

## Overview

Introduction .....	1
Course Objectives.....	2
Prerequisites .....	2
Module Overview .....	3

## Module 1 — Backup and Restore Basics

Objectives .....	1
Evolution of Backup Methods.....	2
Basic Server Backup .....	3
Centralized Server Backup.....	4
Automated Centralized Backup across the LAN .....	5
LAN-Free Backup.....	6
Off-Host Backup (Non-Disruptive Backup) with Snapshots.....	7
Off-Host Backup (Non-Disruptive Backup) with Clones.....	8
Clones and Snapshots Comparison .....	9
Third-Party Copy (Serverless Backup) .....	9
Backup Performance .....	12
Background .....	12
SCSI Evolution .....	13
SCSI-1.....	13
SCSI-2.....	13
SCSI-3.....	14
Backup Speeds and Feeds .....	15
Performance Issues .....	16
Configuration of Primary Storage.....	16
Type of Connection.....	17
Type and Configuration of Tape Drives .....	17
Compression Ratio and Data Verification .....	17
Block Sizes and Formatting of Tapes .....	18
Choosing the Optimal Backup Solution.....	18
Speed.....	19
Capacity .....	22
Hardware and Media Costs.....	23
Hardware Reliability.....	25
Backup Strategies .....	27
Is the Backup Full or Partial?.....	27
Full Backup.....	27
Partial Backup.....	27
Incremental .....	28
Differential.....	28
User-Defined.....	29
Is the Backup Image- or File-Based? .....	30
How Big Are the Backup and Restore Windows? .....	30
How Often Should Backups Take Place? .....	31
Rotating Tapes .....	32
Grandfather-Father-Son Tape Rotation.....	33
GFS Example.....	34
Other Tape Storage Strategies .....	35
Archiving .....	35
Hierarchical Storage Management.....	35
Learning Check .....	36

## Module 2 — Tape Drive Technologies

Objectives .....	1
Tape Drive Technologies .....	2
Helical Scan Recording.....	2
AIT Drives and Media .....	2
Linear Recording.....	4
DLT Drives .....	4
SDLT Drives.....	6
LTO Technology .....	9
Ultrium Generation 1 Drives .....	10
Ultrium Generation 2 Drives .....	17
Ultrium/SDLT Comparison .....	19
Learning Check .....	20

## Module 3 — MSL5000 and SSL2020 Tape Libraries

Objectives .....	1
HP StorageWorks Tape Libraries.....	3
Tape Libraries .....	3
HP StorageWorks Tape Library Nomenclature.....	3
SSL2020 AIT Library.....	5
MSL5000 Library Family Overview .....	7
Learning Check .....	38

## Module 4 — ESL9000 Tape Libraries

Objectives .....	1
ESL9000 Overview .....	2
ESL9000 Family of Tape Libraries.....	2
HP StorageWorks Tape Library Nomenclature.....	2
ESL9000 Family .....	3
Library GUI and Operations for the ESL9000 Family .....	17
SCSI Target Guidelines for Tape Libraries .....	26
Fibre Channel LUN Guidelines for Tape Libraries .....	26
Learning Check .....	27

## Module 5 — Network Storage Routers

Objectives .....	1
Network Storage Router Overview .....	2
Converting Fibre Channel to SCSI.....	4
Converting SCSI to Fibre Channel.....	5
Network Storage Router M2402 .....	6
Hardware Components.....	7
Main Board .....	7
Fibre Channel Module .....	8
SCSI Module .....	8
Network Storage Router N1200 .....	9
Embedded Network Storage Routers .....	10
Network Storage Router e1200.....	10
Network Storage Router e2400.....	10
Visual Manager User Interface.....	11
System Menu.....	13
Modules Menu .....	14
Discovery Menu .....	15
Mapping Menu .....	16
Statistics Menu .....	17
Utilities Menu .....	18
Report Menu .....	19
Reboot Menu.....	19
Optimizing Performance on the Network Storage Router .....	20
Learning Check .....	21

## Module 6—Software Components of the Enterprise Backup Solutions

Objectives .....	1
EBS Software Components .....	2
Maximum Supported Configuration .....	3
Drivers .....	5
HBA Drivers .....	5
Tape Drivers .....	5
Fibre Channel Filter Drivers .....	5
Device Persistency .....	6
Backup Software .....	7
HP OpenView Data Protector Overview .....	8
Key Features .....	9
Data Protector Topology.....	9
Large Data Protector Enterprise Environments .....	10
Backup Exec Overview.....	13
Key Features .....	14
Backup Exec Topology.....	15
ARCserve Overview .....	17
Key Features .....	17
ARCserve Topology .....	19
NetBackup Overview .....	21
Key Features .....	21
NetBackup Topology.....	23
NetWorker Overview .....	25
Key Features .....	25
NetWorker Topology.....	27
NetWorker and AlphaStor .....	29
Tivoli Storage Manager Overview.....	30
Key Features .....	30
Tivoli Storage Manager Topology.....	31
CommVault Galaxy .....	33
Key Features .....	33
CommVault Galaxy Topology .....	34
Component Summary .....	36
Current EBS Caveats .....	37
Computer Associates ARCserve.....	38
VERITAS Backup Exec.....	38
VERITAS NetBackup.....	38
Legato NetWorker.....	39
Tivoli Storage Manager.....	39
Learning Check .....	40

## Module 7 — Designing the Multiserver Backup Solution

Objectives .....	1
MB/s to GB/hr Conversion Formulas.....	2
Performing a Needs Analysis .....	3
Creating a Network Topology Map .....	3
Network .....	3
Servers .....	3
Volumes .....	4
Identifying the Backup Window .....	6
Categorizing Data Importance .....	7
Designing a Solution .....	9
Creating Backup Data Domains.....	9
Calculating the Data Volume .....	9
Grouping Data into Jobs .....	9
Selecting the Appropriate Device .....	10
Step 1: Calculating the Compressed Rate.....	10
Step 2: Adjusting the Compressed Rate for Feed Speed and Lights-Out Operations.....	11
Step 3: Determining the Number of Drives Needed to Meet the Backup Window.....	12
Step 4: Determining the Quantity of Retained Media .....	12
Documenting Your Solution .....	13
EBS Site Survey Form.....	14
HP StorageWorks Backup Sizing Tool .....	16
Other Backup Tools.....	18
StorageWorks Enterprise Volume Manager .....	18
StorageWorks Virtual Replicator .....	19
Microsoft Exchange Backup/Restore Solution .....	19
StorageWorks Data Replication Manager.....	20
Using EBS to Back Up a Cluster.....	21
ARCserve2000 in a Cluster Environment.....	21
Backup Exec and ARCserve in Novell NetWare Cluster Environment .....	22
Cluster Environments in the Enterprise Backup Solution with VERITAS NetBackup DataCenter.....	22
Microsoft Cluster Server.....	22
VERITAS Cluster Software for Solaris.....	23
Legato NetWorker in a Cluster Environment .....	24
Microsoft Cluster Server.....	24
Learning Check .....	25

## Module 8 — Performance Tuning

Objectives .....	1
System View .....	2
Primary Storage Performance .....	3
File System Performance .....	4
Measuring Primary Storage Performance .....	4
Backup Server Performance .....	5
Backup Application Setup .....	6
Switch Performance .....	7
Fibre Channel-to-SCSI Bridge Performance .....	7
Tape Drive Performance .....	8
Measuring Tape Drive Performance .....	8
Restore Performance .....	9
Learning Check .....	10

## Module 9 — Diagnostics and Troubleshooting

Objectives .....	1
HP StorageWorks Library and Tape Tools .....	2
Software Features .....	2
Supported Products and Operating Systems .....	3
Application Window Layout .....	3
Using the Device Information Screen .....	4
Using the Firmware Management Screen .....	7
Using the Tests Screen .....	9
Using the Support Ticket Generation Screen .....	12
Management Support .....	14
Insight Manager 7 .....	15
Features .....	15
EBS Support .....	15
Media Robot Utility .....	16
Features .....	16
Benefits .....	16
Limitations of Fibre Channel .....	17
Common Types of Loop Disruption .....	17
Diagnosis and Repair Using Troubleshooting Flowcharts .....	18
Visual and Physical Inspection .....	19
Checking the Host on the FC-SW .....	20
Device Confirmation .....	21
Topology Confirmation .....	23
Checking the Switch .....	24
Switch Management .....	25
Checking the Fibre Channel-to-SCSI Bridge .....	27
Troubleshooting the Modular Data Router .....	28
Troubleshooting the Network Storage Router .....	31
Checking the Library .....	33

Typical Causes of Library Failure .....	34
SCSI IDs .....	34
Library Troubleshooting Strategy.....	35
Troubleshooting the SSL2020TL and TL891 Libraries .....	36
Troubleshooting the MSL5000 Library Family .....	37
Troubleshooting the TL895 Library and ESL9000 Library Family .....	38
Learning Check .....	51

## **Module 10 — HP OpenView Storage Media Operations**

Objectives .....	1
Overview .....	2
Typical Lifecycle of Physical Media .....	2
Impacts and Challenges.....	3
Storage Media Operations Features and Benefits .....	4
Components of Storage Media Operations .....	5
Storage Media Operations Architecture.....	6
Firewall Support.....	7
Installation Requirements .....	8
Storage Media Operations Server Installation .....	8
Storage Media Operations Server GUI .....	9
Storage Media Operations Database .....	10
Protecting the Storage Media Operations Database .....	10
Storage Media Operations Clients .....	12
Data Protection Integration Agent .....	13
Licensing.....	14
Objects and Organization Units.....	15
Daily Operations.....	18
Premount Jobs .....	18
Vaulting Jobs.....	19
Check-Out Requests.....	23
Scratch-Bin Maintenance .....	24



## Introduction

Exponential growth of information technology in the business arena has led to an enormous volume of data stored on enterprise networks. Backing up that data is increasingly important—and increasingly a challenge. Customers need a scalable, high-performance solution to meet present and future data storage needs.

HP has developed an architecture for consolidated protection of business-critical data through a comprehensive backup solution. HP StorageWorks Enterprise Backup Solutions provide centralized backup of multiple servers, at local backup speeds, over a Fibre Channel storage network to tape libraries.

This course is designed to help you learn the information and skills you need to design and implement a plan for effective centralized backup.

In this course, you will learn practical information about designing, implementing, and troubleshooting Enterprise Backup Solutions, including:

- The benefits of Enterprise Backup Solutions.
- The components of a solution.
- How to design, install, and configure a solution.
- How everything works together to provide customers with a backup storage solution.
- How to maintain a solution.
- How to diagnose and troubleshoot a solution.

In addition, you will participate in hands-on lab activities and case studies that simulate situations you might encounter in the field.

## Course Objectives

After completing this course, you should be able to:

- Identify appropriate applications of HP Enterprise Backup Solutions to meet given business needs.
- Design a solution based on customer requirements.
- Properly install the hardware and software to implement the solution.
- Diagnose Enterprise Backup Solution problems and accurately identify corrective actions.

## Prerequisites

The instructor will deliver this course under the assumption that you have attended, in the order listed, the following courses:

- One of the following, delivered by the respective partner:
  - Computer Associates ARCserveIT 6.6 or later Backup Administrator course
  - VERITAS Backup Exec 7.0 or later course
  - Legato Networker 5.5 or later course
  - VERITAS NetBackup 3.4 or later course
  - Tivoli Storage Manager 3.7 or later course
- StorageWorks Full-Line Technical Training
- SANWorks Full-Line Technical Training
- HP ASE certification or equivalent knowledge or experience

HP also strongly recommends that you have a working knowledge of Microsoft Windows NT , Novell NetWare, Tru64 UNIX, or Sun Solaris.

## Module Overview

Module	Topics to be Covered
Backup and Restore Basics	<ul style="list-style-type: none"> <li>■ The Evolution of Backup Methods</li> <li>■ Backup Performance</li> <li>■ Performance Issues</li> <li>■ Backup Strategies</li> <li>■ Rotating Tapes</li> </ul>
Tape Drive Technologies	<ul style="list-style-type: none"> <li>■ AIT Drives and Media</li> <li>■ DLT Drives</li> <li>■ SDLT Drives</li> <li>■ LTO Technology</li> <li>■ Ultrium Generation 1 Drives</li> <li>■ Ultrium Generation 2 Drives</li> <li>■ Ultrium/SDLT Comparison</li> </ul>
MSL5000 and SSL2020 Tape Libraries	<ul style="list-style-type: none"> <li>■ HP StorageWorks Tape Library Nomenclature</li> <li>■ SSL2020 AIT Library</li> <li>■ MSL5000 Library Family Overview</li> </ul>
ESL9000 Tape Libraries	<ul style="list-style-type: none"> <li>■ ESL9000 Family of Tape Libraries</li> <li>■ Library GUI and Operations for the ESL9000 Family</li> <li>■ SCSI Target Guidelines for Tape Libraries</li> <li>■ Fibre Channel LUN Guidelines for Tape Libraries</li> </ul>
Network Storage Routers	<ul style="list-style-type: none"> <li>■ Network Storage Router Overview</li> <li>■ Converting Fibre Channel to SCSI</li> <li>■ Converting SCSI to Fibre Channel</li> <li>■ Network Storage Router M2402</li> <li>■ Network Storage Router N1200</li> <li>■ Embedded Network Storage Routers</li> <li>■ Visual Manager User Interface</li> <li>■ Optimizing Performance on the Network Storage Router</li> </ul>
Software Components of the Enterprise Backup Solution	<ul style="list-style-type: none"> <li>■ EBS Software Components</li> <li>■ Backup Software</li> <li>■ Current EBS Caveats</li> </ul>
Designing the Multiserver Backup Solution	<ul style="list-style-type: none"> <li>■ MB/s to GB/hr Conversion Formulas</li> <li>■ Performing a Needs Analysis</li> <li>■ Designing a Solution</li> <li>■ EBS Site Survey Form</li> <li>■ HP StorageWorks Backup Sizing Tool</li> <li>■ Other Backup Tools</li> <li>■ Using EBS to Back Up a Cluster</li> </ul>
Performance Tuning	<ul style="list-style-type: none"> <li>■ Primary Storage Performance</li> <li>■ Backup Server Performance</li> <li>■ Switch Performance</li> <li>■ Fibre Channel-to-SCSI Bridge Performance</li> <li>■ Restore Performance</li> </ul>

Diagnostics and Troubleshooting	<ul style="list-style-type: none"><li>■ Management Support</li><li>■ Limitations of Fibre Channel</li><li>■ Diagnosis and Repair Using Troubleshooting Flowcharts</li><li>■ Diagnostics</li></ul>
HP OpenView Storage Media Operations	<ul style="list-style-type: none"><li>■ Overview</li><li>■ Installation Requirements</li><li>■ Objects and Organizational Units</li><li>■ Daily Operations</li></ul>

### Objectives

After completing this module, you should be able to:

- List existing backup methods and explain their evolution.
- Explain the importance of hard drive and controller feed speed on backup performance.
- Identify performance tuning issues of designing a secondary storage solution.
- Define backup and restore strategies.
- Identify backup and restore tape rotation schemes.
- Define archiving and Hierarchical Storage Management.

## Evolution of Backup Methods

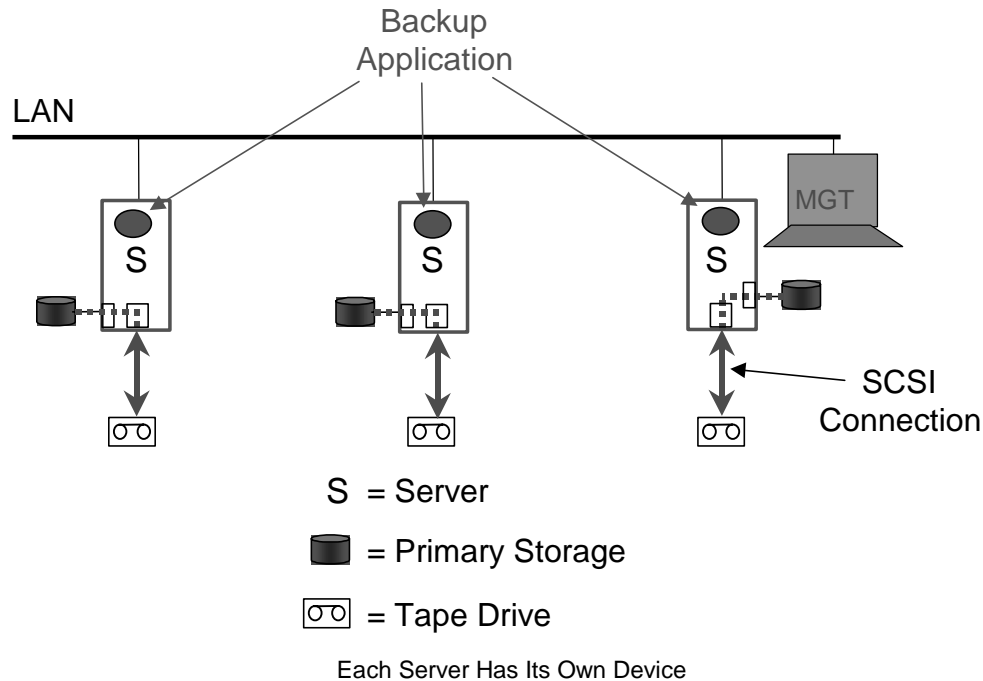
Protecting enterprise data is critical. Backup strategies have evolved to cope with the challenges of larger quantities of data, smaller backup windows, and tighter resource constraints.

