enterprise Storage Server (ESS) and other Seascope solution products.

This glossary includes selected terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172–1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (SIO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

The following cross-reference is used in this glossary:

See This refers the reader to (a) a related term, (b) a term that is the expanded form of an abbreviation or acronym, or (c) a synonym or more preferred term.

A

AAS. See *Advanced Administrative Systems*.

access. (1) To obtain the use of a computer resource. (2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

Advanced Administrative Systems (AAS). A system used by IBM for order entry for machines and external cables either by a regular order or by an MES.

alert. (1) A message or log that a storage facility generates as the result of error event collection and analysis. An alert indicates that you need to perform some service action. (2)

allegiance. In the ESA/390 architecture, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. Refer to implicit allegiance, contingent allegiance, and reserved allegiance.

allocated storage. On an IBM storage server, the space allocated to volumes, but not yet assigned. See *assigned storage*.

American National Standard Code for Information Interchange (ASCII). The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A)

American National Standards Institute (ANSI). An organization consisting of producers, consumers, and general interest groups, that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

ANSI. See *American National Standards Institute*.

APAR. See *authorized program analysis report*.

arbitrated loop. For fibre-channel connection, arbitrated loop is a topology that allows you to interconnect a set of nodes.

array. An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices.

ASCII. See *American National Standard Code for Information Interchange (ASCII)*.

assigned storage. On an IBM storage server, the space allocated to a volume and assigned to a port.

authorized program analysis report (APAR). A report of a problem caused by a suspected defect in a current, unaltered release of a program.

availability. The degree to which a system or resource is capable of performing its normal function. Also see *data availability*.

B

bay. Physical space on an IBM Enterprise Storage server. A bay contains SCSI and ESCON host adapter cards. See *service boundary*.

bit. Either of the digits 0 or 1 when used in the binary numeration system. (T)

block. A group of consecutive bytes.

byte. A group of eight adjacent binary digits that represent one EBCDIC character.

C

cache. A buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

cache fast write. A form of the fast-write operation in which the subsystem writes the data directly to cache where it is available for later destaging.

cache memory. Typically refers to volatile memory that a subsystem uses to improve access times to instructions or data. The cache memory is typically smaller and faster than the primary memory or storage medium. A storage subsystem uses the cache memory to improve the performance of accesses to data that ultimately resides on the storage devices in the storage facility.

cascading. The connecting of network controllers to each other in a succession of levels, to concentrate many more lines than a single level permits. In high-availability cluster multiprocessing (HACMP), pertaining to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails but relinquishes the resource to the primary node upon reintegration of the primary node into the cluster

CCR. See *channel-command retries*.

CCW. See *channel command word*.

CEC. See *computer-electronic complex*.

channel. In the ESA/390 architecture, the part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of controllers.

channel-command retries (CCR). In the ESA/390 architecture, the protocol used between a channel and a controller that allows the controller to request that the channel reissue the current command.

channel command word (CCW). In the ESA/390 architecture, a data structure that is used to specify an I/O operation to the channel subsystem.

channel path. In the ESA/390 architecture, the interconnection between a channel and its associated controllers.

channel subsystem. In the ESA/390 architecture, the part of a host computer that manages I/O communication between the program and any attached controllers.

channel-subsystem image. In the ESA/390 architecture, the logical functions that a system requires to perform the function of a channel subsystem. With EMIF, one channel subsystem image exists in the channel subsystem for each LPAR. Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities.

CKD. See *count key data*.

CLI. See *command-line interface*.

cluster. (1) A partition in a storage server that is capable of performing all function of a storage server. When there are multiple clusters in a storage server, any remaining clusters in the configuration can take over the processing of any failing clusters. (2) On an AIX platform, a cluster is a group of nodes within a complex.

cluster processor complex (CPC). The unit within a cluster that provides the management function for the storage server. It consists of cluster processors, cluster memory, and related logic.

command-line interface. The command-line interface on the open-systems hosts that initiates and manages PPRC and FlashCopy on the ESS.

compact disc, read-only memory. High-capacity read-only memory in the form of an optically read compact disc.

compression. (1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks. (2) Any encoding to reduce the number of bits used to represent a given message or record.

computer-electronic complex (CEC). The set of hardware facilities associated with a host computer.

concurrent copy. Concurrent Copy allows a program such as DFDSS to make a backup of a data set while maintaining continuous read and write access to the data. The data is frozen at the point in time that the commands were issued to the controller.

concurrent download of licensed internal code. Capability to download and install licensed internal code on an ESS while applications continue to run

concurrent maintenance. Service that is performed on a unit while it is operational.

concurrent media maintenance. Service performed on a disk drive module (DDM) without losing access to the data.

configure. To define the logical and physical configuration of the input/output (I/O) subsystem through the user interface provided for this function on the storage facility.

consistent copy. A copy of a data entity (for example a logical volume) that contains the contents of the entire data entity at a single instant in time.

console. A user interface to a server, such as a personal computer.

contingent allegiance. In the ESA/390 architecture, a relationship that is created in a controller between a device and a channel path when unit-check status is accepted by the channel. The allegiance causes the controller to guarantee access; the controller does not present the busy status to the device. This enables the controller to retrieve sense data that is associated with the unit-check status on the channel path with which the allegiance is associated.

controller. In the ESA/390 architecture, a storage controller with ESCON or OEMI interfaces. See *cluster*.

controller image. In the ESA/390 architecture, a logical subsystem that you access with an ESCON I/O interface. One or more control-unit images exist in each controller. Each image appears to be an independent controller, but all images share a common set of hardware facilities.

control unit (CU). In System/390 architecture, a storage control unit with ESCON or OEMI interfaces. Also referred to as a controller.

Copy Services server group. A collection of ESS clusters that participate in Copy Services functions that include the volumes within the clusters.

count field. The first field of a CKD record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD). In the ESA/390 architecture, a logical device that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a track that can contain one or more records, each consisting of a count field, a key field (optional), and a data field (optional).

CPC. See *cluster processor complex*.

CRC. See *cyclic redundancy check*.

CU. See *control unit*.

customer console. See *console*.

CUT. See *Universaile Tempes du Coordonnaire*.

cyclic redundancy check (CRC). A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder. A unit of storage on a CKD device. A cylinder has a fixed number of tracks.