Each existing backup solution presents its own advantages and limitations. Performance can be compared along several dimensions, including:

- Resources consumed
- Exposure to human error
- Difficulty in managing backup media
- Backup speed
- LAN congestion during backup

This section traces the evolution of backup methods beginning with basic server backup and continuing with the trends toward centralization, automation, and virtualization of secondary storage.

## Basic Server Backup



This diagram shows a basic server backup environment (also referred to as *local backup*) in which each server connects to its own backup device through a SCSI bus. The operator loads a dedicated instance of the backup software for each server that has a backup requirement. The backup software reads the data from primary storage and writes the data to the backup device.

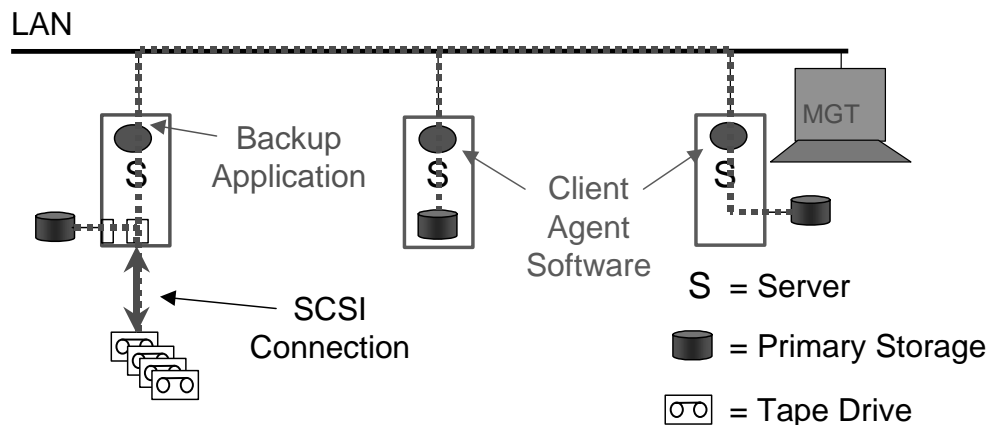
The operator controls the application locally or remotely depending on the remote management capabilities of the application. The storage media for each server is managed locally and manually. In installations where large quantities of servers exist, the management of backup media becomes quite a task.

In this arrangement, the speed of the backup device is usually the bottleneck to backup performance. Unlike centralized backup solutions, backup data and network traffic each travel on dedicated lines. One advantage of this popular backup method is that backups do not consume LAN bandwidth.

### Basic Server Backup Advantages and Disadvantages

Advantages	Disadvantages
Fast	Relatively expensive
Does not consume LAN bandwidth	Must manage each server individually

## Centralized Server Backup



Servers Backed Up Over the Network

With the introduction of client “push” agents, backup devices no longer need to be attached directly to the server in need of a backup; they can be located on a different server attached to the LAN. The backup application runs on the server hosting the backup devices, and client agents push the data over the LAN to the server running the backup application. Media management difficulties decrease with the consolidation into one backup device.

However, in installations where large quantities of servers exist, the LAN becomes the performance bottleneck for backup. Also, the additional traffic on the LAN consumes precious bandwidth that could otherwise be used for business productivity. This backup traffic places new constraints on when backups are performed and the scheduling of server backup windows becomes critical to the smooth operation of the business.

The following table displays maximum and typical LAN speeds, which can be regarded as the upper limits to backup data transfer rates over a LAN.

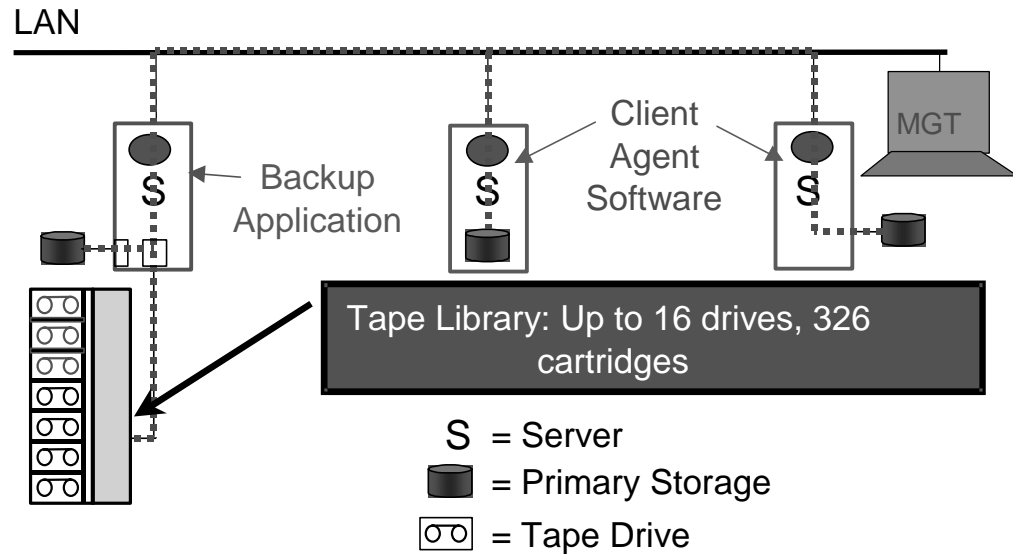
LAN Type	Maximum Speed	Typical Speed
10Base-T	3.6GB/hr	2GB/hr
100Base-T	36GB/hr	15–20GB/hr
Gigabit Ethernet	360GB/hr	36-100GB/hr
FDDI	Similar to 100Base-T	Similar to 100Base-T
Fibre Channel	360GB/hr	280GB/hr

### Centralized Server Backup Advantages and Disadvantages

Advantages	Disadvantages
Centralized management	Consumes LAN bandwidth
	Single point of failure (backup server)



## Automated Centralized Backup across the LAN



Extra Capacity and Automation Using a Tape Library

Using a tape library adds both capacity and automation to further reduce the media management problems. The library can contain up to 16 SDLT drives, each drawing data at a rate of up to 22MB/s with 2:1 data compression. However, with more consolidation and faster backup devices such as the SDLT drive, the LAN bottleneck worsens, even when using a 100Mb/s dedicated backup LAN.

With all data flowing through one server, backup speed is limited by the:


- Capabilities of the LAN.
- I/O capabilities of the server.

Because backups require most of the network bandwidth, they are scheduled during off-peak hours or during scheduled outage windows.

### Automated Centralized Backup across the LAN Advantages and Disadvantages

Advantages	Disadvantages
Centralized management	Consumes LAN bandwidth
Tape automation	Single point of failure (backup server)
High speeds from backup server to tape device	

## LAN-Free Backup




By dedicating storage to a Fibre Channel Storage Area Network (SAN), data traffic between storage devices is offloaded from the LAN. System performance is dramatically increased because data and communications traffic no longer compete for the limited bandwidth on the standard LAN. With a SAN, storage resources can be consolidated and shared by many servers. In addition, the SAN solution can be scaled by cascading Fibre Channel fabric switches.

Industry-leading independent software vendors (ISVs) have developed backup applications that intelligently manage when and where data is backed up in a SAN solution.

Initial cost of Fibre Channel components is the main disadvantage of LAN-free backup. However, storage consolidation and centralized management can reduce the total cost of ownership.

Advantages	Disadvantages
Centralized management	Initial cost of Fibre Channel components
Tape automation	
High performance over dedicated storage area network	
Consolidation of storage resources	
Scalability	

## Off-Host Backup (Non-Disruptive Backup) with Snapshots



Although LAN-free backup provides excellent performance, many customers have strict availability requirements for data. Users may be accessing servers and data on a 24x7 basis. If users are updating files or database records while the backup process is copying data, the integrity of the data copied to tape may be inconsistent.


Off-host backup, also known as non-disruptive backup, mounts a record of the data to be backed up onto a dedicated backup server by means of snapshots. A snapshot is record of the blocks of data at a particular instant. In order to ensure that buffers and cached data are flushed to disk, applications must be quiesced prior to taking the snapshot. The quiescing will result in some downtime, which should be minimal.

Once the snapshot is mounted on the dedicated backup server, data can be backed up. During the backup, data movement is between the backup server, primary storage, and tape library. Therefore, the server where the data was originally mounted can resume with running applications. In addition, off-host backup ensures point-in-time data integrity.

Finally, off-host backups also help provide quick restores. Instead of restoring data from tape, a file can be restored by accessing the snapshot on disk.

Advantages	Disadvantages
Point-in-time integrity of data being backed up	Applications need to be quiesced, resulting in some downtime
Minimal downtime for application servers	If hard drives lost, both original data and snapshot data will be lost
Applications can resume after snapshot	Some overhead on original application server since both backup server and original server accessing same disks
Quick restores from snapshots	

## Off-Host Backup (Non-Disruptive Backup) with Clones



Another type of off-host backup mounts a copy of the data to be backed up onto a dedicated backup server by means of clones. A clone is a physical copy of data at a particular instant. In order to ensure that buffers and cached data are flushed to disk, applications must be quiesced after creating the clone. The quiescing will result in some downtime, which should be minimal. The clone can be split off and mounted on a dedicated backup server.

Once the clone is mounted on the dedicated backup server, data can be backed up. During the backup, data movement is between the backup server, primary storage, and tape library. Therefore, the server where the data was originally mounted can resume with running applications with little added overhead during the backup. In addition, off-host backup ensures point-in-time data integrity.


Finally, off-host backups also help provide quick restores. Instead of restoring data from tape, a file can be restored by accessing the clone on disk.

Advantages	Disadvantages
Point-in-time integrity of data being backed up	Applications need to be quiesced, resulting in some downtime
Minimal downtime for application servers	Creation of clone at 60GB/hr
Applications can resume after clone mounted	
Little overhead from backup process on original application server because clones use different disks	
Can access data on clone if original disks lost	
Quick restores from clones	

## Clones and Snapshot Comparison

Clones	Snapshots
Physical copy	Virtual copy
60GB/hr creation time	Creation time in seconds
Fast restores	Fast restores
No performance impact on original disks during backups	Some performance impact on original disks during backups
Good fault tolerance	Data is lost if disks are lost

## Third-Party Copy (Serverless Backup)



Under many backup methods, the server must be actively involved during a backup. Below is a synopsis of the movement of data during a backup of a server with an external disk subsystem.

1. Server initiates backup.
2. External disk subsystem sends data to server.
3. Server sends data to tape library.

The above backup process consumes resources on the server as well as bandwidth between the server and the storage devices.

The T10 Technical Committee has proposed the EXTENDED COPY command for inclusion in SCSI standards. The EXTENDED COPY command uses a copy manager device to copy data from one set of logical units to another or the same set of logical units.

Instead of actively engaging the server during the backup process, movement of data is managed by a “third-party” device. This third-party device can issue EXTENDED COPY commands to move data between the external disk subsystem and tape library. Data is copied on a block basis, not on a file basis. Below is a synopsis of third-party copy:

1. Snapshot is taken.

2. Server initiates backup by sending list of blocks to be backed up to the third-party device.
3. Third-party device issues commands to external disk subsystem.
4. External disk subsystem sends data to tape library.

Although the server is still involved with the backup process, the third-party device manages the movement of data between disk and tape. This reduces traffic between the server and storage devices and reduces overhead on the server.

However, the server and ISV software will be actively involved in a restore to catalog file structure and to recover files.

HP does not currently implement the third-copy backup.

Advantages	Disadvantages
Less overhead from backup process on server	Server is still involved in backup process
Less traffic over the SAN between the server and storage devices	Server and ISV application are actively involved in restores
	Extended Copy specifications are being standardized
	Current implementations are limited

## Backup Performance

An important aspect of tuning the backup and restore solution is the interface between the backup device and the server. SCSI is an I/O bus widely used in the computer industry to attach peripheral devices such as hard disk and tape drives to network servers.

## Background

One of the original objectives of SCSI technology was to provide device independence within a class of devices. Standard compliant SCSI devices from different manufacturers are capable of attaching to the same host computer without requiring hardware or software changes.

Standardization of the SCSI specification has been ongoing in the industry for many years. In 1986 the American National Standards Institute (ANSI) approved the SCSI specification as a standard defining the mechanical, electrical, and functional requirements for the attachment of SCSI peripheral devices. Over the years, performance improvements in both processors and peripheral devices pushed the original SCSI specification to its limit and drove the need for additional versions of the specification.

Subsequent versions of the specification have been created to increase performance and improve functionality of the interface while maintaining the industry-standard approach.



## SCSI Evolution

### SCSI-1

The maximum data transfer speed for this implementation of SCSI is 2 to 4MB/s (actual average is around 2.5MB/s), using a limited instruction set. Under SCSI-1, all devices use different commands.

### SCSI-2

SCSI-2 (referred to as plain SCSI) is the second-generation SCSI standard. It consists of the basic SCSI-1 standard with many additions and some deletions. Two alternative signaling systems are available when implementing SCSI-2:

- **Single-ended interface** — This is “regular” SCSI and uses the type of conventional signaling that is used on other buses.
- **Differential interface** — The differential SCSI bus minimized the potential bottleneck created by bus length limitations experienced with single-ended SCSI.

These two alternatives are incompatible, resulting in two main groups of SCSI devices and controllers that cannot be mixed on the same bus. It is possible to use special converter hardware to transform a single-ended bus into a differential one (and vice versa).

Single-ended implementations are the most common. They are suitable for internal cabling. Differential interfaces are used externally.

Two SCSI bus widths are defined by the SCSI-2 standard:

- **Narrow** — 8-bit data pathway incorporating a 50-pin connector and cable.
- **Wide** — 16-bit data pathway incorporating a 68-pin connector and cable. *Wide* can indicate 16-bit or 32-bit buses. The 32-bit variant is rarely used, so *Wide* normally means 16-bit.

Two bus speeds are defined in SCSI-2 protocols:

- Regular
- Fast

## SCSI-3

SCSI-3 defines new physical-level transports, IEEE 1394 and Fibre Channel, as a means of transporting SCSI data packets. SCSI-3 defines a new low-voltage differential (LVD) SCSI specification.

LVD SCSI is a technology that combines the advantages of both its predecessors. LVD uses differential signaling techniques instead of single-ended, making the bus more stable. It will support up to 15 devices on one cable and enables the use of external SCSI cabling up to 12m long.

Bus speeds are defined in the SCSI-3 protocols:

- **Ultra** — Transfer rate of 20MHz (also called *Fast-20* or *F20*). Ultra SCSI buses have a maximum transfer rate of 20MB/s for Narrow SCSI or 40MB/s for Wide SCSI.
- **Ultra2** — Transfer rate of 40MHz (also called *Fast-40* or *F40*). Ultra2 SCSI buses use LVD and has a maximum transfer rate of 40MB/s for Narrow SCSI or 80MB/s for Wide SCSI.
- **Ultra3** — Ultra160 SCSI is the generation of high-performance SCSI technology that offers data transfer speeds of up to 160MB/s.
- **Ultra320** — Ultra320 SCSI is the generation of high-performance SCSI technology that offers data transfer speeds of up to 320MB/s.

## Backup Speeds and Feeds

The feed speed is the rate at which data is transferred to a tape drive. Feed speed is dependent on many factors, which are discussed later in this module. The write speed is the rate at which a tape drive transfers data to a tape.

The HP engineering team in Houston quantified the performance of the entire backup solution and its components. Basic speeds and feeds of the solution were tested. In order to achieve optimal backup performance, a 3:1 feed speed to write speed ratio is necessary for tape drives. If the feed speed to write speed ratio is less than 2:1, the tape drive performance may be halved.

## Performance Issues

High-performance enterprise backup solutions have been developed to meet the demand for reliable secondary storage. For these backup solutions to function with the performance their critical role demands, several aspects of their implementation must be optimized, including:

- Configuration of the primary storage system.
- Type of connection to the tape storage system.
- Type and configuration of the tape drives.
- Compression ratio and data verification.
- Block sizes and formatting of tapes.

After these factors have been considered and optimized, you can choose the optimal backup solution.

## Configuration of Primary Storage

The configuration of primary storage has an impact on secondary storage performance.

The type of hard drives used in the primary storage solution and their transfer rate help determine how quickly data can be moved from the drives to the secondary storage system. The faster the transfer rate of the hard drive, the faster the data can be accessed.

The number of spindles in the primary storage system also affects secondary storage performance.

The controller for the primary storage helps determine how quickly data can be retrieved from the system. The following table lists transfer rates for some of the array controllers.