D

DA. See *device adapter* and *SSA adapter*.

daisy chain. A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

DASD. See *direct access storage device* and *disk drive module*.

DASD-fast write. A function of a storage controller that allows caching of active write data without exposure of data loss by journaling of the active write data in nonvolatile storage.

data availability. The degree to which data is available when needed. For better data availability when you attach multiple hosts that share the same data storage, configure the data paths so that data transfer rates are balanced among the hosts.

data compression. A technique or algorithm that you use to encode data such that you can store the encoded result in less space than the original data. This algorithm allows you to recover the original data from the encoded result through a reverse technique or reverse algorithm. See *compression*.

Data Facility Storage Management Subsystem. An operating environment that helps automate and centralize the management of storage. To manage storage, SMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field. The third (optional) field of a CKD record. You determine the field length by the data length that is specified in the count field. The data field contains data that the program writes.

data record. A subsystem stores data records on a track by following the track-description record. The subsystem numbers the data records consecutively, starting with 1. A track can store a maximum of 255 data records. Each data record consists of a count field, a key field (optional), and a data field (optional).

data sharing. The ability of homogenous or divergent host systems to concurrently utilize information that they store on one or more storage devices. The storage facility allows configured storage to be accessible to any, or all, attached host systems. To use this capability, you need to design the host program to support data that it is sharing.

DDM. See *disk drive module*.

DDM group. See *disk drive module group*.

dedicated storage. Storage within a storage facility that is configured such that a single host system has exclusive access to the storage.

demote. The action of removing a logical data unit from cache memory. A subsystem demotes a data unit in order to make room for other logical data units in the cache. It could also demote a data unit because the logical data unit is not valid. A subsystem must destage logical data units with active write units before they are demoted.

destaging. The moving of data from an online or higher priority device to an offline or low priority device.

device. In the ESA/390 architecture, a disk drive.

device adapter (DA). A physical subunit of a storage controller that provides the ability to attach to one or more interfaces used to communicate with the associated storage devices.

device address. In the ESA/390 architecture, the field of an ESCON device-level frame that selects a specific device on a control-unit image.

device interface card. A physical subunit of a storage cluster that provides the communication with the attached DDMs.

device number. (1) In the ESA/390 architecture, a four-hexadecimal-character identifier, for example 13A0, that you associate with a device to facilitate communication between the program and the host operator. (2) The device number that you associate with a subchannel.

device sparing. A subsystem function that automatically copies data from a failing DDM to a spare DDM. The subsystem maintains data access during the process.

direct access storage device (DASD). A mass storage medium on which a computer stores data.

disk drive module (DDM). The primary nonvolatile storage medium that you use for any host data that is stored within a subsystem. The number and type of storage devices within a storage facility may vary.

disk drive module group. A group of eight disk drive modules (DDMs). Each DDM group is a RAID array.

domain. (1) That part of a computer network in which the data processing resources are under common control. (2) In TCP/IP, the naming system used in hierarchical networks.

drawer. A unit that contains multiple DDMs and provides power, cooling, and related interconnection logic to make the DDMs accessible to attached host systems.

drive. (1) A peripheral device, especially one that has addressed storage media. (2) The mechanism used to seek, read, and write information on a storage medium.

duplex. A communication mode in which data can be sent and received at the same time.

dynamic sparing. The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

EBCDIC. See *extended binary-coded decimal interchange code*.

EC. See *engineering change*.

ECKD. See *extended count key data*.

electrostatic discharge (ESD). An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO). A means of turning off power during an emergency, usually a switch.

EMIF. See *ESCON multiple image facility*.

enclosure. A unit that houses the components of a storage subsystem, such as a controller, disk drives, and power source.

end of file. A coded character recorded on a data medium to indicate the end of the medium. On a CKD direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC). An update to a machine, part, or program.

Enterprise Systems Architecture/390 (ESA/390). An IBM architecture for mainframe computers and peripherals. Processor systems that follow this architecture include the ES/9000[®] family.

Enterprise Systems Connection (ESCON). (1) An ESA/390 computer peripheral interface. The I/O interface uses ESA/390 logical protocols over a serial

interface that configures attached units to a communication fabric. (2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO. See *emergency power off*.

ERP. See *error recovery procedure*.

error recovery procedure (ERP). Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used in conjunction with programs that record information on machine malfunctions.

ESA/390. See *Enterprise Systems Architecture/390*.

ESCD. See *ESCON director*.

ESCON. See *Enterprise Systems Connection*.

ESCON director (ESCD). An I/O interface switch that allows the interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON multiple image facility (EMIF). In the ESA/390 architecture, a function that allows LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

ESD. See *electrostatic discharge*.

ESS Specialist. See *StorWatch Enterprise Storage Server Specialist*.

extended binary-coded decimal interchange code (EBCDIC). A coded character set of 256 8-bit codes.

extended count key data (ECKD). An extension of the CKD architecture.

Extended Remote Copy (XRC). A function of a storage controller that assists a control program to maintain a consistent copy of a logical volume on another storage facility. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume.

F

fabric. In fibre-channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *casca*ded.

failover. The routing of all transactions to a second controller when the first controller fails. See *cluster*.

fast write. A write operation at cache speed that does not require immediate transfer of data to a DDM. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FBA. See *fixed-block address*.

FC-AL. See *Fibre Channel - Arbitrated Loop*.

FCP. See *fibre-channel protocol*.

FCS. See *fibre-channel standard*.

Fibre Channel - Arbitrated Loop (FC-AL). An implementation of the fibre-channel technology that uses a ring topology for communication. In this topology two or more fibre-channel end points are interconnected through a looped interface. The ESS supports this topology.

fibre-channel protocol (FCP). For fibre-channel communication the protocol has five layers. The layers define how fibre-channel ports interact through their physical links to communicate with other ports.

fibre-channel standard (FCS). An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols. Refer to ANSI X3.230-199x.

FICON. Fibre-channel connection for mainframe computers and peripherals that use ESCON architecture, such as the IBM ES/9000 family.

field replaceable unit (FRU). An assembly that is replaced in its entirety when any one of its components fails. In some cases, a field replaceable unit may contain other field replaceable units.

FIFO. See *first-in-first-out*.

firewall. A protection against unauthorized connection to a computer or a data storage system. The protection is usually in the form of software on a gateway server that grants access to users that meet authorization criteria.

first-in-first-out (FIFO). A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

fixed-block address (FBA). An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device.

The logical data unit is a block. All blocks on the device are the same size (fixed size). The subsystem can access them independently.

FlashCopy. An optional feature for the ESS. With FlashCopy you can make an instant copy of data, that is, a point-in-time copy of a volume.

FRU. See *field replaceable unit*.

full duplex. See *duplex*.

G

GB. See *gigabyte*.

gigabyte (GB). A gigabyte of storage is 10^9 bytes. A gigabyte of memory is 2^{30} bytes.

group. See *DDM group*.

H

HA. See *host adapter*.

hard disk drive (HDD). (1) A storage medium within a storage server used to maintain information that the storage server requires. (2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

HDA. See *head and disk assembly*.

HDD. See *hard disk drive*.

hdisk. An AIX term for storage space.

head and disk assembly (HDA). The portion of an HDD associated with the medium and the read/write head.

home address. A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

host adapter (HA). A physical subunit of a storage controller that provides the ability to attach to one or more host I/O interfaces.

host processor. A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See *host system*.

host system. (1) A data processing system used to prepare programs and operating environments for use on another computer or controller. (2) The data processing system to which a network is connected and with which the system can communicate. (3) The controlling or highest level system in a data communication configuration.

hot plug. Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

I

IBM product engineering (PE). The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

ID. See *identifier*.

identifier (ID). A unique name or address that identifies things such as programs, devices, or systems.

implicit allegiance. In the ESA/390 architecture, a relationship that a controller creates between a device and a channel path when the device accepts a read or write operation. The controller guarantees access to the channel program over the set of channel paths that it associates with the allegiance.

initial program load (IPL). The initialization procedure that causes an operating system to commence operation.

initiator. A SCSI term for the part of a host computer that communicates with its attached targets.

i-node. The internal structure in an AIX operating system that describes the individual files in the operating system. It contains the code, type, location, and owner of a file.

input/output (I/O). Pertaining to (a) input, output, or both or (b) a device, process, or channel involved in data input, data output, or both.

Internet Protocol (IP). In the Internet suite of protocols, a protocol without connections that routes data through a network or interconnecting networks and acts as an intermediary between the higher protocol layers and the physical network.

invalidate. The action of removing a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal may be the result of a failure within the storage controller or a storage device that is associated with the device.

I/O. See *input/output*.

I/O device. An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface. An interface that you define to allow a host to perform read and write operations with its associated peripheral devices.

IP. See *Internet Protocol*.

IPL. See *initial program load*.

K

KB. See *kilobyte*.

key field. The second (optional) field of a CKD record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field. The subsystem uses this data to identify or locate a given record.

kilobyte (KB). (1) For processor storage, real, and virtual storage, and channel volume, 2^{10} or 1024 bytes. (2) For disk storage capacity and communications volume, 1000 bytes.

KPOH. See *thousands of power-on hours*.

L

LAN. See *local area network*.

last-in first-out (LIFO). A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

least recently used (LRU). (1) The algorithm used to identify and make available the cache space that contains the least-recently used data. (2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LCU. See *logical control unit*.

LED. See *light-emitting diode*.

LIC. See *licensed internal code*.

licensed internal code (LIC). Microcode that IBM does not sell as part of a machine, but licenses to the customer. LIC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wired circuitry.

LIFO. See *last-in first-out*.

light-emitting diode (LED). A semiconductor chip that gives off visible or infrared light when activated.

link address. On an ESCON interface, the portion of a source or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel-subsystem or controller-link-level functions that are attached to the switch port.

link-level facility. The ESCON hardware and logical functions of a controller or channel subsystem that allow communication over an ESCON write interface and an ESCON read interface.

local area network (LAN). A computer network located on a user's premises within a limited geographic area.

local e-mail. An e-mail configuration option for storage servers that are connected to a host-system network that does not have a domain name server (DNS).

logical address. On an ESCON interface, the portion of a source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical control unit (LCU). See *controller image*.

logical data unit. A unit of storage that is accessible on a given device.

logical device. A file for conducting input or output with a physical device, or a file for mapping user I/O between virtual and real devices.

logical partition (LPAR). A set of functions that create the programming environment that is defined by the ESA/390 architecture. ESA/390 architecture uses this term when more than one LPAR is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also the LPAR does not depend on an operating system to create the virtual machine environment.

logical subsystem (LSS). The logical functions of a storage controller that allow one or more host I/O interfaces to access a set of devices. The controller aggregates the devices according to the addressing mechanisms of the associated I/O interfaces. One or more logical subsystems exist on a storage controller. In general, the controller associates a given set of devices with only one logical subsystem.

logical unit. The open-systems term for a logical disk drive.

logical unit number (LUN). The SCSI term for the field in an identifying message that is used to select a logical unit on a given target.

logical volume. The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. A logical volume is referred to on an AIX platform as an hdisk, an AIX term for storage space. A host system logical volume can be a physical volume.

logical volume manager (LVM). A set of system commands, library routines, and other tools that allow

the user to establish and control logical volume storage. The LVM maps data between the logical view of storage space and the physical disk drive module (DDM).

LPAR. See *logical partition*.

LRU. See *least recently used*.

LSS. See *logical subsystem*.

LUN. See *logical unit number*.

LVM. See *logical volume manager*.

M

machine level control (MLC). A database that contains the EC level and configuration of products in the field.

maintenance analysis procedure (MAP). A hardware maintenance document that gives an IBM service representative a step-by-step procedure for tracing a symptom to the cause of a failure.

management information base (MIB). The physical and logical characteristics of a system make up a collection of information called a management information base (MIB). The individual pieces of information that comprise an MIB are called MIB objects and they reside in the SNMP Agent. See *simple network management protocol*.

MAP. See *maintenance analysis procedure*.

MB. See *megabyte*.

MCA. See *Micro Channel architecture*.

mean time between failures (MTBF). (1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions. (2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (1) (A)

medium. For a storage facility, the disk surface on which data is stored.

megabyte (MB). (1) For processor storage, real and virtual storage, and channel volume, 2²⁰ or 1 048 576 bytes. (2) For disk storage capacity and communications volume, 1 000 000 bytes.

MES. See *miscellaneous equipment system*.