<b>Array Controller</b>	<b>Simultaneous Drive Transfer Channels</b>	<b>Total Channel Transfer Rate (MB/s)</b>
SA 5300	4	528 (maximum)
HSG80 static load balancing	2	110 (average)
HSG80 Secure Path	4	220 (average)
HSV110	4	Up to 2100
XP1024	4	Up to 3700

The array configuration, for example, the RAID level, also affects secondary storage performance. Use of striped sets will allow data to be read from all the drives concurrently. Use of multiple volumes helps overall backup performance by permitting additional simultaneous backup jobs.

## Types of Connection

The type of connection between the servers and clients to be backed up and the secondary storage system affects the backup performance. This connection is typically one of the following:

- **Directly connected SCSI tape device.** Devices connected directly to the server through a SCSI connection are very fast at backing up that server.
- **Network connection between client and backup server.** The LAN bandwidth affects the speed at which data can be transmitted between the client devices and the backup server.
- **Fibre Channel connection between backup server and tape device.** Data transmitted over a Fibre Channel connection to the tape device is very fast.

Typically, a combination of these connections is found in the backup environment.

## Types and Configuration of Tape Drives

The type and configuration of tape drives also affects backup performance. Newer tape technologies typically allow faster data transfer rates.

The way in which the tape drives are configured, for example, in a Redundant Array of Independent Tapes (RAIT), also affects performance.

## Compression Ratio and Data Verification

The optimal backup solution is one that performs backups while other tasks are processing on the system. It must do this with minimal processor use to enable those other tasks to continue processing. This means that the backup solution should have hardware (rather than software) data compression and read-while-write data verification to ensure data integrity without affecting the performance of the system or the backup devices.

*Hardware data compression* means that compressing data for efficient storage on the backup device does not require use of the server's processor, which would slow performance of other transactions on the server.

*Read-while-write data verification* means that when data is recorded to tape, parity checks and other forms of data verification occur concurrently with the writing of the data itself. This enables smooth streaming of data and continuous recording, even when errors are being discovered and corrected.

## Block Sizes and Formatting of Tapes

When formatting tapes, you must set the block size for the tape. For optimal performance, the block size on the tape should match the block size for the primary storage that is being backed up. Formatting the primary storage with the largest block size available will provide the fastest data transfers.

## Choosing the Optimal Backup Solution

The basic considerations for choosing the optimal backup solution are:

- **Speed** — How much speed is required to back up the data within the time allotted for the backup?
- **Capacity** — How much storage capacity will be needed to avoid unnecessary manual media changes during the backup operation? Is an automated tape rotation over the long term a requirement? How much capacity will this require and what level of automation?
- **Hardware and media costs** — The right backup solution can lower administrative costs, but too large of a system adds overhead without additional benefit for the company.
- **Hardware reliability** — A tape drive that must be driven beyond the duty cycle it was designed for to achieve the desired backups will suffer maintenance problems and excessive wear and tear. For the long term, the solution that is properly sized for the job costs less.

**Speed**

$$\frac{\text{DATA (Gigabytes)}}{\text{Backup Window (Hours)}} = \text{Required Transfer Rate (GB/hr)}$$

Required Transfer Rate Formula

Performance needs are determined by dividing the amount of information (in gigabytes) that must be backed up by the size of the backup window (in hours). This simple calculation yields the required performance as an overall transfer rate expressed in GB/hr.

For you to begin working with the required transfer rate formula, the table on the next page shows the following specifications on some of the HP tape drives:

- The maximum capacity of uncompressed data each drive can store on a cassette
- The transfer rates for reading and writing uncompressed data with each drive
- Transfer rates when performing a local backup using each drive in a system
- The estimated time for that system to perform a 100GB backup using each of the drives
- The number of tape cartridges that will have to be loaded and unloaded during the backup operation

By comparing the results of the required performance calculation with the typical backup performance rates listed in this table, the administrator can determine if the required backup performance is achievable. If one of the listed drive types will perform adequately, an appropriate choice would be local, offline backups.

## HP Tape Drive Specifications

Performance Feature	AIT	DLT	SDLT		Ultrium	
Capacity (GB)	50	40/80	110/220	160/320	230	460
Drive's native (uncompressed) capacity (GB)	50	40	110	160	100	200
Transfer rate without data compression (GB/hr)	20	18.6	39.6	57.6	54.0	108
Time for a 100GB backup	300 min	320 min	150 min	104 min	111 min	56 min
Number of tapes for a 100GB backup	2	2	1	1	1	1



### Note

When you need to calculate whether a particular drive will meet your backup needs, use transfer rates less than the published transfer rate. It is also a good practice to calculate backup system requirements based on native (uncompressed) data figures. The 2:1 compression ratio assumed might not be realistic because compressibility is variable in real-world conditions. Zero compression (native data) is a safe assumption.

In the case of a network backup, LAN speed is another factor that can limit backup performance. The following table indicates maximum and typical LAN speeds, which can be regarded as upper limits to backup performance.

## LAN Speeds

LAN Type	Maximum Speed	Typical Speed
10Base-T	3.6GB/hr	2GB/hr
100Base-T	36GB/hr	15–20GB/hr
FDDI	Similar to 100Base-T	Similar to 100Base-T
Gigabit Ethernet	360GB/hr	36–100GB/hr
Fibre Channel	720GB/hr	560GB/hr



## The Backup Window

After identifying backup requirements, determine the *backup window*—the amount of time the company can afford to take servers offline each night for performing backups (if indeed they need to perform an offline backup).

In the following table, four HP servers are chosen as examples to illustrate how the required transfer rate formula can be used to choose a backup tape device.

Server	Amount of Data to Back Up	Backup Window	Suggested Backup Device
ProLiant ML350	25GB	7 hours	20/40GB DLT
ProLiant ML370	30GB	2 hours	Ultrium 230
ProLiant DL740	80GB	2.5 hours	MSL5060L1
ProLiant DL760	160GB	4.3 hours	MSL5060L1

### Example

You need to provide local backup support for the ProLiant ML370 Server shown in the table to have 30GB of data to back up. Your backup window is four hours. What performance is required to deliver this kind of backup within the available time?

$$30\text{GB} \div 4 \text{ hours} = 7.5\text{GB/hr performance required}$$

Three HP tape drives can transfer the data quickly enough to complete the backup in the allowable backup window: the HP AIT 35GB Tape Drive (10GB/hr), the HP AIT 50GB Tape Drive (20GB/hr), and the HP 40/80GB DLT Drive (18.6GB/hr). All three are good choices, but the DLT drive might offer a smoother migration path toward bigger storage systems.

## Capacity



Unattended backups can significantly reduce these administrative costs. Deploying sufficient backup capacity and automation, if needed, to make “lights-out” backups possible is often a significant cost saver for the long term.

For unattended backups, the tape drive and media must meet the backup requirements in terms of overall storage capacity and backup performance.

As the preceding table indicated, both performance and capacity vary with the type of drive technology selected. The newer SDLT technology has the highest capacities and fastest transfer rates.

- **DAT** drives are usually the appropriate choice for servers with 4GB to 20GB capacity. An autoloader for DAT tapes is available from HP. Autoloaders reduce administrative costs by using a robotic mechanism to load and unload tapes.
- **AIT** drives have the high capacity and reliability of DLT drives to back up servers up to 50GB without compression. With greater data transfer rates than the 35/70GB DLT, they are an attractive upgrade from high-end DAT drives and a good alternative to DLT for lower cost solutions. AIT also offers a tape array and library option.
- **DLT** drives use simultaneous multichannel and multihead read/write technology to achieve capacities up to 40GB without compression. A DLT drive is the appropriate high-end backup solution for larger servers, especially those with storage of 40GB or more, owing to a wide range of tape array, mini-library, and large library options.
- **SDLT** (Super DLT) drives offer backward-compatibility with DLT tapes and increase storage capacity and transfer rate by an order of magnitude. Storage capacity starts at 110GB without compression. The SDLT drive is appropriate for enterprise-class servers, owing to its capacity and transfer rates, as well as improved reliability over DLT.
- **Ultrium** drives offer outstanding reliability, excellent capacity, and data rate matching technology to optimize performance. Cartridge memory helps improve media management and reduce wear.

## Hardware and Media Costs



The drive performance and drive capacity must be evaluated in view of current drive and media costs and the budget of the company. A realistic budget for the purchase of the appropriate drives and media is important. If too little is budgeted for drives, the company might incur increased labor costs or possibly higher hardware costs in the future.

An all-too-common example of this involves a system administrator who must work overtime every evening just to change tapes on a drive without an autoloader.

Alternatively, the company might end up with a drive that is simply too slow to complete the backup during the backup window. This could result in degradation in server performance outside the backup window. It will certainly result in premature wear of the tape device.

A drive with 10 times the required capacity might have the advantage of low cost per gigabyte, but the initial purchase price would be difficult to justify.

Because the costs of backup hardware and media constantly change, the preceding graphics illustrate only the relative prices of different drives and media types.

## Hardware Reliability

The reliability of a backup device is related directly to its duty cycle (the number of hours per day that the device is in use). For example, if a tape drive designed for 1GB backups is being used to back up a 10GB server, the results on that drive will be:

- Premature aging
- Reliability problems
- A need for accelerated preventive maintenance—especially head cleaning



### Note

This applies primarily to older DAT drives. DLT read and write heads make very little actual contact with the tape and require much less cleaning. In DAT drives, the tape is wrapped around the read-and-write heads; therefore, the drives need additional head cleaning. However, state-of-the-art DAT and AIT technology such as that provided by current HP products automates much of the cleaning, reducing the need for regular maintenance.

The best method for building hardware reliability into a backup strategy is to ensure that the backup hardware is matched to the servers. The following table relates various servers with the appropriate tape drives. If the company needs special, partial backups in addition to the routine backups, it might be appropriate to select the next larger drive size.

Server Capacity	Recommended Drive Format	Recommended HP Drive
<20GB	DAT	20/40GB
<50GB	AIT	50GB
>50GB	SDLT	110/220GB or 160/320
	Ultrium	Ultrium 230 or Ultrium 460

Another factor to consider is media life. The following table lists the typical life for the various media discussed in this module.

Media Type	Rated Passes	Rated Backups <sup>a</sup>
DAT	2,000	300
AIT (50GB)	20,000	10,000
SDLT I	1,000,000	17,850
Ultrium 200GB	1,000,000	20,000
Ultrium 400GB	1,000,000	20,000

---

**Note**

*a* Under optimum environmental conditions of 50 percent relative humidity and 22 degrees C. Assumes that drive is streaming and that the entire tape capacity is used for each backup.

---

## Backup Strategies

Backup strategies are determined by asking the following questions:

- Is the backup full or partial?
- Is the backup image- or file-based?
- How big are the backup and restore windows?
- How often should backups take place?

### Is the Backup Full or Partial?

Trade-offs must be considered when choosing between full and partial backups:

- Amount of data that will be backed up each day
- Time required to complete the backup session
- Number of tapes required to recover data to the server after a disaster

### Full Backup

A full backup is a complete backup of the entire server or PC client hard drive. For a server, this includes all volumes, directories, and files. For a PC client, this includes all drives, directories, and files.

Full backups can be further categorized as *normal* or *copy* backups:

- **Normal backup** — Backs up files and resets the archive bit. The archive bit is used to determine if the file has been backed up or not.
- **Copy backup** — Is similar to a normal backup except that it does not reset the archive bit.

### Partial Backup

A partial backup can be any of the following:

- Incremental
- Differential
- User-defined

## **Incremental**

An incremental backup copies all files that were changed after the last backup, regardless of what kind of backup it was. This type of backup is used when each revision of a file must be maintained. If the same tapes are used for consecutive incremental backups, the newer versions of backed-up files are not allowed to overwrite earlier versions. Rather, the newer files are usually appended to the backup medium. This backup strategy is recommended if different files are accessed each day.

Typically, backup programs reset the archive bit following each incremental backup. An incremental backup works only on the log files and therefore only when circular logging is disabled. Like a normal backup, an incremental backup also purges log files after backing them up, providing yet another way to rid log files from your drive without compromising recoverability.

To restore an incremental backup, you must return to your last normal backup set, which contains your database files. Restore those database files, restore every incremental backup set made after the normal backup, and then start the service. Do not start the service until you have restored all the backup sets; otherwise, any logs restored after the backup set will not be played forward.

## **Differential**

A differential backup copies all files that were changed since the last complete backup. Differential backups are useful when it is important to have the latest version of each file. If the same tapes are used for consecutive differential backups, the newer versions of backed-up files are often allowed to overwrite older versions of the same file on the tape. Typically, backup programs do not reset the archive bit after a differential backup; the archive bit remains enabled until the next complete backup. This backup strategy is recommended if the same files are accessed each day.

Because a differential backup also works on log files, you must disable circular logging to use it. Unlike an incremental backup, however, a differential backup does not delete the log files.

To restore a differential backup set, return to the last normal backup and restore your differential backup set that contains any log files generated after your last normal backup. As with an incremental backup, do not start the service until you have restored all the backup sets.



## User-Defined

A user-defined backup copies a user-defined set of files. Often this is a special backup requested by a group of employees on a mission-critical project.



---

### Note

When backing up databases such as Microsoft SQL, most third-party backup applications include a special module that works in tandem with the database software. Before you can activate the backup of any or all your SQL databases, you must have administrative rights or be a database owner for the database you want to back up.

---

## Is the Backup Image- or File-Based?

Two types of software backup techniques are generally available, depending on the version of software used:



- **Image-based** — In image-based backups, the backup program views the disk drive or file system as a single unit and takes a snapshot of it, backing it up in one session. Data is transferred to tape in blocks without regard to filenames or objects. Image-based is usually the fastest method of backup and restore for an entire file system. In environments with large numbers of small files present, image-based backup is the method of choice. However, some image-based backups can be restored only as an entire unit, meaning that the entire backup must be restored to obtain a single file.



- **File-based** — File-based backups are made on a file-by-file basis, retaining the information needed to retrieve a single file from the backup tape and restore it to the correct location on the system. The file-based method usually takes much longer than the image-based method because the full path information is saved with each file. In environments with large numbers of small files, the overhead for the operating system of reading the file allocation table many times can reduce backup throughputs by as much as 90%. The file-based method's advantage is that it enables random access to the files for faster retrieval of individual files.

## How Big Are the Backup and Restore Windows?

Administrators typically perform backups when user demands on the server are at their lowest. Ideally, this time period, known as the *backup window*, is when user access can be restricted or the server shut down. As more companies move to 24 x 7 operation, backup windows are shrinking.

For many companies with worldwide operations accessing their servers, no clear backup window exists. The system administrator must determine how to get the backup done without affecting the productivity of users or seriously degrading network performance.

The length of the restore window is increasingly important. The restore window helps determine how quickly you must be able to restore files after a disaster.

As more critical data is placed on networks, businesses cannot afford to have their data offline for long periods of time. After a disaster, data must be restored and available as quickly as possible.

## **How Often Should Backups Take Place?**

Backups must be performed regularly. The actual frequency of backups will be determined by considerations such as:

- The acceptable amount of work that could be lost, if any, in the event of a catastrophic failure.
- The allowable downtime for recovery from this failure.
- The volume of update transactions that normally take place.
- Administrative and resource costs of making backups.

In determining the frequency of backups, an effective backup strategy should also incorporate redundancy for security.

## Rotating Tapes

Because a company's data is valuable and difficult to replace, backups should be viewed not as something to be done when needed, but as a disciplined, scheduled regimen, carefully designed to address issues that include:

- Redundancy
- Storage space limitations
- Speed of recovery after a disaster
- Efficient use of media

The information gathered in the backup needs assessment enables you to design a regular backup regimen for the client. Such a regimen involves establishing a schedule for backups, including decisions about when to use partial or full backups and whether partial backups should be incremental or differential.

Although tapes are relatively inexpensive, users will reuse tapes periodically. When tapes are reused, the data that currently exist on the tapes will likely be overwritten. Tape rotation schemes are used to ensure the integrity of the data being restored.

Tape rotation schemes must fit a customer's business needs. Tape rotation schemes may need to match a business's cycles. Some businesses will need to archive certain data for preservation while keeping other data on a weekly basis.

Tape rotation schemes are configured through the backup application.

## Grandfather-Father-Son Tape Rotation

The Grandfather-Father-Son (GFS) tape rotation scheme is the most commonly used and requires a weekly backup capacity of at least double the server storage capacity. It uses three levels of backup to provide redundancy and security.

Among other things, this scheme allows for different levels of data retention. The system administrator can select which generation of tapes to store temporarily and which to archive.

The name *GFS* refers to three levels of backup:

- Monthly *grandfathers*
- Weekly *fathers*
- Daily *sons*

Typically, the system administrator performs a full backup every Monday (father) and does incremental backups on Tuesdays, Wednesdays, and Thursdays (sons). The administrator performs another full backup at the end of the week (father) and yet another at the end of the month (grandfather).

The GFS tape rotation scheme is intended to ensure that a company can always restore lost data within a day of a disaster. Businesses dedicated to helping get a company back online after a disaster can make computer equipment, tape drives, phone equipment, and so forth available to a company very quickly.

The tapes containing the weekly and monthly backups are usually stored in a location away from the site of the server. To help reduce media costs, many companies reuse older, weekly backup tapes.