MIB. See *management information base*.

Micro Channel architecture (MCA). The rules that define how subsystems and adapters use the Micro Channel bus in a computer. The architecture defines the services that each subsystem can or must provide.

mirrored pair. Two units that contain the same data. The system refers to them as one entity.

mirroring. In the AS/400 host systems, the process of writing the same data to two disk units within the same auxiliary storage pool at the same time.

miscellaneous equipment specification.

MLC. See *machine level control*.

mobile service terminal (MoST). The mobile terminal used by IBM service personnel.

MoST. See *mobile service terminal*.

MTBF. See *mean time between failures*.

multiple virtual storage (MVS). Consisting of MVS/System Product Version 1 and MVS/370 Data Facility Product operating on an IBM System/370 processor.

MVS. See *multiple virtual storage*.

N

node. The unit that is connected in a fibre-channel network. An ESS is a node in a fibre-channel network.

non-RAID. A collection of disk drives in a non-RAID group or array.

non-removable medium. A recording media that you cannot add to or remove from a storage device.

non-retentive data. Data that the control program can easily recreate in the event it is lost. The control program may cache non-retentive write data in volatile memory.

nonvolatile storage (NVS). (1) Typically refers to nonvolatile memory on a processor rather than a nonvolatile storage device. On a storage facility, nonvolatile storage is used to store active write data to avoid data loss in the event of a power loss. (2) A storage device whose contents are not lost when power is cut off.

NVS. See *nonvolatile storage*.

O

OEMI. See *original equipment manufacturer's information*.

open system. A system whose characteristics comply with standards made available throughout the industry

and that therefore can be connected to other systems complying with the same standards.

organizationally unique identifier (OUI). An identifier that identifies an organization according to IEEE standards.

original equipment manufacturer's information (OEMI). A reference to an IBM guideline for a computer peripheral interface. More specifically, refer to *IBM System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufactures' Information*. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multi-drop bus topology.

OUI. See *organizationally unique identifier*.

P

parallel access volume (PAV). Created by associating multiple devices of a single control-unit image with a single logical device. Up to 8 device addresses can be assigned to a parallel access volume.

parity. A data checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to recreate data if a disk drive fails.

path group. The ESA/390 term for a set of channel paths that are defined to a controller as being associated with a single LPAR. The channel paths are in a group state and are online to the host.

path-group identifier. The ESA/390 term for the identifier that uniquely identifies a given LPAR. The path-group identifier is used in communication between the LPAR program and a device. The identifier associates the path-group with one or more channel paths, thereby defining these paths to the control-unit as being associated with the same LPAR.

PAV. See *parallel access volume*.

PCI. See *program-controlled interruption*.

PE. See *IBM product engineering*.

Peer-to-Peer Remote Copy (PPRC). A function of a storage controller that maintains a consistent copy of a logical volume on the same controller or on another controller. You access the controllers over an ESCON I/O interface. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume.

peripheral component interconnection. *see program-controlled interruption*.

point-to-point connection. A fibre-channel topology that allows you to interconnect ports directly.

POST. See *power-on self test*.

power-on self test (POST). A diagnostic test run by servers or computers when they are turned on.

PPRC. See *Peer-to-Peer Remote Copy*.

predictable write. A write operation that can cache without knowledge of the existing format on the medium. All writes on FBA DASD devices are predictable. On CKD DASD devices, a write is predictable if it does a format write for the first data record on the track.

product engineering. See *IBM product engineering*.

program. A generic term for the software that controls the operation of a host computer. Typically, the program is an operating system that allows sharing of the host resources between multiple tasks.

program-controlled interruption (PCI). An interruption that occurs when an I/O channel fetches a channel command word with the program-controlled interruption flag on. PCI is also called peripheral component interconnection.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of a program

promote. The action of adding a logical data unit to cache memory.

PTF. See *program temporary fix*.

R

R0. See *track-descriptor record*.

rack. See *enclosure*.

RAID. See *redundant array of independent disks and array*.

random access. A mode of accessing data on a medium in a manner that requires the storage device to access nonconsecutive storage locations on the medium.

rank. See *array*.

redundant array of independent disks (RAID). A collection of disk drives that operate independently. The IBM Enterprise Storage Server protects all storage with redundant arrays of independent disks (RAID). The ESS attaches serial storage architecture (SSA) disk drive modules (DDMs) in RAID-5 configurations. The implementation of RAID-5 configurations distributes (stripes) parity across all DDMs in the array. Also see *array*.

remote technical assistance information network (RETAIN). Initial Service Tracking for IBM service support.

reserved allegiance. In the ESA/390 architecture, a relationship that is created in a controller between a device and a channel path when a Sense Reserve command is completed by the device. The allegiance causes the control-unit to guarantee access (busy status is not presented) to the device. Access is over the set of channel paths that are associated with the allegiance; access is for one or more channel programs, until the allegiance ends.

RETAIN. See *remote technical assistance information network*

S

SAM. See *sequential access method*.

SBCON. See *Single-Byte Command Code Sets Connection*.

SCSI. See *small computer system interface*.

SCSI ID. A unique identifier assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs. A SCSI device is either an initiator or a target.

Seascape architecture. A storage system architecture developed by IBM for open-systems servers and S/390 host systems. It provides storage solutions that integrate software, storage management, and technology for disk, tape, and optical storage.

sequential access method (SAM). A sequential access method to find information.

self-timed interface (STI). An interface that has one or more conductors that transmit information serially between two interconnected units without requiring any clock signals to recover the data. The interface performs clock recovery independently on each serial data stream and uses information in the data stream to determine character boundaries and inter-conductor synchronization.

server group. See *Copy Services server group*.

sequential access. A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

serial storage architecture (SSA). An IBM standard for a computer peripheral interface. The interface uses a SCSI logical protocol over a serial interface that configures attached targets and initiators in a ring topology.

server. (1) A type of host that provides certain services to other hosts that are referred to as clients. (2) A functional unit that provides services to one or more clients over a network.

service boundary. Physical space on an ESS rack. It consists of a bay that contains SCSI and ESCON host adapter cards and the attached devices (DDMs). If a single component in a service boundary is being serviced, none of its components are available to the ESS.

service-information message (SIM). A message sent by a storage server to service personnel.

service processor. A dedicated processing unit used to service a storage facility.

service support representative (SSR). Individuals or a company that you have authorized to service your ESS. This term also refers to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs your ESS.

shared storage. Storage within a storage facility that is configured such that multiple homogenous or divergent hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. You need to design the programs to handle the effects of concurrent access.

short wave. A technology that allows you to connect a variety of communication protocols.