### Note

Never overwrite a recently used backup tape. If a hard drive crash should occur during the backup, not only will all the data on the drive have been lost, but the tape will no longer be useful for a complete restoration.

---

## GFS Example

The implementation of this GFS tape rotation scheme varies from company to company. One engineering and construction firm has developed the plan shown in the following generic calendar. In addition to GFS backups, this plan incorporates special, user-defined backups.

### Typical Monthly GFS Backup Schedule

	Monday	Tuesday	Wednesday	Thursday	Friday
<b>Week 1</b>	<b>FULL</b> store off-site (father)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>FULL</b> store off-site (father)
	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site
<b>Week 2</b>	<b>FULL</b> store off-site (father)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>FULL</b> store off-site (father)
	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site
<b>Week 3</b>	<b>FULL</b> store off-site (father)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>FULL</b> store off-site (father)
	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site
<b>Week 4</b>	<b>FULL</b> store off-site (father)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>INCR.</b> (son)	<b>End of Month</b> <b>FULL</b> store off-site (grandfather)
	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site	<b>SPECIAL</b>	<b>SPECIAL</b> store off-site

The company that developed this plan has an engineering staff who works on several design-and-development projects simultaneously. One of these projects is mission-critical.

The project management team has requested that the network administrators perform a full, special backup of their applications and data every day. (These are the daily special backups shown in the calendar.) These managers have also requested that their backup tapes be stored off-site every Monday, Wednesday, and Friday. This special backup requirement is in addition to the routine, company-wide GFS backups.

No weekend activity is shown in the calendar. However:

- If any mission-critical development is done over a weekend, the network administrator is notified to schedule one or two additional special backups.
- If the last day in the month is a Saturday or Sunday, the grandfather backup takes place on that day.

If any work is done on any of the other engineering projects on Saturday or Sunday, the normal Monday father backup will cover that work.

Notice that the applications and data of the mission-critical project are backed up twice every Monday and Friday, the same days when father backups (which are complete backups) take place. The company management does not see this as an issue. To them, this redundancy is additional insurance against a catastrophic loss of important project information.

## Other Tape Storage Strategies

### Archiving

When reusing tapes in a rotation scheme, previous data on tapes eventually are overwritten. However, customers may need to store certain data on a long-term basis. The process of storing a permanent copy of data is known as archiving.

Archiving can be user-initiated or automated. After data is archived, the data can then remain on disk or be deleted from disk to free space.

### Hierarchical Storage Management

Another scheme to free disk space is Hierarchical Storage Management (HSM). Data is migrated from disk to optical media and then to tape or from disk to tape. The policies for migrating data from fastest media to slowest media are configured through the backup application. Such policies may be based on amount of free disk space and when a file was last accessed.

After a file has been migrated from one media to another, a marker is used in place of the file. When a user accesses the file, the file is automatically retrieved from the appropriate media.

## Learning Check

1. The backup solution should have hardware data compression rather than software data compression for optimal performance.  
☐ True  
☐ False
2. Write the required transfer rate formula.  
.....
3. List the three adverse results expected when using a 1GB tape drive to back up a 10GB server.  
.....  
.....  
.....
4. The recent generation of Wide-Ultra SCSI-3 increases the bandwidth of the storage subsystem.  
☐ True  
☐ False
5. What potential issues exist when each server has its own backup device?
  - a. Congestion on the LAN
  - b. Tape cartridge schemes
  - c. Management capabilities
  - d. Connectivity to the server



6. Centralized server backups involve potential LAN congestion.
  - ☐ True
  - ☐ False
7. In a high-speed backup solution, what does the term *feed speed* mean?
 

.....

.....
8. How is a differential backup different from an incremental backup?
 

.....

.....

.....
9. When referring to GFS tape rotation schemes, the “F” stands for \_\_\_\_\_ fathers.
  - a. Monthly
  - b. Weekly
  - c. Daily
  - d. Hourly
10. Which device manages the movement of data during the backup process for third-party copy (serverless backup)
  - a. The server
  - b. The client
  - c. The third-party device
  - d. The tape library
11. Off-host backup (non-disruptive backup) with snapshots provide zero downtime before and during the backup process.
  - ☐ True
  - ☐ False
12. Which of the following is NOT a disadvantage of LAN-free backup?
  - a. Must manage each server individually
  - b. Consumes LAN bandwidth
  - c. Relatively low performance
  - d. All of the above




### **Objectives**

After completing this module, you should be able to:

- Describe the features of AIT, DLT, SDLT, and LTO drives


## Tape Drive Technologies

### Helical Scan Recording




Helical scan recording has both tape and head movement. The read and write heads are attached to a rotating drum which is tilted at an angle to the tape. The tape is pulled out of the cassette and wrapped around the rotating heads. Data is recorded in diagonal stripes. This approach results in high capacity and relatively slow tape movement. Helical scan recording is used in DAT and AIT drives.

### AIT Drives and Media



Advanced Intelligent Tape (AIT) drives offer an exceptional combination of data integrity, speed, price, reliability, and storage capacity. AIT drives use a self-cleaning head, eliminating the need for periodic cleaning.





AIT tape uses the durable Advanced Metal Evaporated (AME) tape, which includes a 100 percent pure cobalt magnetic layer and relies on a Diamond Like Carbon (DLC) protective coating for increased strength and durability. AME media is highly resistant to oxidation and moisture, which minimizes drive contamination and results in less frequent cleaning of read/write heads.

AIT tapes have a Memory-In-Cassette(MIC) chip to store and retrieve a tape's data log, history, and directory information before the tape starts to move.

### AIT 50GB Drive Specifications

Category	Specifications
Storage capacity per drive	50GB native data, 100GB compressed 2:1
Data transfer rate per drive	Maximum: 20GB/hr native data, 40GB/hr compressed 2:1 Typical: 16GB/hr native data, 36GB/hr compressed 2:1
Media	755ft (230 m) Advanced-Metal Evaporated (AME) tape with memory in cartridge (MIC)
Tape format	8mm helical scan
Head life	60,000 tape contact hours
Cartridge life	20,000 passes (10,000 full backups)

## Linear Recording



Linear (DLT)

Linear recording divides the tape into parallel, horizontal tracks and records data by moving the tape at high speeds past heads that remain stationary.

### DLT Drives

Digital Linear Tape (DLT) drives utilize a linear serpentine recording pattern. The drive writes along one set of tracks from the beginning of the tape until the end of the tape. After reaching end of the tape, the tape moves in the opposite direction and another set of tracks is used for writing data. DLT 35/70GB drives and DLT 40/80GB drives write data in angled pattern, allowing higher density per cartridge. This angled pattern is known as Symmetric Phase Recording. For every backup, a direct track access (DTA) directory is recorded at the beginning of the tape. The directory includes the location of each file on the tape, including track numbers.

DLT tape uses metal particulate media (MP-2). Metal particulate is prepared by double-coating a thick nonmagnetic underlayer with an extremely thin magnetic overlayer saturated with tiny metal particles.

Within the DLT drive, six precision rollers provide a gentle path for the tape. The recorded side of the tape never touches the rollers.


### 35/70GB DLT Drive Specifications

Category	Specifications
Storage capacity per drive	35GB native data, 70GB compressed 2:1
Data transfer rate per drive	Maximum: 18GB/hr native data, 36GB/hr compressed 2:1 Typical: 14.4GB/hr native data, 28.8GB/hr compressed 2:1
Media	DLT Tape IV for full capacity and performance; backward compatible with legacy DLT tape formats
Tape format	DLT 7000 format: 208 track Symmetric Phase Recording
Head life	30,000 hours
Cartridge life	1,000,000 passes (19,230 full backups)

### 40/80GB DLT Drive Specifications

Category	Specifications
Storage capacity per drive	40GB native data, 80GB compressed 2:1
Data transfer rate per drive	Maximum: 21.6GB/hr native data, 43.2GB/hr compressed 2:1 Typical: 16.2GB/hr native data, 32.4GB/hr compressed 2:1
Media	DLT Tape IV for full capacity and performance; backward compatible with legacy DLT tape formats
Tape format	DLT 8000 format: 208 track Symmetric Phase Recording
Head life	50,000 hours
Cartridge life	1,000,000 passes (15,000 full backups)

## SDLT Drives




SuperDLT (SDLT) technology increases tape capacity and transfer rate by an order of magnitude over DLT technology. In addition, each HP SDLT drive has a second head to provide backward read-compatibility for DLT tapes. SDLT technological advancements ensure reliable, durable, economical tape storage. SDLT drives have either capacities of 110/220GB (SDLT 220) or SDLT 160/320 (SDLT 320). Both the SDLT 220 and SDLT 320 utilize similar optical and magnetic technologies.

Under normal circumstances, as tape moves through the roller guides in the drive, the tape wobbles slightly. The servo system in a tape drive is the scheme that is used to position the head in the center of the track. For DLT drives, the heads have a tolerance for the slight wobbling due to the width of the tracks. However, the width of the tracks also limits the capacity of the tapes.


For SDLT drives, the servo system consists of an optical servo track on the unused back side of the tape and a laser optics system to track. SDLT drives have two fewer roller guides than DLT drives. This reduces the tension applied to the media as the tape moves, thereby improving tape life.

Magneto-Resistive Cluster (MRC) heads deliver higher data transfer rates and greater data density than traditional heads of equal size. The MRC heads are densely packed using advanced thin-film technology. In addition, MRC heads are less susceptible to temperature and humidity conditions.





One of the leading causes for failures in DLT tapes and drives is the unbuckling of the drive leader from the tape leader. In SDLT drives, the Positive Engagement Tape Leader Buckling Mechanism uses a solid metal pin that is attached to the drive leader to link with molded clips that are permanently attached to the tape leader inside the cartridge. The Positive Leader Link design makes the buckling of SDLT media a highly reliable mechanical process.



Advanced Metal Powder (AMP) media uses durable metal powder technology for recording at very high densities. The backside of the AMP media contains a specially formulated back coating to receive the optical servo tracks, allowing the entire data-bearing side of the media to be used for recording data and eliminating the need for pre-formatting.

## 110/220GB SDLT Drive Specifications

Category	Specifications
Storage capacity per drive	110GB native data, 220GB compressed 2:1
Data transfer rate per drive	Maximum: 39.6GB/hr native data, 79.2GB/hr compressed 2:1
Media	SuperDLT Tape 1 for full capacity and performance; backward read-compatible with DLT IV tape formats
Tape format	SuperDLT Tape 1: 448 track Serial Serpentine Recording
Head life	30,000 hours
Cartridge life	1,000,000 passes (19,230 full backups)

## 160/320GB SDLT Drive Specifications

Category	Specifications
Storage capacity per drive	160GB native data, 320GB compressed 2:1
Data transfer rate per drive	Maximum: 57.6GB/hr native data, 115.2GB/hr compressed 2:1
Media	SuperDLT Tape 1 for full capacity and performance; backward read-compatible with DLT IV tape formats
Tape format	SuperDLT Tape 1: 448 track Serial Serpentine Recording
Cartridge life	1,000,000 passes (17,850 full backups)

## **LTO Technology**

Linear Tape-Open (LTO) technology was conceived by and continues to be developed jointly by HP, IBM and Seagate.

LTO technology is an "open format" technology, which means that users will have multiple sources of product and media.


The open nature of LTO technology also provides a means of enabling compatibility between different vendors' offerings as well competitive pricing. In contrast, the SDLT and AIT drives have single sources.

The three LTO founding companies set up three processes — specification, licensing, and compliance. The specification process is the design vehicle for the tape formats. The licensing process extends to any company that wishes to develop an LTO-compliant product access to the design requirements and the right to submit a developed product for verification testing. The compliance process describes the steps through which a product must pass before it can receive certification.

The Ultrium format is the high capacity implementation of LTO technology. Ultrium technology is a single-reel design that uses half-inch tape. The design objectives called for maximizing the amount of tape surface area (hence higher capacity) while maintaining a small form-factor tape cartridge. The Ultrium format is a scaleable format with a 4-generation roadmap. It has no restrictions of older formats while incorporating best of breed features from current storage products.

## Ultrium Generation 1 Drives

### Ultrium-1 Drive Design Overview




The mechanism is split into two sections by the chassis. The tape handling, and tape motion areas are designed not to have a cooling airflow. The main PCA is mounted underneath the chassis, with vents in the front panel and housing to ensure maximum airflow is drawn over the PCA. This ensures the maximum cooling for the main PCA and reel motors and the lowest contaminant build up in the tape handling area. Significant testing has been performed on the mechanism to ensure resistance to dust contamination. The drive has an extremely high tolerance to dust contamination.

The mechanism load motor drives the powered load and unloads of the cartridge as well as the leader pin grab mechanics. With a powered cartridge load/unload and leader grab it is important to have a method to extract the cartridge should the drive power fail. Although not shown in this illustration one of the gears at the right side of the mechanism can be turned as a thumbwheel allowing manual mechanical load, unload, and leader pin grab/ungrab. This could be useful in the case of a power or PCA failure event with a cartridge loaded. The supply reel motor can be manually turned using a Torx-10 screwdriver through the access points in the mechanism casing.


The HP Ultrium-1 mechanism has no speed encoder as found in DLT tape devices with the tape speed being calculated both from hall sensors in the reel motors, and also from the servo signal on tape. The servo controller processes these to ensure the tape is maintained at the required speed.

## Leader Pin and Tape Path


With a single reel tape device, the interlock between the takeup reel and the media is the most important part of the drive for reliability. As specified in the Ultrium format, with the there is a leader pin attached to the end of the media. The media is wrapped around the steel pin, and fixed in place by an elastomer and C-Clip. The leader pin can easily be checked on a piece of media by sliding open the cartridge door, and checking the pin is seated in the cartridge as shown in the diagram below.



Within the drive there is a leader block attached to the take-up leader that is engaged directly onto the leader pin. The leader block mechanically engages the leader pin, with a spring retainer to maintain the grab in case of power failure / tension loss. The grab actuator (shown in the illustration below) will only pull the leader block through the drive and onto the take-up reel when opto-mechanical sensors at the top and bottom edges of the leader block confirm that the leader pin is in place. A third sensor tells the drive when the actuator has rotated to the correct position for leader pin grab.




Once successfully engaged, the leader block (light blue) is drawn into the mechanism, and guided along the load tracks (yellow), away from the head and onto the take up reel.



The leader block is designed so that it completes the spindle of the take up reel. The illustration below shows how the leader block fits into place, when pulled into the drive.




Previous tape technologies have fixed very precisely the position of the tape during motion (called the tape path), and used a stepper motor to move the head across the tape. The HP Ultrium-1 drive uses an active servo system to rapidly track the head to any movement of the tape. Because of this the two guide rollers do not need to 'pinch' the tape into an exact position but allow the tape to 'float' causing less wear to the media. The illustration below shows the tape path and head actuator assembly.



The left hand guide roller also makes up part of the leader pin grab assembly. When the leader block is loaded into the drive, rails in the mechanism guide it away from the head assembly preventing any damage occurring. As can be seen above when the block passes the second guide roller the tape then pulls flat across the MR head and actuator assembly.

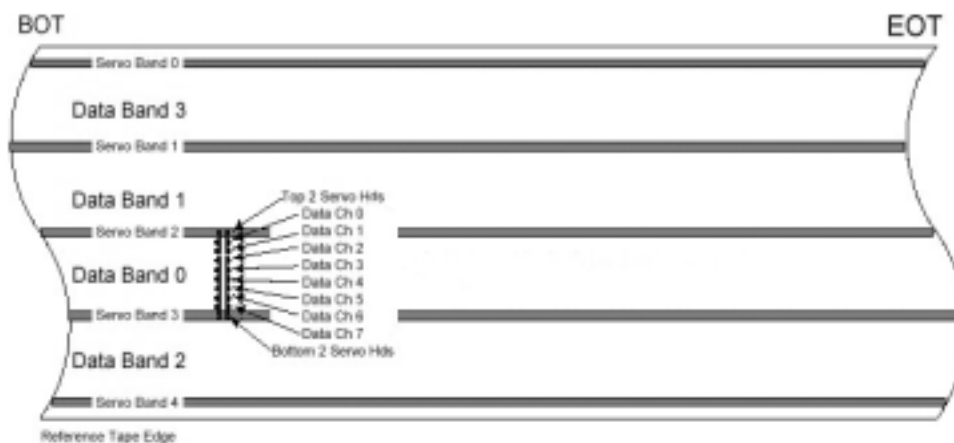
A reliability concern with any linear tape device is media and/or roller wear distorting the tape path. As already mentioned, the tape is not pinched by the rollers, significantly reducing edge wear. In addition the rollers are coated with wear resistant tungsten carbide to prevent any channels being eroded into the surface. HP uses a SEM to check that roller wear is within acceptable limits.



## Tape Format


The 4 data bands are shown between the 5 servo bands in the illustration below. Each consecutive data band will be filled before the drive will start to write into a different data band. The drive will write the data bands in order 0 thru 3. If a tape is damaged it is normally the edges of the tape which become affected. This is why the first two data bands are in the center of the tape, improving data reliability.