SIM. See *service-information message*.

simple network management protocol (SNMP). (1) An SNMP Manager is part of a network administration software product, such as Tivoli TME 10 Netview. (2) SNMP permits you to define management information base (MIB) extensions, or enterprise-specific MIBs. By loading a MIB file on a manager station, you can monitor the MIB objects. The SNMP Agent accesses MIB objects when a request is received from an SNMP Manager. (3) In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

Single-Byte Command Code Sets Connection (SBCON). The ANSI standard for the ESCON I/O interface.

small computer system interface (SCSI). (1) An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface uses a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology. (2) A standard hardware

interface that enables a variety of peripheral devices to communicate with one another.

SMIT. See *System Management Interface Tool*.

SMP. See *System Modification Program*.

SNMP. See *simple network management protocol*.

software transparency. Criteria applied to a processing environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

spare. A disk drive that is used to receive data from a device that has experienced a failure that requires disruptive service. A spare can be pre-designated to allow automatic dynamic sparing. Any data on a disk drive that you use as a spare is destroyed by the dynamic sparing copy process.

SSA. See *serial storage architecture*.

SSA adapter. The adapter that connects stored data on devices, such as disk drive modules for access and control by the storage server. See *serial storage architecture*.

SSID. See *subsystem identifier*.

SSR. See *service support representative*.

stacked status. In the ESA/390 architecture, the condition when the control-unit is holding status for the channel and the channel responded with the stack-status control the last time the control-unit attempted to present the status.

staging. The action of moving data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

STI. See *self-timed interface*.

storage complex. Multiple storage facilities.

storage controller. A physical unit that provides an interface between one or more storage devices and a host computer by providing the function of one or more logical subsystems. The storage controller may provide functions that are not provided by the storage device. The storage controller has one or more clusters.

storage device. A physical unit that provides a mechanism to store data on a given medium such that it can be subsequently retrieved. See *disk drive module*.

storage facility. (1) A physical unit that consists of a storage controller integrated with one or more storage devices to provide storage capability to a host computer. (2) A storage server and its attached storage devices.

storage server. A unit that manages attached storage devices and provides access to the storage or storage-related functions for one or more attached hosts.

StorWatch Enterprise Storage Server Specialist. The Web-based management interface to the IBM Enterprise Storage Server.

striping. A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subchannel. A logical function of a channel subsystem associated with the management of a single device.

subsystem identifier (SSID). A number that uniquely identifies a logical subsystem within a computer installation.

support catcher telephone number. The telephone number that service personnel use to send heartbeat call-home records that are non-problem related, directly to the support catcher server rather than to the RETAIN catcher server. These records might include such information as a trace or a dump package.

synchronous write. A write operation that stores data in cache and in nonvolatile storage (NVS). Later the data is stored on a device, using normal destage algorithms.

System Management Interface Tool (SMIT). An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program (SMP). A program used to install software and software changes on MVS systems.

T

target. A storage device to which information is written.

TB. See *terabyte*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

terabyte (TB). (1) For processor storage, real and virtual storage, and channel volume, 2^{40} or 1 099 511 627 776 bytes. (2) For disk storage capacity and communications volume, 1 000 000 000 000 bytes.

thousands of power-on hours (KPOH). A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO). An operating system option, the option provides interactive time sharing from remote terminals.

TPF. See *transaction processing facility*.

track. A unit of storage on a CKD device that can be formatted to contain a number of data records. See *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0). A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF). A high-availability, high-performance system, designed to support real-time, transaction driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and update to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol/Internet Protocol (TCP/IP). (1) The Transmission Control Protocol and the Internet Protocol, which together provide reliable end-to-end connections between applications over interconnected networks of different types. (2) The suite of transport and application protocols that run over the Internet Protocol.

transparency. See *software transparency*.

TSO. See *time sharing option*.

U

UFS. UNIX filing system.

Ultra-SCSI. An enhanced small computer system interface.

unit address. The ESA/390 term for the address associated with a device on a given controller. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a controller and device pair on the interface.

Universaile Tempes du Coordinaire (UTC). The basis for coordinated world wide time. This system, formerly known as Greenwich Mean Time (GMT), was officially renamed to UTC, for 'Universaile Tempes du Coordinaire' because the world clock is now in Paris, France. It is also known as coordinated universal time (CUT).

UTC. See *Universaile Tempes du Coordinaire*.

utility device. The ESA/390 term for the device used with the Extended Remote Copy facility to access information that describes the modifications performed on the primary copy.

V

virtual machine (VM). A virtual data processing machine that appears to be for the exclusive use of a particular user, but whose functions are accomplished by sharing the resources of a real data processing system.

vital product data (VPD). Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM. See *virtual machine*.

volume. In the ESA/390 architecture, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the terms may also refer, indirectly, to the storage device that you associate with the volume. When you store multiple volumes on a single storage medium transparently to the program, you may refer to the volumes as logical volumes.

VPD. See *vital product data*.

W

world-wide port name (WWPN). A unique 64-bit identifier for a fibre-channel port that is assigned in an implementation and protocol-independent manner.

write hit. A write operation in which the data that you requested is in the cache.

write penalty. The performance impact of a classical RAID write operation.

WWPN. See *world-wide port name*.

X

XRC. See *Extended Remote Copy*.