The servo heads and servo tracks allow the head to be positioned at 12 unique positions within each data band. With an 8 channel head this gives us 96 tracks that can be written within a data band. As we have four data bands we get a total of 384 tracks written onto the tape.



Within the data band the head serpentine, completing a full travel to the end of the tape before changing the head position and writing a track in the reverse direction back down the tape.

The illustration below shows the order in which the drive will write data tracks within the data band (note that only 3 of the data channels are shown in the image).

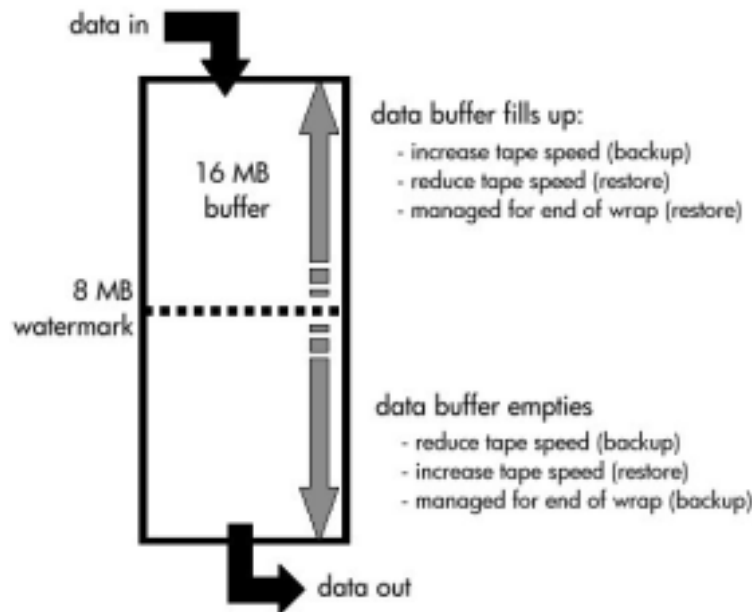




## Adaptive Tape Speed

Adaptive tape speed is a new feature specific to HP Ultrium-1 tape drives. This feature may also be referred to as ‘Data Rate Matching’. Both these titles refer to the same functionality within the HP Ultrium-1 drive.

The HP Ultrium-1 drive has a 16 MB data buffer. In order to ‘stream’ the drive successfully in either the read or write mode, we have to maintain data in this buffer. The benefit of this is that we can prolong or even eliminate repositions




caused by the buffer emptying of data.

The illustration shown above explains how the ATS works. Essentially the drive can vary its transfer rate from its maximum speed of 15 MB/s (native) to a minimum of 6 MB/s (native). By varying the speed of the tape it attempts to maintain 8 MB of data in the data buffer. In steady state operation the data buffer is maintained at half full (8MB). During a backup, as the buffer starts to fill, the tape speed is increased. This prevents us from having to disable host transfer if the data is coming in too quickly, helping us match the data rate of the server.

If the buffer starts to empty then the tape speed is reduced. To achieve optimum performance, the data buffer is emptied as the end of wrap is approached. While the tape reverses and the head repositions the buffer is filled by the host. The buffer is then emptied to watermark level when writing starts on the next wrap. For data restore the operation is reversed at the end of the wrap.

By reducing the number of repositions that occur during normal operation for the drive we can improve the overall reliability.

## Cartridge Memory (LTO-CM)



Previous tape technologies have always stored important information about the tape, calibration, and contents in what was called the ‘system area’ at the beginning of a tape. Whenever the tape was used the drive would always have to read & write this same section of tape. As a result this section of tape would often become heavily worn, and become a common failure area on tape.

With Cartridge Memory, the tape identity, type (cleaning tape or data tape), usage and contents are stored in an on-board memory located in the cartridge rather than on tape. This eliminates this heavy wear area at the beginning of the tape.

The Ultrium CM is a non-contact proximity RF chip which again improves reliability by removing physical contact. The CM chip as shown in the photograph on the left is mounted at 45 degrees within ALL Ultrium cartridges. This is to allow it to be read both by a tape drive from underneath the cartridge, and a library picker from in front of the cartridge.

The reader, as can be seen in the photograph on the right is an inductive coil built into the drive PCA. The passive RF interface gives a range of up to 20mm from the reader to the cartridge memory chip. The cartridge memory stores 4 KB of information in its memory organized as 128 x 32 byte blocks. Data is transferred to and from the CM in 32 byte blocks.

Examples of useful information contained in the CM:

- Unique Cartridge Identifier
- Cartridge manufacturer & batch id
- Cartridge manufacturer’s servo writer & head assembly
- Any tape alert flags set with this cartridge
- Thread count
- Tape directory
- Usage data from the last 4 loads of this cartridge (cyclic buffer).

A slimmed down version of the cartridge memory contents is also contained at the beginning of the tape (called the FID). If for whatever the reason the CM fails, it is still possible to recover data from that tape.

## Ultrium 230 Drive Specifications

Category	Specifications
Storage capacity per drive	100GB native data, 200GB compressed 2:1
Data transfer rate per drive	Maximum: 54GB/hr native data, 108GB/hr compressed 2:1
Media	LTO Ultrium Tape
Tape format	LTO Ultrium: 384 track Serial Serpentine Recording
Cartridge life	1,000,000 passes

## Ultrium Generation 2 Drives

### Ultrium-2 Capacity and Performance

Ultrium-2 builds upon Ultrium-1 technology while increasing the native capacity of the tape cartridge as well as improving performance. The Ultrium 460 drive increases performance with a native sustained throughput of up to 30MB/s. The Ultrium-2 format is similar to the Ultrium-1 format but doubles the native capacity of a tape cartridge to 200GB by increasing the number of data tracks to 512 and increasing the linear data density along the tracks by 33 percent. The Ultrium-2 drive can read and write Ultrium-1 format data cartridges. However, an Ultrium-1 drive cannot read nor write Ultrium-2 format data cartridges.

### Ultrium-2 Read Channel

The Ultrium-1 drive uses a peak detect method of recovering the digital data from the analogue, signal and although well capable at this data density is not sufficient for Ultrium-2 data densities. The read channel of the HP StorageWorks Ultrium 460 uses Partial Response Maximum Likelihood (PRML) technology to recover the data signal from the tape. This well established technology is essential in order to recover data at the densities used in the Ultrium 2 format. PRML has been used to increase the capacity of disk drives for many years.


### Ultrium-2 SCSI Interface

The Ultrium 460 drive uses a SCSI-3 interface commonly referred to as Ultra-160 SCSI. This interface is low voltage differential (LVD) and capable of transferring data bursts at 160MB/s. In order to optimize performance, the SCSI interface at the host or Fibre Channel to SCSI router must be capable of Ultra-160 SCSI as well.

## Cartridge Presence Sensor

A new cartridge presence sensor has been added to the Ultrium 460 tape mechanism. In the previous Ultrium 230 tape drive a cartridge was only detected when it was pushed into the mechanism. This is advantageous for tape drives used in automation as it removes the need for the 'picker' mechanism to push in the cartridge. The tape cartridge can be placed in the slot and then the drive can be instructed via the ACI port to 'load' the cartridge. This can be used to increase the reliability of tape automation mechanisms. The sensor is located at the front of the drive on the cartridge memory reader assembly. This is a great benefit in automation and can identify mechanical faults to the library controller.

## Ultrium-2 Tape Mechanism



The Ultrium 460 tape mechanism is based on the Ultrium 230 mechanism but has several major modifications. The tape path has been changed with new zirconium nitride coated rollers and a new take-up spool with a reversed tape direction. These improvements reduce tape wander at the higher speeds of the Ultrium 2 tape and produce a better tape 'pack'. (Tape pack is the name for tape wound onto a spool) This is beneficial to tape life and reliability. The above shows the new tape path in the Ultrium 2 drive. The improved tape path also enables 'retuning' of the servo system necessary for the increased data density of the Ultrium 2 format. The external mounting of the tape mechanism chassis has been improved with new flexible mounts, which provide better vibration isolation. The Ultrium 460 drive has a better shock & vibration specification.

## Ultrium 460 Drive Specifications

Category	Specifications
Storage capacity per drive	200GB native data, 400GB compressed 2:1
Data transfer rate per drive	Maximum: 108GB/hr native data, 216GB/hr compressed 2:1

Media	LTO Ultrium-2 Tape
Tape format	LTO Ultrium-2: 512 track Serial Serpentine Recording
Head life	30,000 hours
Cartridge life	1,000,000 passes

## Ultrium/SDLT Comparison

<b>Ultrium</b>	<b>SDLT</b>
3 major suppliers	Single source supplier
Dual redundant servo	Servo not redundant
Critical information on media and cartridge memory	Critical information on media
2 rollers only; reduces wear	4 rollers
Drive leader engages media leader pin, not media	Drive leader engages media
Not compatible with DLT	Backward read compatible with DLT IV
Drive not constrained by previous technology	Requires presence of additional DLT read head
Active head cleaning	
Adaptive tape speed	
Provides manual unload mechanism	

## Learning Check

1. What is the purpose of the Memory-In-Cassette for AIT tapes?  
.....  
.....  
.....
2. What information does the directory track access contain? Where is the directory track access located?  
.....  
.....  
.....
3. What is the purpose of a servo system?  
.....  
.....  
.....
4. What type of servo system does the SDLT drive use?  
.....  
.....  
.....
5. What is adaptive tape speed?  
.....  
.....  
.....

### Objectives

After completing this module, you should be able to:

- Describe the features of the SSL2020TL AIT Library
- Describe the features of the MSL5000 family of tape libraries

## How the Tape Libraries Work

### Transport Assembly

The tape libraries use a robotic device called the *transport assembly* to move the tapes from the storage slots to a drive and back again. The robotics installed in the tape library perform a number of tasks that free the network administrator from routine searches.

The robotics:

- Store media in selectable storage slots.
- Load and eject media to and from the drives.
- Communicate with the host and execute requested move instructions.
- Report element status and error information to the host.

The robotics use “sense of touch” with motion feedback systems and a controller that monitors expected and actual force for move operations. The drives load inserted media and give the host read/write access to the data.

### Automation

Tape cartridges in the tape library are labeled with bar codes, which are read by the robotics’ bar code reader. Working together, the tape library and backup software enable customers to catalog media and develop backup strategies without manual intervention. Locations of files and scheduling are managed by the backup applications. With the proper hardware and backup applications, multiple servers can share the same tape library.



## HP StorageWorks Tape Libraries

### Tape Libraries

HP supports the following tape library options in SAN environments:

- StorageWorks SSL2020TL AIT Library
- StorageWorks MSL5026 DLX Library
- StorageWorks MSL5026 SDLT 320 Library
- StorageWorks MSL5026SL Library
- StorageWorks MSL5052SL Library
- StorageWorks MSL5060 Ultrium 230 Library
- StorageWorks ESL9198SL SDLT Library
- StorageWorks ESL9322S2 SDLT 320 Library
- StorageWorks ESL9322L1 Ultrium 230 Library
- StorageWorks ESL9595SL SDLT Library
- StorageWorks ESL9595S2 SDLT 320 Library
- StorageWorks ESL9595L1 Ultrium 230 Library

### HP StorageWorks Tape Library Nomenclature

The names of the HP StorageWorks tape libraries provide an indication to the capacity and type of tape drives as follows:

- SSL2xyz – small storage library
- MSL5xyz – medium storage library
- ESL9xyz – enterprise storage library
- xyz – maximum possible media slots
- DLX – LVD 40/80 DLT drive
- TL – LVD AIT 50/100 drive
- SL – SDLT 220 drive
- S2 – SDLT 320 drive
- L1 – Ultrium-1 drive

All the libraries employ sophisticated robotics to automate tape-changing functions and enable “lights-out” backups of hundreds (or even thousands) of gigabytes of data. The HP library mechanisms place and remove tape cartridges with minimum contact to surfaces through a precision-grip cartridge handling system that emulates the human hand.

Competing libraries on the market, including the StorageTek libraries, slide the tape cartridges across guides that cause the outside of the tapes to develop wear patterns on the cartridges. The worn cartridges release small dust particles that can

contaminate the tape drives or other cartridges, causing damage to tapes and possible library downtime. Reliability is a key to success, especially when multiple servers rely on fewer tape devices.

## SSL2020TL AIT Library




The SSL2020TL AIT Library is offered as either a stand-alone tabletop unit or a scalable/stackable rack-mount unit with a removable 19-cartridge magazine, an integrated bar code reader, and two HP 50GB AIT Drives. The library offers 1TB (native) of storage. This library also offers scalability by stacking modules up to five high to give the customer up to 10TB of storage within a 20U space.

The AIT drives require an LVD SCSI connection. Therefore, the tape controller must use a LVD module.


### SSL2020TL AIT Library Specifications

Category	Specifications
Maximum data transfer rate with one drive	20GB/hr native data 40GB/hr with 2:1 compression
Maximum data transfer rate with two drives	40GB/hr native data 80GB/hr with 2:1 compression
Storage capacity	1TB native data 2TB with 2:1 compression
Mean time between failures (MTBF)	250,000 hours (30% duty cycle)
Media	170 to and 230-meter AME tape with MIC
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with HP Insight Manager for remote monitoring and administration
Options	Additional AIT drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## SSL2020TL AIT Library Components



## MSL5000 Library Family Overview



The MSL5000 tape library family provides a high performance mission-critical automated mid-range backup and restore solution. Some key features include

- High availability
  - True hot-plug tape drives; SCSI bus quiescence not required
  - Redundant hot-plug power supplies with MSL5052/5060
- Multi-unit scalability for capacity expansion and investment protection
- Easy to use touch screen GUI for management
- Mail slot for rapid load/unload
- cPCI I/O slots for future addition of cards
- Designed to accommodate drive upgrades for investment protection

## MSL5026/5030 Library Family




The MSL5026/30 Library family, available in a tabletop or a rack mount form, utilizes one or two hot-pluggable tape drives and two removable magazines. The library offers multi-unit scalability of up to seven units for future growth.

The four members of the MSL5026/5030 Library family are:

- MSL5030L1 (LVD 100/200GB Ultrium 230 drives)
- MSL5026S2 (LVD 160/320GB SDLT drives)
- MSL5026SL (LVD 110/220GB SDLT drives)
- MSL5026DLX (LVD 40/80GB DLT drives)

The MSL5026/5030 Library family has the following specifications:

- One or two hot-pluggable tape drives in a single library
- Up to 26 or 30 data cartridges, using two magazines
- Wide Ultra 2 Low Voltage Differential (LVD) SCSI interface
- Compact 5U form factor with a single module
- Field-scalable modular design that can be configured with up to eight modules
- 2 cPCI slots for future card upgrades
- Single power supply



### MSL5030 Ultrium 230 Specifications

Category	Specifications
Maximum data transfer rate with one drive	54GB/hr native data 108GB/hr with 2:1 compression
Maximum data transfer rate with two drives	108GB/hr native data 216 GB/hr with 2:1 compression
Storage capacity	3.0TB native data 6.0TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	LTO Ultrium cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with Insight Manager for remote monitoring and administration
Options	Additional Ultrium 230 Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

### MSL5026 SDLT320 Specifications

Category	Specifications
Maximum data transfer rate with one drive	57.6GB/hr native data 115.2GB/hr with 2:1 compression

Maximum data transfer rate with two drives	115.2GB/hr native data 230.4 GB/hr with 2:1 compression
Storage capacity	4.16TB native data 8.32TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	SDLT cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with Insight Manager for remote monitoring and administration
Options	Additional SDLT Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## MSL5026SL Specifications

Category	Specifications
Maximum data transfer rate with one drive	39.6GB/hr native data 79.2GB/hr with 2:1 compression
Maximum data transfer rate with two drives	79.2GB/hr native data 178.4GB/hr with 2:1 compression
Storage capacity	2.86TB native data 5.72TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	SDLT cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with HP Insight Manager for remote monitoring and administration
Options	Additional SDLT Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only



## MSL5026DLX Specifications

Category	Specifications
Maximum data transfer rate with one drive	21.6GB/hr native data 43.2GB/hr with 2:1 compression
Maximum data transfer rate with two drives	43.2GB/hr native data 86.4GB/hr with 2:1 compression
Storage capacity	1.04TB native data 2.08TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	DLT IV cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with HP Insight Manager for remote monitoring and administration
Options	Additional DLT Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## MSL5052/5060 Library Family



The MSL5052/5060 Library family, available in a tabletop or a rack mount form, utilizes up to four hot-pluggable tape drives and holds up to a total of 52 or 60 tapes in four removable magazines. The library offers multi-unit scalability of up to four units for future growth.