Index

Numerics

10BASE-T cable 18
524-byte sector format 26

A

AAS 57
access
 definition 57
 ESS Specialist 6
 ESS Web Copy Services 6
adapter
 device x
 host xi
address
 fixed-block (definition) 61
 for communication 6
 home (definition) 62
 link (definition) 63
 TCP/IP 18
 unit 68
Advanced Administrative Systems
 definition 57
agreement for Licensed Internal
 Code 53
AIX commands
 backup 25
 dump 25
 restore 25
alert (definition) 57
allegiance (definition) 57
allocated storage
 definition 57
 FB and CKD 7
American National Standards Institute
 (definition) 57
analyzing
 problems that occur during
 operation 39
ANSI (definition) 57
APAR (definition) 57
arbitrated loop
 definition 57
 illustration of 11
archive data 29
array
 definition 57
 ESS definition ix
 RAID 1
AS/400
 data migration 27
 support for SCSI attachment 3
AS/400 systems
 data migration for 27
ASCII (definition) 57
assigned storage (definition) 57
attaching the ESS Network to your local
 area network 17
authorized program analysis report
 (definition) 57
availability
 continuous 2

availability (*continued*)
 definition 57

B

backup and restore 29
backup and restore methods for data
 migration 35
battery disposal vii
bay (definition) 58
bit (definition) 58
block (definition) 58
browser
 MIB 41
byte (definition) 58

C

cable
 10BASE-T 18
 connection to the ESSNet 5
 length from the ESS to the modem 9
cache
 definition 58
 fast write (definition) 58
 memory (definition) 58
 options 1
Canadian compliance statement 51
cascaded 58
caution notice vii
CCR (definition) 58
CCW (definition) 58
CEC (definition) 58
channel
 command retry (definition) 58
 command word (definition) 58
 definition 58
 ESCON x
 path (definition) 58
 subsystem (definition) 58
 subsystem image (definition) 58
CKD
 definition ix
 disk groups 6
CKD (definition) 58
CLI (definition) 58
cluster
 complex (definition) 58
 definition 58
command-line interface (definition) 58
communications statement 51
compact disc (definition) 58
compliance statement
 Canadian 51
 European 51
 German 51
 Japanese 52
 Korean 52
 radio frequency energy 51
 Taiwan 52
compression (definition) 58

computer-electronic complex
 (definition) 58
concurrent
 download 58
 download of licensed internal
 code 13
 maintenance (definition) 58
 media maintenance (definition) 59
Concurrent Copy (definition) 58
configure
 definition 59
 with ESS Specialist 6
connection
 fibre-channel 10
 Web 5
consistent copy (definition) 59
console
 customer (definition) 59
 definition 59
contingent allegiance (definition) 59
control panel
 expansion enclosure 17
 indicators 15
 switches 15
controller
 definition 59
 image ix
 image (definition) 59
Copy Services
 access 6
 CLI 7
 command-line interface 7
 description 7
 panels 8
 server group 7
 Web-based interface 7
Copy Services server group
 (definition) 59
copying a complete logical volume 29
count field (definition) 59
count key data (definition) 59
CPC (definition) 59
cpio command 25
CRC (definition) 59
CU (definition) 59
customer console (definition) 59
cyclic redundancy check (definition) 59
cylinder (definition) 59

D

DA (definition) 59
daisy chain (definition) 59
danger notice vii
DASD (definition) 59
DASD-fast write (definition) 59
DASD SIM 40
data
 archive 29
 availability (definition) 59
 compression (definition) 59
 field (definition) 59

- data (*continued*)
 - migrating 25
 - record (definition) 59
 - sharing 60
 - Data Facility Storage Management
 - Subsystem (definition) 59
 - data migration
 - affect on data availability 34
 - AS/400 systems 27
 - backup and restore methods 35
 - direct copy method 35
 - for a VM environment 36
 - for a VSE environment 36
 - for an MVS environment 35
 - for MVS, VM, and VSE operating systems 35
 - from System/390 operating systems 35
 - publications 35
 - services 37
 - UNIX systems 27
 - DDM 60
 - DDM (definition) 60
 - dedicated storage (definition) 60
 - demote (definition) 60
 - destaging (definition) 60
 - device
 - adapter x
 - adapter (definition) 60
 - address (definition) 60
 - definition 60
 - fixed block x
 - interface card (definition) 60
 - logical xi
 - number (definition) 60
 - sparing (definition) 60
 - direct access storage device (definition) 60
 - direct copy method for data migration 35
 - disk drive module
 - definition 60
 - group (definition) 60
 - disposal, product vii
 - documents, ordering xiii
 - domain (definition) 60
 - drawer (definition) 60
 - drive (definition) 60
 - duplex (definition) 60
 - dynamic sparing (definition) 60
- E**
- EBCDIC (definition) 60
 - EC (definition) 60
 - ECKD (definition) 60
 - edition notice ii
 - electronic emission notices 51
 - emergency power off (definition) 60
 - emergency power switch 15
 - EMIF (definition) 60
 - enclosure
 - definition 60
 - expansion 2
 - end of file (definition) 60
 - engineering changes 60
 - Enterprise Storage Server Network
 - connection 5
 - Enterprise Storage Server Network (*continued*)
 - description 5
 - workstation 5
 - Enterprise Storage Server overview 1
 - enterprise systems
 - Architecture/390 (definition) 60
 - connection (definition) 60
 - connection architecture x
 - environmental notices vii
 - EPO (definition) 61
 - ERP (definition) 61
 - error message 39
 - error recovery procedure (definition) 61
 - ESA/390 (definition) 61
 - ESCD (definition) 61
 - ESCON
 - channel x
 - definition 61
 - director (definition) 61
 - multiple image facility (definition) 61
 - ESD (definition) 60, 61
 - ESS
 - Copy Services 7
 - Expert 9
 - host systems supported by 3
 - overview 1
 - publications xii
 - Specialist 7
 - Specialist (definition) 61
 - ESSNet
 - description 5
 - European Community Compliance statement 51
 - expansion
 - enclosure 2
 - expansion enclosure
 - control panel 17
 - Expert
 - description 9
 - extended
 - binary-coded decimal interchange code (definition) 61
 - count key data (definition) 61
 - Extended Remote Copy (definition) 61
 - extent of warranty 55
- F**
- fabric
 - definition 61
 - illustration of 11
 - failover 12
 - failover (definition) 61
 - fast write (definition) 61
 - FB
 - disk groups 6
 - FBA (definition) 61
 - FC-AL (definition) 61
 - FCP (definition) 61
 - FCS (definition) 61
 - Federal Communications Commission (FCC) statement 51
 - fibre-channel
 - arbitrated loop (definition) 61
 - connection 10
 - host systems 3
 - protocol (definition) 61
 - fibre-channel (*continued*)
 - standard (definition) 61
 - topologies 10
 - FICON
 - definition 61
 - field replaceable unit (definition) 61
 - FIFO (definition) 61
 - firewall
 - definition 61
 - for Copy Services 8
 - first-in-first-out (definition) 61
 - fixed block
 - devices x
 - fixed-block
 - address (definition) 61
 - FlashCopy
 - definition 62
 - frequently used terms ix
 - FRU (definition) 62
 - full duplex (definition) 62
- G**
- GB (definition) 62
 - German compliance statement 51
 - getting service 47
 - gigabyte
 - capacity x
 - definition 62
 - glossary 57
 - group (definition) 60
 - guidelines for defining the logical control unit 22
 - guidelines for migrating data to an ESS 35
- H**
- HA (definition) 62
 - hard disk drive (definition) 62
 - HDA (definition) 62
 - HDD (definition) 62
 - hdisk
 - AIX definition xi
 - definition 62
 - head and disk assembly (definition) 62
 - home address (definition) 62
 - host
 - adapter (definition) 62
 - processor (definition) 62
 - system (definition) 62
 - host system
 - ESCON x
 - fibre channel 3
 - S/390 4
 - SCSI xi, 3
 - supported by the ESS 3
 - host system views of the IBM Enterprise Storage Server 26
 - hostname 18
 - hot plug (definition) 62
- I**
- i-node (definition) 62
 - I/O
 - definition 62