The MSL5052/5060 Library family has the following members:

- MSL5060L1 (LVD 100/200GB Ultrium 230 drives)
- MSL5052S2 (LVD 160/320GB SDLT drives)
- MSL5052SL (LVD 110/220GB SDLT drives)

The MSL5052/5060 Library family has the following specifications:

- Up to four hot-pluggable tape drives in a single library
- Up to 52 or 60 data cartridges, using four magazines
- Wide Ultra 2 Low Voltage Differential (LVD) SCSI interface
- 10U form factor with a single module
- Field-scalable modular design that can be configured with up to four modules
- 5 cPCI slots for future card upgrades

### MSL5060L1 Specifications

Category	Specifications
Maximum data transfer rate with two drives	108GB/hr native data 216GB/hr with 2:1 compression

Maximum data transfer rate with four drives	216GB/hr native data 432GB/hr with 2:1 compression
Storage capacity	6.0TB native data 12TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	LTO Ultrium cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with Insight Manager for remote monitoring and administration
Options	Additional Ultrium 230 Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

### MSL5052 SDLT 320 Specifications


Category	Specifications
Maximum data transfer rate with two drives	115.2GB/hr native data 230.4GB/hr with 2:1 compression
Maximum data transfer rate with four drives	230.4GB/hr native data 460.8GB/hr with 2:1 compression
Storage capacity	8.32TB native data 16.64TB with 2:1 compression
SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	SDLT cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with Insight Manager for remote monitoring and administration
Options	Additional SDLT Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

### MSL5052SL Specifications

Category	Specifications
Maximum data transfer rate with two drives	79.2GB/hr native data 178.4GB/hr with 2:1 compression
Maximum data transfer rate with four drives	178.4GB/hr native data 356.8GB/hr with 2:1 compression
Storage capacity	5.72TB native data 11.84TB with 2:1 compression

SCSI Interface	Wide Ultra2 Low Voltage Differential
MSBF	1,000,000 swaps
MTBF	250,000 hrs (100% duty cycles)
Media	SDLT cartridges
Status indicators	Front panel activity lights and 4-line, 20-character display; compatible with HP Insight Manager for remote monitoring and administration
Options	Additional SDLT Drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only


### MSL5026/30 Top View



The MSL5026/30 Library contains the following main components as shown in the accompanying photo with the top covers removed:

- Front panel
- Rear panel
- DLT/SDLT/LTO tape drives
- Magazines
- Power supply
- Robotics
- Library controller board

### MSL5052/60 Top View



The MSL5052/60 Library contains the following main components as shown in the accompanying photo with the top covers removed:


- Front panel
- Rear panel
- DLT/SDLT/LTO tape drives
- Magazines
- Power supply
- Robotics
- Library controller board
- Front vertical axis assembly
- Rear vertical axis assembly

**Note:** The MSL5052/MSL5060 Library has a lower level (not shown) that accommodates:


- Two additional tape drives (purchased separately)
- Two additional magazines
- One additional power supply

The MSL5052/MSL5060 Tape Library also contains two (front and rear) vertical axis assemblies that enable the robot to access both the lower level and upper level magazines.

## MSL5000 Rear View



MSL5026/30 Rear View



MSL5052/5060 Rear View

The rear panel provides access to the:


- Power supply
- Tape drives

- Card cage
- PCI add-in card slots
- Library controller board
- Cable connections


Some models of the library ship with an optional embedded Fibre Channel board in the slot next to the library controller board.



## MSL5000 Magazines



## MSL5026 Magazines



## MSL5060 Magazines

The MSL5000 Series Library contains two or four removable tape cartridge magazines. The magazine doors on the front panel of the library provide easy access to remove the magazines in order to insert and remove tape cartridges. The magazines slide in and out of the library chassis on the magazine guides which assure precise positioning for the library robotics.


The left magazines include a mail slot, which is accessible when that magazine slot door is open. Pivoting forward, this mail slot feature lets you insert or remove a single tape without removing the entire magazine.

The MSL5052/MSL5060 Library has an added feature that lets you restrict access to the mail slot when the magazine door is opened. This capability is controlled by GUI touch screen settings.


For configuration purposes:

- Fixed slots are identified as storage elements.
- The mail slot is identified as an import/export element.

## MSL5000 Power Supply



MSL5026/30 Power Supply




MSL5052/60 Power Supply

The power supply includes the following:

- **On/Off power supply switch** manually controls the supply of AC power to the module. Library power is normally controlled from the GUI touch screen and the manual On/Off power supply switch.
- **AC power inlet** supplies AC to the library using a single IEC-compatible socket on the rear panel, which can be connected to any properly grounded outlet with an AC power cord.
- **Power LED** illuminates on the rear panel after touching the GUI touch screen with the power on.

- **Power supply release latch** is located to the right of the Power On LED on the rear panel.
- **Power supply bracket** is used to prevent the power supply from being removed accidentally.
- **PTM Interface Connector** is used to connect the master library to the pass-through mechanism (PTM) on a multimodule configuration.

## MSL5000 Shuttle Assembly (Robotics)



The library robotics consist of a cartridge shuttle, brake, motor hardware, motor drives, and other support electronics.

These robotics are capable of picking and placing tapes throughout a 180 degree arc by moving back and forth along a track that is part stationary and part rotating. For example, the robot can access tapes in the left magazine, rotate to access the tape drive, and rotate again to access tapes in the right magazine.

The MSL5052/MSL5060 Library has vertical axis assemblies that enable the robot to access tapes in the upper as well as the lower level of the library.


The robot will also have access to the optional pass-through mechanism (PTM).

### Bar Code Reader and Labels

The cartridge shuttle assembly includes a mounted bar code reader for scanning tape cartridges installed in the magazines and tape drives. It reads bar code labels attached to each cartridge, and maintains the bar code data in memory as part of the library system map. This allows backup application software packages to automatically perform backups and restores without having to access the data on each tape.

Both a full bar code reader scan and a physical scan are conducted each time the library is initially powered up or each time a tape magazine is exchanged.

## MSL5000 Controller Board




The library controller board contains a single microprocessor and associated logic devices to control robotics operations and manage overall library functions. The microprocessor also manages the SCSI interface between the library and the host system.

The library controller board is installed in a card cage at the rear of the library and can be serviced without requiring special tools.

The photograph above illustrates the library controller board and its J ports.

- J1 — VHDCI library SCSI connector
- J3 — VHDCI library SCSI connector
- J4 — Card cage backplane connector
- J5 — 10-BaseT Ethernet connector
- J6 — RS-232 trace connector
- J8 — Card cage backplane connector

## MSL5000 Pass-Through Mechanism



MSL5000 Series Library can consist of stacked libraries, using either the same model or a combination of different models, to form a scalable multimodule library configuration. Through the use of a rear-mounted pass-through mechanism (PTM), all libraries in a multimodule library configuration can operate together as a single virtual library system.

There are two PTM upgrade kits available:

- The original HP branded PTM
  - Part Number 231824-B21
  - Used with HP libraries using DLT or SDLT drives.
  - Includes a seven-port router/hub (seven port connectors and one Wide Area Network (WAN) connector)
- A new HP branded PTM
  - Part Number 303050-001
  - Used with HP libraries (LTO compatible) using SDLT or LTO drives.
  - Includes an eight-port router/hub (eight port connectors and one WAN connector)

Stacked libraries are interconnected through their rear panel Ethernet connections and a supplied router, which is mounted to the storage cabinet. The router also provides an additional connector when libraries are combined in their maximum stacked height.

The maximum number of stacked libraries allowed is:

- Four MSL5052/MSL5060 Libraries.


- Seven MSL5026 Libraries with the original seven-port router.
- Eight MSL5026/MSL5030 Libraries with the eight-port router.
- A combination of libraries, not exceeding 40 Units (40U) in total stacked height.

A multimodule library configuration appears to the host computer system and library control software as a single library. For multimodule software applications, the top library becomes the primary master unit and all other lower libraries are treated as slave units. The PTM will continue to function each time a slave library is physically removed from the storage cabinet configuration during normal library operations.

### Mixed Media

It is recommended that the MSL5000 Series Library in a multimodule configuration use the same drive type. If you want to mix media, check with your application software vendor to manage partitioning by tape technology. Mixed media using LTO Ultrium cartridges is not supported on libraries with firmware levels of 4.04 or earlier.

### Pass-Through Mechanism Cabling



To install the PTM:

1. Hold the PTM unit by the handle while supporting the PTM unit from the bottom.
2. Carefully insert the PTM base unit into the top library, making sure the alignment pins are fully engaged.



3. If extensions were used, engage the extension pins in the lower libraries.  
**Note:** If the lower pins are difficult to align, push the library slightly forward and then pull it back towards the PTM extension.
4. Using the PTM thumbscrews that are positioned at the top and bottom of each library, secure the PTM base unit and extenders to the libraries.  
**Note:** If you still have difficulty aligning a module, do the following:
  - Slightly loosen the rear rail bracket screws for the bottom library.
  - Reposition the library.
  - Secure the library to the PTM base unit.
  - Retighten the rear rail bracket screws.

Once you have installed the PTM, you are now ready to cable the PTM and router.



---

**Caution**

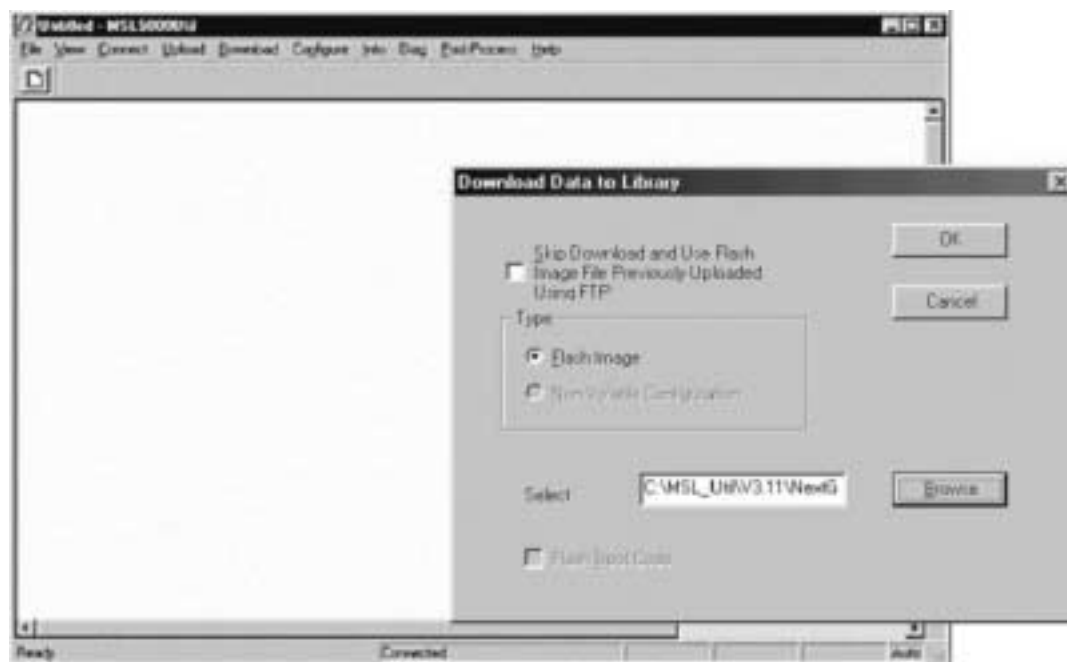
Do not apply power to any library at this time. Connecting power to a library at this point can damage the library and the PTM motor controller board.

---

#### Multimodule Configuration (Master and Slaves)

1. Connect the PTM motor cable to the master library and the PTM base unit.
2. Connect the 10-BaseT connector cable to the master library controller board and to the router.
3. Connect the 10-BaseT connector cable to the slave library controller board and to the router.
4. Repeat step 3 for each additional slave library.

## MSL5000 Firmware Updates — MSL5000 Util



The MSL5000 Series Library can be flashed using one of the following tools:

- HP StorageWorks Library and Tape Tools
- MSL5000Util Diagnostic Utility
- Remote Management Interface

To upgrade the firmware using the MSL5000Util Diagnostic Utility, you will need:

- MSL5000Util.exe
- Cable assembly, serial RJ11 to DB9F

---

### Note

Use the appropriate special or vendor provided cable.

---

- PC with available serial port

The firmware flashing procedures are as follows:

1. Install the MSL5000util.exe on the PC.
2. Copy the firmware image (bmfxxx.img) into the same directory as the MSL5000Util.exe.
3. Connect the library diagnostic port to the PC using the supplied cable assembly.
4. Start the MSL5000Util.exe.  
**Note:** The initial screen, as shown to the right, will appear.
5. Click **Connect**.

6. Select the PC serial port to which the cable assembly is connected.
7. Click **OK**.
8. Click **Download** > **Browse** and navigate to the directory with the firmware image bmfxxx.img.
9. Click **OK**. The library firmware automatically updates.  
**Note:** If the flash process does not complete, select a lower baud rate option, available under the Connect window (Step 5), and try again.
10. Cycle power on the library and verify that the firmware level has been updated.

### MSL5000 Firmware Updates — Remote Management Interface



To upgrade the firmware using the Remote Management Interface, you will need one of the following:


- A LAN connection to the library
- An Ethernet connection to the World Wide Web (WWW).

Once you have established a LAN connection to the library or an Ethernet connection to the Web, the firmware flashing procedures are as follows:

1. Login to the Remote Management Interface.  
**Note:** For information on setting up and using the Remote Management Interface see the *hp StorageWorks MSL5000 series library user guide*.
2. Click **Functions**.
3. Scroll down to **Library Flash Operation**.
4. Select **Flash Library from Remote File** or **Flash Library from Local File**.  
**Note:** If you Flash Library from Local File, the \*.bin file must be copied to the /upload directory, which is located on the library at the IP Address/upload.
5. Click **Start the Flash**. The library firmware automatically updates.
6. The library will reboot.

7. Verify that the firmware level has been updated successfully.

### **Tape Drive Upgrades**



The MSL5000 library family is designed for simple drive upgrades for investment protection. The following are steps required for drive upgrades:

1. Verify newer type of drive supported for library
2. Update library firmware
3. Replace drive

## MSL5000 Management GUI

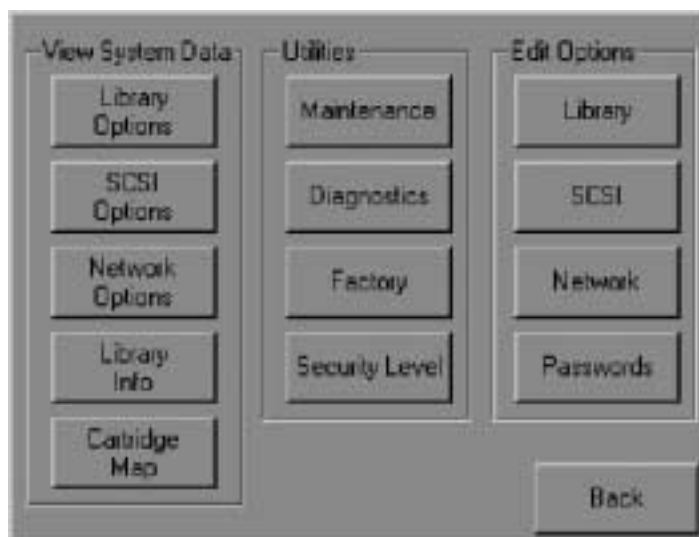


The front panel of the MSL5000 library family provides a convenient means for managing the tape library. Touching the option buttons on this screen gives you the ability to view information and operate the library.

Following are the available options on the main menu for the MSL5000 Series Library:

- Technical Support Information (HP)
- Mail Slot Access
- Magazine Access
- Move Media
- LCD Contrast Controls (Up and Down Arrows)
- Menu
- Online
- Status
- Power


## Management GUI > Menu



Touching **Menu** on the library Status screen displays a screen with three distinct areas:

- View System Data - lets you view the library data.
- Utilities - lets you maintain, diagnose, and secure the library.
- Edit Options - lets you set library, SCSI, and network options.

## Management GUI > Menu > Library Options



View Library Options:

Drive and Slot Numbering:	Zero Based
Library Mode:	Random
Sequential Mode:	Unavailable
LCD Contrast Adjust:	16
Mail Slot Mode:	Mail Slot Enabled

▲ ▼ Back

View Library Options:

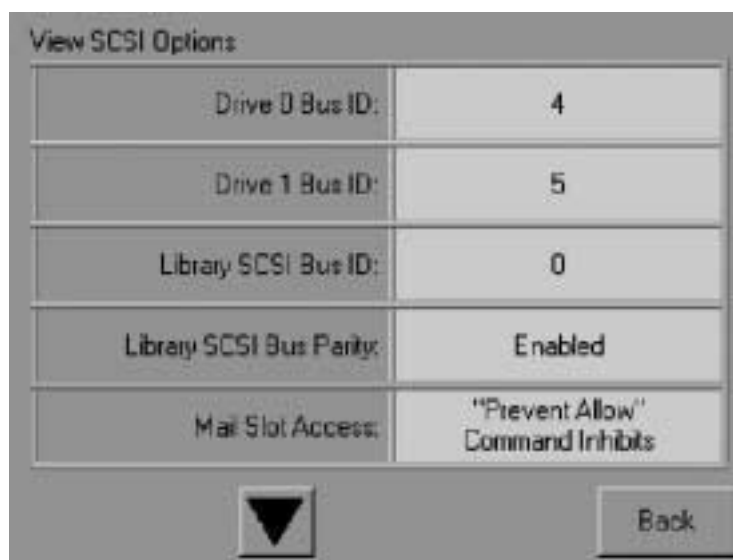
Barcode Label Size:	8 Chars
Barcode Label Alignment:	Left Align
Barcode Label Check Digit:	Disabled
Barcode Reader:	Retries Enabled
Module Configuration:	Standalone

▲ Back

To navigate the library options:

1. From the Menu screen, touch **Library** in the Edit Options area.
2. Use the navigation buttons to progress through the screens.
3. To edit, touch the desired library option button and follow the instructions.