I/O (*continued*)
 device (definition) 62
 interface (definition) 62
 IBM product engineering (definition) 62
 IBM Subsystem Device Driver
 overview 12
 Web site xv, 15
 ID (definition) 62
 ID, SCSI xii
 identifier (definition) 62
 image
 controller ix
 implicit allegiance (definition) 62
 Industry Canada Compliance
 statement 51
 information message, service 39
 initial program load (definition) 62
 initiator
 definition 62
 input/output (definition) 62
 input/output configuration program
 (IOCP)
 overview 20
 reference publications 20
 interface
 service 9
 Web 4, 6
 Internet Protocol (definition) 62
 invalidate (definition) 62
 IOCP
 example 20
 overview 20
 IP (definition) 62
 IPL (definition) 63

J

Japanese Voluntary Control Council for
 Interference (VCCI) statement 52

K

KB (definition) 63
 key field (definition) 63
 kilobyte (definition) 63
 Korean government Ministry of
 Communication (MOC) statement 52
 KPOH (definition) 63

L

LAN (definition) 63
 LAN. See *local area network* 15
 LCU
 guidelines for defining 22
 least recently used (definition) 63
 LED (definition) 63
 liability, limitation of 56
 LIC
 agreement 53
 licensed internal code
 agreement 53
 definition 63
 LIFO (definition) 63
 light, message 39
 light-emitting diode (definition) 63
 limitation of liability 56

limited warranty statement (SOLW) 54
 link address (definition) 63
 link-level facility (definition) 63
 local area network
 attach to ESNet 17
 problems 17
 local area network (definition) 63
 local e-mail (definition) 63
 logical
 address (definition) 63
 control unit (definition) 63
 data unit (definition) 63
 device xi
 device (definition) 63
 partition (LPAR) (definition) 63
 subsystem xi
 subsystem (definition) 63
 unit xi
 unit (definition) 63
 unit number xi
 unit number (definition) 63
 volume xi
 volume (definition) 63
 logical control unit
 guidelines for defining 22
 logical volume management software
 using 28
 logical volume manager (definition) 63
 Logical Volume Manager (LVM) 29
 loop
 arbitrated 11
 LPAR (definition) 64
 LRU (definition) 64
 LSS (definition) 64
 LSS. See *logical subsystem*. xi
 LUN
 definition 64
 LUN. See *logical unit number*. xi
 LVM 64
 AIX 29
 using 28

M

machine level control (definition) 64
 machines, warranty for 54
 mail server 18
 maintenance analysis procedure
 (definition) 64
 maintenance or warranty support 47
 management information base
 (definition) 64
 manuals, ordering xiii
 MAP (definition) 64
 MB (definition) 64
 MCA (definition) 64
 media SIM 40
 medium (definition) 64
 megabyte (definition) 64
 MES (definition) 64
 message
 error 39
 service information 39
 message light 39
 MIB (definition) 64
 MIB browser 41

Micro Channel architecture
 (definition) 64
 migrating data from UNIX, Windows NT,
 and OS/400 25
 migrating data to the IBM Enterprise
 Storage Server 25
 mirror
 definition 64
 pair (definition) 64
 mirroring 31
 miscellaneous
 definition 64
 MLC (definition) 64
 mobile service terminal
 definition 64
 external connection 9
 MoST (definition) 64
 MoST. See *mobile service terminal*. 9
 MTBF (definition) 64
 multiple virtual storage (definition) 64
 MVS (definition) 64

N

network
 local area 17
 node (definition) 64
 non-RAID (definition) 64
 non-removable medium (definition) 64
 non-retentive data (definition) 64
 nonvolatile storage (definition) 64
 notices ii
 caution vii
 danger vii
 electronic emission 51
 environmental vii
 IBM 49
 safety vii
 NVS (definition) 64

O

OEMI (definition) 64
 open system (definition) 64
 operating system
 AIX 3
 HP-UX 3
 Novell Netware 3
 OpenVMS 3
 OS/400 3
 Solaris 3
 Tru64 UNIX 3
 Windows NT 3
 ordering publications xiii
 organizationally unique identifier
 (definition) 65
 original equipment manufacturer's
 information (definition) 65
 OUI (definition) 65
 overview
 IBM Enterprise Storage Server 1
 input/output configuration program
 (IOCP) 20
 parallel access volumes 19

P

parallel access volume (definition) 65

- parallel access volumes
 - base and aliases 19
 - overview 19
 - requirements 19
- parity (definition) 65
- path group
 - definition 65
 - identifier (definition) 65
- path recovery algorithms 13
- path selection policy 13
- PAV 65
- PCI (definition) 65
- PE (definition) 65
- Peer-to-Peer Remote Copy
 - (definition) 65
- peripheral component interconnection
 - definition 65
- personnel, service xii
- physical volume migration 30
- picture
 - of ESS Copy Services server group connections 8
 - of the expansion enclosure 2
 - of the Models E10, E20, F10, and F20 front and rear views 2
- point-to-point
 - connection (definition) 65
 - illustration of 10
- POST (definition) 65
- power
 - complete 15
 - emergency switch 15
 - LED 16
 - local 16
 - switch 16
 - turn off 15
 - turn on 15
- power-on self test (definition) 65
- PPRC 65
- predictable write (definition) 65
- problem analysis 39
- problem state 44
- problems
 - that occur during operation 39
- problems on the local area network 17
- product
 - disposal vii
 - recycling vii
- product engineering (definition) 62, 65
- program
 - controlled interruption (definition) 65
 - definition 65
 - temporary fix (definition) 65
- promote (definition) 65
- PTF (definition) 65
- publications
 - ESS xii
 - library xii
 - ordering xiii
 - related xiii
- publications about data migration 35

R

- R0 (definition) 65
- rack (definition) 65

- radio frequency energy compliance
 - statement 51
- RAID (definition) 65
- random access (definition) 65
- rank (definition) 65
- recycling, product vii
- redundant array of independent disks
 - (definition) 65
- reformatting existing disk drive
 - modules 26
- related publications xiii
- remote service 47
- remote services support 9
- remote support 9
 - access 9
- repair 47
- replacing existing storage 26
- requirements
 - parallel access volumes 19
- reserved allegiance (definition) 66
- restore 29
- RETAIN (definition) 66

S

- S/390
 - host systems 4
 - operating system 4
- safety notices vii
- SAN Data Gateway
 - Web site xvi
- SBCON (definition) 66
- SCSI
 - definition 66
 - host systems xi, 3
 - ID xii
 - ID (definition) 66
- Seascape architecture (definition) 66
- sector
 - fixed-byte 26
- security
 - Web connection 5
- self-timed interface (definition) 66
- sequential access (definition) 66
- sequential access method (definition) 66
- serial storage architecture (definition) 66
- server
 - definition 66
 - mail 18
 - storage xii
- server group (definition) 66
- service
 - boundary (definition) 66
 - getting 47
 - information message (definition) 66
 - interface 9
 - personnel xii
 - processor (definition) 66
 - remote 47
 - remote support 9
- service-information message
 - System/390 39
- service support representative
 - (definition) 66
- service warranty 55
- setting up and operating the Enterprise Storage Server 15
- shared storage (definition) 66

- short wave
 - definition 66
- SIM
 - DASD 40
 - media 40
 - severity level 40
 - storage server 40
 - System/390 39
- SIM (definition) 66
- simple network management protocol
 - (definition) 66
- Single-Byte Command Code Sets
 - Connection (definition) 66
- sites, Web browser xv
- small computer system interface
 - (definition) 66
- SMIT (definition) 67
- SMP (definition) 67
- SNMP (definition) 67
- software
 - open-systems hosts 3
 - S/390 4
 - transparency (definition) 67
 - volume management 25
- SOLW (limited warranty statement) 54
- spare (definition) 67
- Specialist 6
 - access 6
- splittvcopy command 35
- SSA adapter (definition) 67
- SSID (definition) 67
- SSR (definition) 67
- stacked status (definition) 67
- staging (definition) 67
- statement of compliance
 - Canadian 51
 - European 51
 - Federal Communications Commission 51
 - German 51
 - Japanese 52
 - Korean government Ministry of Communication (MOC) 52
 - Taiwan 52
- statement of limited warranty 54
- STI (definition) 67
- storage
 - complex (definition) 67
 - controller (definition) 67
 - device (definition) 67
 - facility (definition) 67
 - partitioning 26
 - replacing existing 26
 - server xii
 - server (definition) 67
 - shared (definition) 66
- storage server SIM 40
- StorWatch Enterprise Storage Specialist
 - (definition) 67
- striping (definition) 67
- subchannel (definition) 67
- subsystem
 - identifier (definition) 67
- Subsystem Device Driver
 - introduction 12
 - Web site xv
- support 47

- support 47 (*continued*)
 - maintenance 47
 - remote services 9
 - warranty 47
- support, technical 47
- support catcher telephone number (definition) 67
- switch
 - power 16
- synchronous write (definition) 67
- System/390
 - data migration 35
- System Management Interface Tool (definition) 67
- System Modification Program (definition) 67
- systems
 - AS/400 27
 - ESCON host x
 - S/390 host 4
 - SCSI host xi, 3
 - UNIX 27

T

- Taiwan class A compliance statement 52
- target
 - host system SCSI IDs 27
- target (definition) 67
- TB (definition) 67
- TB. See *terabyte*. xii
- TCP/IP (definition) 67
- TCP/IP address 18
- technical support 47
- telephone
 - support catcher number (definition) 67
- terabyte
 - capacity xii
 - definition 67
- terms, frequently used ix
- thousands of power-on hours (definition) 67
- time sharing option(definition) 67
- topology
 - arbitrated loop 11
 - point-to-point 10
 - switched fabric 10
- TPF (definition) 68
- track (definition) 68
- track-descriptor record (definition) 68
- trademarks 50
- transaction processing facility (definition) 68
- Transmission Control Protocol/Internet Protocol (definition) 68
- transparency (definition) 68
- TSO (definition) 68
- turn power off 15
- turn power on 15
- turning the power on and off 16

U

- UFS (definition) 68
- Ultra SCSI (definition) 68
- unit
 - logical xi

- unit address (definition) 68
- Universaile Tempes du Coordinaire (definition) 68
- UNIX systems
 - data migration for 27
- using
 - direct copy 28
 - migratepv command 30
 - mirroring for data migration 31
 - mirrorvg command 33
- using logical volume management software 28
- using parallel access volumes for System/390 hosts 19
- using the control panel 15
- utility device (definition) 68

V

- virtual machine (definition) 68
- vital product data (definition) 68
- VM (definition) 68
- volume
 - definition 68
 - logical xi
 - management methods for data migration 34
- VPD (definition) 68

W

- warranty
 - extent 55
 - for machines 54
 - limited 54
 - service 55
- Web browser
 - connection security 5
 - interface 4, 6
 - Microsoft Internet Explorer 5
 - Netscape Navigator 6
- Web interfaces 4
- Web site
 - ESS publications xv
 - host systems supported by the ESS xv
 - IBM storage servers xv
 - IBM Subsystem Device Driver xv
 - SAN Data Gateway xvi
- world-wide port name
 - definition 68
- write
 - hit (definition) 68
- write penalty 68

X

- XRC (definition) 68



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC26-7295-02