## Management GUI > Menu > SCSI Options



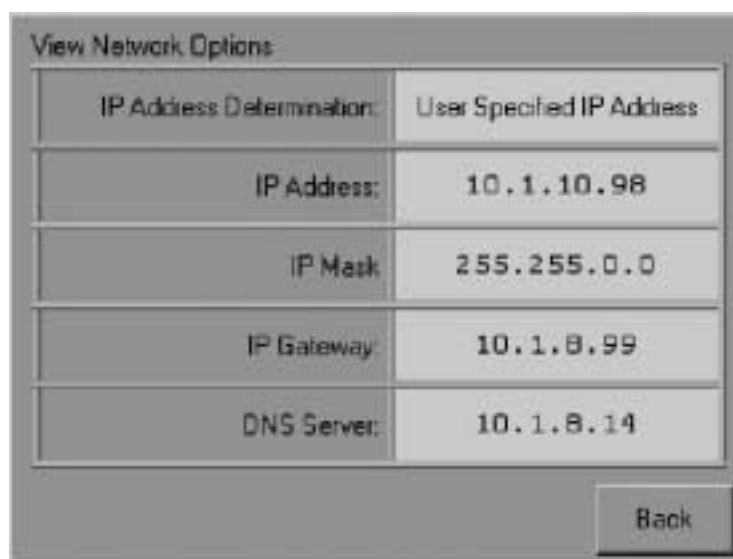
View SCSI Options	
Drive 0 Bus ID:	4
Drive 1 Bus ID:	5
Library SCSI Bus ID:	0
Library SCSI Bus Parity:	Enabled
Mail Slot Access:	"Prevent Allow" Command Inhibits

▼ Back

**Selecting** this option lets you view the current SCSI options. There are several screens associated with SCSI Options. The first of several screens is displayed to above. To view the SCSI options:

1. From the Menu screen, touch **SCSI Option** in the View System Data area.
2. Use the navigation buttons to progress through the screens.

## Management GUI > Menu > Network Options




View Network Options	
IP Address Determination:	User Specified IP Address
IP Address:	10.1.10.98
IP Mask:	255.255.0.0
IP Gateway:	10.1.8.99
DNS Server:	10.1.8.14

Back

The network options allows the user to assign an IP address to the library for multi-unit configurations.




## Management GUI > Menu > Library Info



The Library Info screen displays:

- Firmware revision
- Local IP address
- Ethernet address
- Serial Number

## Management GUI > Menu > Cartridge Map



Selecting this option provides you with a layout of the slots and cartridges that are currently being used in the library. Barcode labels are displayed if enabled and in use.

- Slots that are blank have no cartridges.

- Slots that are black have cartridges without barcode labels.

---

**Note**

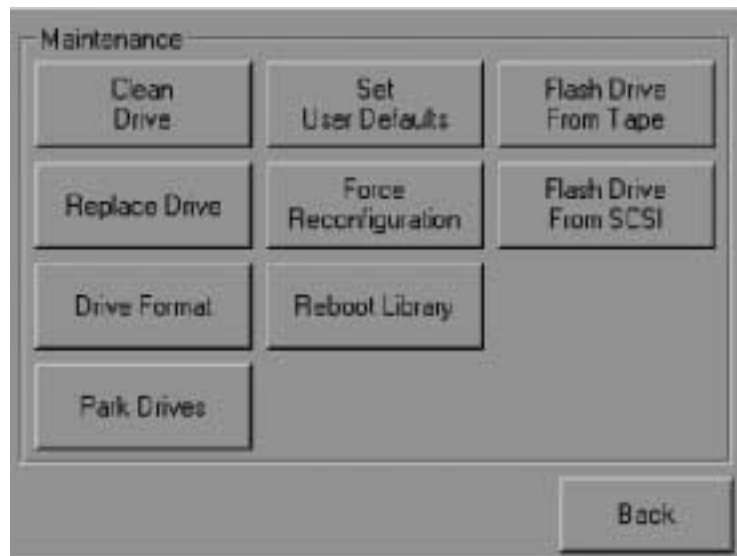
When the robot inventories the slots, it makes two passes:

- One to identify cartridges with barcode labels
  - Another to determine if the unlabeled slots are empty or full
- 

To display the layout:

1. From the Menu screen, touch **Cartridge Map** in the View System Data area.
2. Use the navigation buttons to progress through the screens.

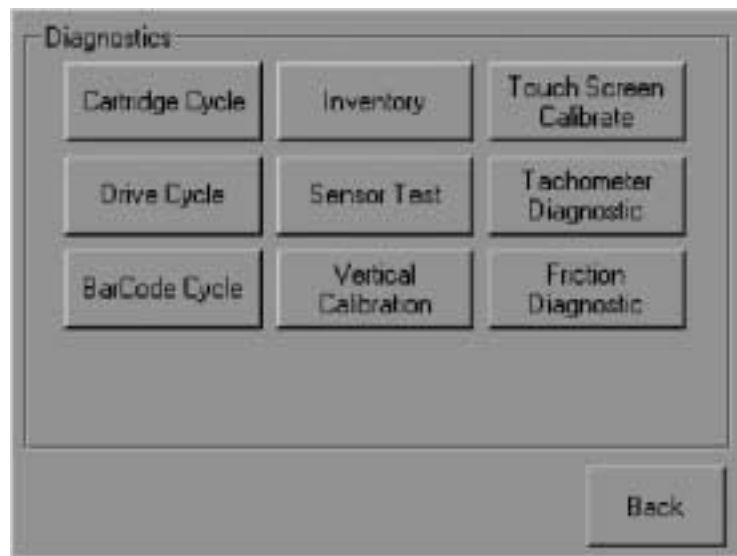
Management GUI > Menu > Maintenance



Selecting this option lets you replace or perform system maintenance on components. Following are the maintenance options:

- Clean Drive
- Replace Drive
- Drive Format
- Park Drives
- Set User Defaults
- Force Reconfiguration
- Reboot Library
- Flash Drive From Tape
- Flash Drive From SCSI

## Management GUI &gt; Menu &gt; Diagnostics



Selecting this option lets you perform system diagnostics and calibration on available components. Available options are dependent on the specified user level. Following are the diagnostics options:

- Cartridge Cycle
- Drive Cycle
- BarCode Cycle
- Inventory
- Sensor Test
- Vertical Calibration
- Touch Screen Calibrate
- Tachometer Diagnostic
- Friction Diagnostic

## Learning Check

1. Which type of SCSI interface does the SSL2020 require?  
.....
2. How many power supplies do the MSL5052/5060 have? How many power supplies do the MSL5026/5030 have?  
.....  
.....
3. What is the advantage of using the mail slot over removing the magazine for importing and exporting tapes?  
.....  
.....  
.....
4. In a multi-unit configuration, what is the purpose of the 10-BaseT hub?  
.....  
.....
5. List three ways to upgrade the firmware of the MSL5000 library.  
.....  
.....  
.....

### Objectives

After completing this module, you should be able to:

- Describe the features of the ESL9000 family of tape libraries

## ESL9000 Overview

The ESL9000 family of tape libraries provides highly reliable backup and restore solutions for mission critical data storage requirements of high-end enterprise customers, providing component level redundancy, high availability, and enterprise level capacity for direct SCSI and SAN environments. Key features of the ESL9000 family include:

- Hot-pluggable, redundant power supplies and fans
- Hot-pluggable tape drives
- Multi-unit scalable for capacity expansion
- Easy-to-use touch screen for management
- Load port for easy import/export of tapes

## ESL9000 Family of Tape Libraries



**ESL9198  
SDLT220  
Library**

198 slots



**ESL9322  
Ultrium 1  
& SDLT320  
Library**

222/322  
slots



**ESL9595  
SDLT220 &  
SDLT320 &  
Ultrium 1  
Library**

400/500/  
595 slots

## HP StorageWorks Tape Library Nomenclature

The names of the HP StorageWorks tape libraries provide an indication to the capacity and type of tape drives as follows:

- MSL5xyz – medium storage library
- ESL9xyz – enterprise storage library
- xyz – maximum possible media slots
- SL – SDLT 220 drive
- S2 – SDLT 320 drive
- L1 – Ultrium-1 drive
- L2 – Ultrium-2 drive

## ESL9000 Family

The ESL9000 family of tape libraries includes:

- ESL9198SL
- ESL9322L1
- ESL9322S2
- ESL9595L1
- ESL9595S2
- ESL9595SL

The ESL9198 and ESL9322 libraries are multi-unit scalable to five units. The ESL9595 libraries are multi-unit scalable to four units.


---

**Note**

Multi-unit scalability is not supported for ESL9000 libraries with SDLT320 drives at initial release.


---

## ESL9322 License Key



The ESL9322 has 322 physical slots. Libraries are sold with licenses for either 222 or 322 slots. A customer who wishes to upgrade from 222 to 322 slots must enter a license key through the touch panel.

## ESL9595 License Key

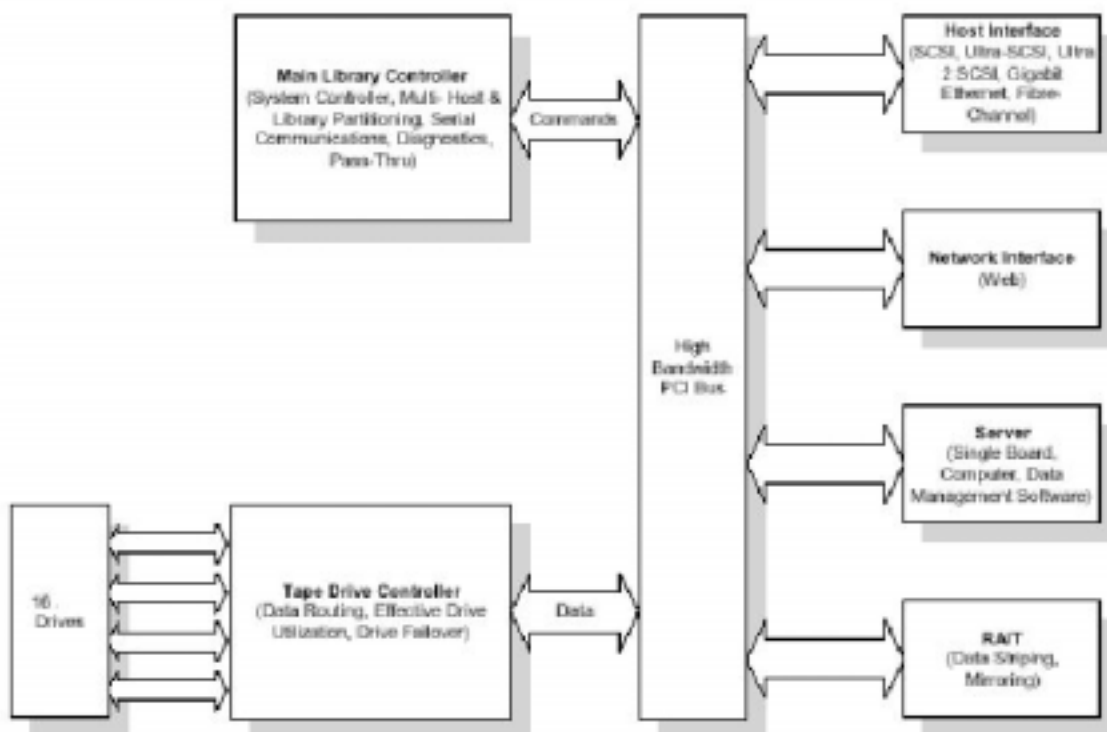


The ESL9595 has 595 physical slots. Libraries are sold with licenses for either 400, 500, or 595 slots. A customer who wishes to upgrade the number of slots must enter a license key through the touch panel.

The license key for both the ESL9322 and ESL9595 is unique for a given library serial number.




## ESL9000 Architecture



Utilizing the Prism bus architecture, the PCI backplane provides easy addition of new technologies in the future. Six cPCI expansion slots are available for future additions. The PCI backplane is network ready, Fibre ready, server ready, and tape array ready.


The main library controller responds to commands from a host computer to control the robotics and movement of cartridges. Components that the main library controller communicates with include the gripper, control panel, fans, bar code scanner, and the load ports.

### ESL9000 Front View




The load port holds up to 12 tapes for import/export purposes. The touch panel provides management of the tape library.




## ESL9000 Rear View



## ESL9000 Internal View




## ESL9000 Extension Axis Assembly



The extension axis assembly is mounted on the rotary axis and contains the extension motor/encoder assembly and the gripper assembly. Its function is to grip a single cartridge at a time so that the cartridge can be moved from a storage location to a tape drive and back.

## ESL9000 Vertical Axis Assembly



The vertical axis assembly provides Y-axis motion to the extension axis assembly through motors, belts, pulleys, and sensors.

The Y-axis stepper motor is mounted on the left side of the horizontal carriage assembly, at the base of the cabinet. It is identical to the X-axis stepper motor. They both are 38V DC, 11 amp stepper motors utilizing 9-conductor cable harnesses. It engages the vertical belt, driving the vertical carriage up and down the vertical rails.

### ESL9000 Horizontal Axis Assembly




The horizontal axis assembly provides X-axis motion to the extension axis assembly through motors, belts, pulleys, and sensors. The X-axis stepper motor is mounted in the lower left corner at the base of the cabinet. It is identical to the Y-axis stepper motor. They both are 38V DC, 11 amp, stepper motors utilizing 9-conductor cable harnesses. It engages the Horizontal Belt, which traverses three sides of the library cabinet (the bottom, right and top sides), and is approximately 30 feet long. These components drive the horizontal carriage between the left and right end of the cabinet.

### ESL9000 Rotary Stepper Motor



The rotary stepper motor protrudes down from the vertical carriage. It drives the 180-degree rotation of the extension axis. It drives the rotary axis to spin the extension axis to the front or rear of the cabinet

## ESL9000 Control Electronics Bay




The control electronics bay houses the following components:

- Robotics Backplane
- Robotics Controller
- Actuator Drivers (2 identical Printed Wiring Assemblies — PWAs)
- DC-DC Converters (2 identical PWAs)
- SCSI Interface



## ESL9000 AC Power Distribution Unit



One or two (redundant configurations) Power Input assemblies filter and route AC power to the Power System Backplane Printed Wiring Assembly (PWA).

AC power input is provided through the power system backplane PWA from one or two identical wide-ranging AC power distribution assemblies that accept single-phase input power of 90 VAC to 264 VAC at 47 Hz to 63 Hz. The figure above show a single AC power distribution assembly. The library can be supplied with a redundant AC power distribution unit. The library will, in that case, also have two power supply cords. One AC distribution assembly can supply AC power to up to five DC power supplies. The AC power distribution assemblies are located in the rear lower left corner of the cabinet.

## ESL9000 DC Power Supply Units




The DC power supply units are hot-swappable and 2N redundant.

## ESL9000 Fans



The fans are hot-swappable and redundant. Each fan has three associated status LEDs to indicate Normal, Slow, or Fault.

## ESL9000 Multi-Unit Scaling



Multi-scaling allows a customer to connect multiple libraries using the pass-through mechanism to form a virtual library. The host would see the virtual library as one library. The ESL9198 and ESL9322 libraries support multi-module configurations of up to five modules, and the ESL9595 libraries support up to four modules. The ESL9595, ESL9326, ESL9322, and ESL9198 modules can also be mixed in the multi-module configurations.


---

**Note**

Multi-unit scaling is not supported for ESL9000 libraries with SDLT320 drives at initial release.

---

## ESL9000 Multi-Scaling Cabling



The QSPI bus is a full-duplex synchronous serial interface for communicating with peripherals and other micro-controller units. The QSPI bus is a standard Motorola bus that ESL9000 multi-unit library systems use as a custom control language to communicate between master and slave libraries.

---


### Note

The QSPI bus must be terminated at both ends for proper operation.

---

1. Install a QSPI terminator on the left QSPI connector 2 for the left library.
2. Connect one end of the interconnect cable to the right QSPI connector 3 on the left library.
3. Route the interconnect cable through the bottom of the library and then connect it to the left QSPI connector on the right library.
4. If there are no other libraries in the configuration, install a terminator in the right QSPI connector of the right library. If there are additional libraries in the configuration, continue connecting the libraries using the interconnect cable. Note that the final library in the configuration will have a QSPI terminator on the right QSPI connector.

## Library GUI and Operations for the ESL9000 Family



Initial Screen of ESL9322 GUI

In most circumstances, the ESL9000 libraries are operated from the host using a backup application software package. However, the on-board Control Panels feature a touch-screen menu system that enables you to determine library status, configure the library, and perform certain diagnostic functions.

---

**Note**

The system must be off-line (standby) in order to perform many of the tasks from the Control Panel. Be sure backup software is not running.

---

The menu screen consists of a:

- Vertical toolbar (left column)
- Horizontal toolbar (top row)
- Main display area

## Vertical Toolbar

The vertical toolbar consists of the following items (from top to bottom):

- **System status display** — Displays the current status of the library system and important messages relating to library operation
- **Standby button** — Switches the library between online and standby status
- **Load Port button** — Opens the load port for inserting or removing a tape cartridge
- **Security level indicator** — Shows one of three security levels currently available at the Control Panel:
  - User (U)
  - Operator (O)
  - Service (S)
- **Stop button** — Halts and resumes all activity of the library by cutting and restoring power to the library robotics

## Horizontal Toolbar

The horizontal toolbar consists of the following menu navigation buttons for use with the tabbed screens in the main display area:

- The Back, Forward, and Home buttons enable navigation through previous screens in the order they were selected, much like the same buttons on a web browser.
- The Contrast buttons adjust the contrast of the screen.

## Main Display Area

The main display area is headed by four tabs:

- Overview
- Tapes
- Operator
- Service

Touching the appropriate tab on the screen gives access to the following main menu functions.

Overview	Tapes	Operator	Service
Tape Drive Status	Tape Drives	Configure Library	Statistics
Robot Activity	Storage	Configure Options	Actuator
Load Port	Transport	Move Cartridges	Sys Test Library Results
		Inventory Tapes	Auto Clean
		Calibrate Library	Sys Test Library
		Exercise Library	Operate Axes
		Unload Drives	Initialize Nonvol Statistics
		Unload Imp/Exp	Initialize Nonvol Configuration
			Change Password



ESL9000 GUI Overview Tab



ESL9000 GUI Tapes Tab



ESL9000 GUI Operator Tab





ESL9000 GUI Service Tab

## **ESL9000 Drive Upgrade**

ESL9000 tape libraries are designed for easy upgrades to tape drives. In general, the following steps must be taken:

1. Verify that the newer type of drive is supported for a particular tape library.
2. Update the firmware for the library.
3. Replace or insert the drive.

Other modifications may be necessary, depending the tape library and drive type.

## ESL9198SL Specifications

Category	Specifications
Maximum data transfer rate with 8 drives (maximum configuration)	316.8GB/hr native data 633.6GB/hr with 2:1 compression
Storage capacity	21.78TB native data 43.56TB with 2:1 compression
SCSI Interface	Ultra2 SCSI, LVD
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	SDLT cartridges
Status indicators	Full-function GUI; compatibility with Compaq Insight Manager for remote monitoring and administration
Options	Additional SDLT drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## ESL9322L1 Specifications

Category	Specifications
Maximum data transfer rate with two drives	108GB/hr native data 216GB/hr with 2:1 compression
Maximum data transfer rate with eight drives (maximum configuration)	432GB/hr native data 864GB/hr with 2:1 compression
Storage capacity	32.2TB native data 64.4TB with 2:1 compression
SCSI Interface	Low-Voltage Differential, Ultra2 SCSI
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	LTO Ultrium cartridges
Status indicators	Full-function GUI; compatibility with Insight Manager for remote monitoring and administration
Options	Additional Ultrium 230 drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## ESL9322S2 Specifications

Category	Specifications
Maximum data transfer rate with two drives	115.2GB/hr native data 230.4GB/hr with 2:1 compression
Maximum data transfer rate with eight drives (maximum configuration)	460.8GB/hr native data 921.6GB/hr with 2:1 compression
Storage capacity	51.5TB native data 103TB with 2:1 compression
SCSI Interface	Low-Voltage Differential, Ultra2 SCSI
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	SDLT cartridges
Status indicators	Full-function GUI; compatibility with Insight Manager for

	remote monitoring and administration
Options	Additional SDLT 320 drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## ESL9595L1 Ultrium 230 Specifications

Category	Specifications
Maximum data transfer rate with two drives	108GB/hr native data 216GB/hr with 2:1 compression
Maximum data transfer rate with 16 drives (maximum configuration)	864GB/hr native data 1728GB/hr with 2:1 compression
Storage capacity	59.5TB native data 119TB with 2:1 compression
SCSI Interface	Ultra2 SCSI, LVD
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	LTO Ultrium cartridges
Status indicators	Full-function GUI; compatibility with Insight Manager for remote monitoring and administration
Options	Tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## ESL9595S2 SDLT 320 Specifications

Category	Specifications
Maximum data transfer rate with two drives	115.2GB/hr native data 230.4GB/hr with 2:1 compression
Maximum data transfer rate with 16 drives (maximum configuration)	921.6GB/hr native data 1843GB/hr with 2:1 compression
Storage capacity	95.2TB native data 190.4TB with 2:1 compression
SCSI Interface	Ultra2 SCSI, LVD
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	SDLT cartridges
Status indicators	Full-function GUI; compatibility with Insight Manager for remote monitoring and administration
Options	Tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## ESL9595SL Specifications

Category	Specifications
Maximum data transfer rate with two drives	79.2GB/hr native data 158.4GB/hr with 2:1 compression

Maximum data transfer rate with 16 drives (maximum configuration)	633.6GB/hr native data 1.26TB/hr with 2:1 compression
Storage capacity	65.45TB native data 130.9TB with 2:1 compression
SCSI Interface	Ultra2 SCSI, LVD
MSBF	1 million swaps
MTBF	250,000 hours (20% duty cycle)
Media	SDLT cartridges
Status indicators	Full-function GUI; compatibility with Compaq Insight Manager for remote monitoring and administration
Options	Additional SDLT drive kits, tape cartridges, cleaning tapes, and bar code labels for cartridges
Warranty	Three-year limited warranty: first year on-site parts and labor, second and third years parts only

## **SCSI Target Guidelines for Tape Libraries**

In general, the following guidelines should be observed in setting SCSI target IDs:

- Assign a unique SCSI ID to the robot (changer) and to each tape drive.
- Assign the robot a SCSI ID lower than the SCSI IDs of the tape drives.
- Assign SCSI IDs of the tape drives in ascending order with physical location.

## **Fibre Channel LUN Guidelines for Tape Libraries**

In general, the following guidelines should be observed in setting Fibre Channel LUNs on the Modular Data Router:

- Assign LUNs in ascending order, starting with zero for the SCC controller.
- Assign each robot a LUN lower than the LUNs of its drives.
- If multiple drives are controlled by a robot, assign LUNs to drives based on ascending SCSI Bus-Target-LUN order.
- If multiple libraries are attached to the MDR, assign LUNs to the robot and associated drives of the lowest SCSI bus first, then assign LUNs to the robot and associated drives of the next lowest SCSI bus, and so forth.

## Learning Check

1. The ESL9198 family of tape libraries
  - a. Accommodates up to eight tape drives
  - b. Accommodates up to sixteen tape drives
  - c. Stores up to 198 tape cartridges
  - d. Stores up to 326 tape cartridges
2. Explain the license keys for the ESL9322 and ESL9595 libraries.  
.....  
.....  
.....  
.....
3. What is the purpose of the load port on the ESL9000 libraries?  
.....  
.....
4. In general, what are the steps involved in upgrading the tape drives in the ESL9000 libraries?  
.....  
.....  
.....  
.....
5. Which of the following contains the bar code reader?
  - a. Vertical Axis Assembly
  - b. Horizontal Axis Assembly
  - c. Extension Axis Assembly
  - d. Control Electronics Bay
  - e. Rotary Stepper Motor






### **Objectives:**

After completing this module, you should be able to:


- Describe the processes of converting Fibre Channel to SCSI and vice versa.
- Describe the features and benefits of the Network Storage Router M2402.
- Describe the hardware components of the Network Storage Router M2402.
- Describe the features of the Network Storage Routers N1200, e1200, and e2400.
- Describe the management features of the Visual Manager User Interface.
- Describe recommendations for optimizing performance with the network storage routers.

## Network Storage Router Overview



LAN-free backup and restore

The Network Storage Router family supports LAN-free backup and restore for high performance through movement of data across a SAN. The router translates Fibre Channel to SCSI and vice-versa.



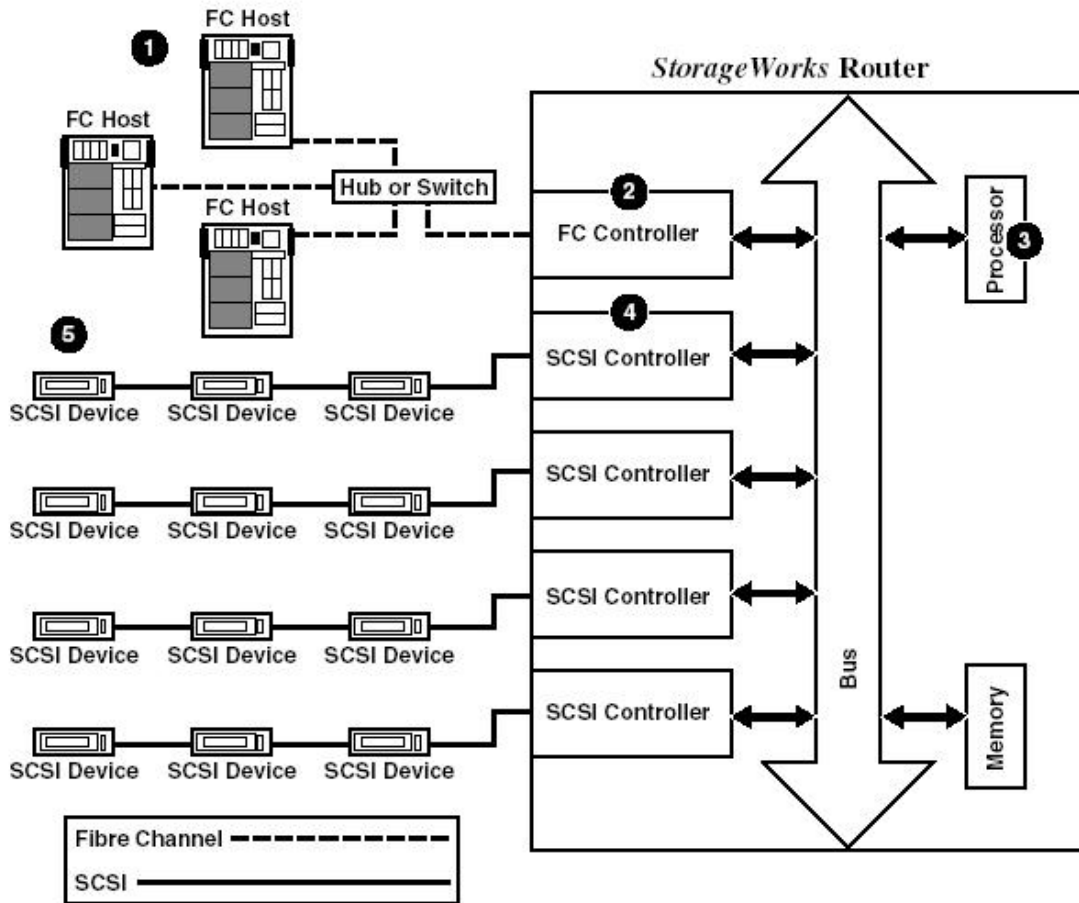
Server-free data movement

In addition, the Network Storage Router family supports third-party copy, which allows server-free data movement for reducing overhead on servers during backups.

Other features of the Network Storage Router family include:

- High Performance
  - 2Gb Fibre Channel Support
  - Ability to support up to 300 MB/s of information throughput
- Web Based Management
- Mapping and masking at the device level

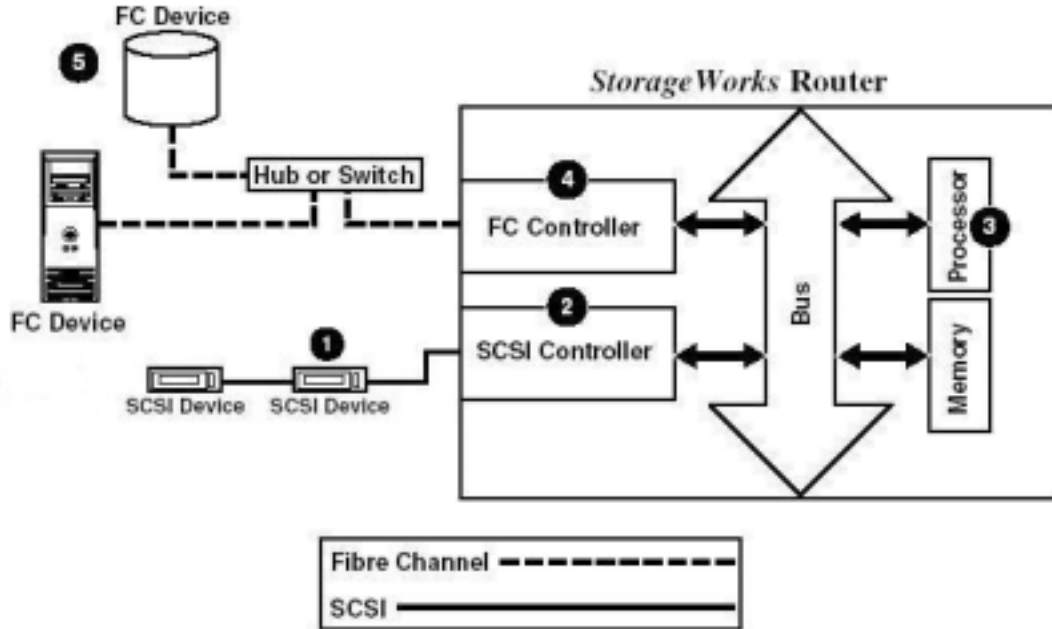
## Converting Fibre Channel to SCSI



The following sequence describes the sequence the network storage router uses to convert Fibre Channel host protocol to SCSI device protocol.

1. A Fibre Channel host issues an encapsulated FCP protocol command packet to the router.
2. The router Fibre Channel controller interprets the Fibre Channel information and places the packet in buffer memory.
3. The router interprets the Fibre Channel information packet and programs the router SCSI controller to process the transaction.
4. The router SCSI controller sends the command to the SCSI device (target).
5. The SCSI device interprets the command and executes it.

## Converting SCSI to Fibre Channel



The following sequence describes the sequence the network storage router uses to convert SCSI device protocol to Fibre Channel device protocol.

1. A SCSI host issues a command to the router.
2. The SCSI controller in the router interprets the command and places it in buffer memory.
3. The router processor interprets data and programs the router Fibre Channel controller to process the transaction.
4. The router Fibre Channel controller translates data into an FCP protocol packet and sends it to the Fibre Channel target.
5. The Fibre Channel target interprets the FCP protocol packet and executes the command.

## Network Storage Router M2402



The HP StorageWorks Network Storage Router M2402 is a 1U Fibre Channel-to-SCSI router with a dual-port Fibre Channel module and a four-port SCSI module on the base model. The Network Storage Router (NSR) allows multiple host servers to communicate with a SCSI tape device over a Fibre Channel link. It offers support for LAN-free backup, serverless backup, 1Gb and 2Gb Fibre Channel and has web based management allowing for remote management.


The NSR is modular and can be upgraded by purchasing additional Fibre Channel or SCSI modules. Up to four modules of any combination of Fibre Channel and SCSI modules can be installed in the NSR.

### Features and Benefits

- **Increased performance** – This router supports a throughput of up to 300MB/s. The router can stream data to up to 8 SDLT 110/220GB tape drives.
- **Serverless backup support** – The router supports third-party copy commands to reduce the impact on the hosts during backups for backup software that support third-party copy.
- **Web-based management** – Visual Manager enables an administrator to view statistics, change settings and configurations, as well as running various utilities from a browser.
- **Modular design** – The modular design allows a customer to buy additional Fibre Channel and/or SCSI modules as needed. In addition, the modular design allows upgrades to other modules that may become available in the future. The NSR supports up to three SCSI modules and up to three Fibre Channel modules.
- **Reliability, Availability, and Serviceability** –
  - Hot swappable fan module consists of three fans in a single module and can tolerate the loss of any single fan. The fan module can be replaced without taking the router down, avoiding unplanned down time.
  - By purchasing a second power supply, the NSR can tolerate the failure of a single power supply. The power supplies are hot swappable.
  - Enhanced reliability features include automatic shutdown on over temperature conditions.

## Hardware Components


### Main Board



The NSR M2402 uses the Crossroads 64bit/66Mhz PCI Modular Interface™ (CPMI). All I/O slots are the same, guaranteeing module and slot interoperability. A dual port memory controller regulates PCI and memory bandwidth.


Through the Crossroads Electronic Branding Technology, the NSR identifies modules by vendor identification during POST. Mismatched vendors between a module and the base unit will cause a mismatch device error and prohibit the module from functioning. The NSR accepts only HP brand modules.

## Fibre Channel Module



The Fibre Channel module has two independent Fibre Channel ports for load balancing. Each port has an SFP optical interface, supporting 1Gb and 2Gb Fibre Channel. Each port must be configured by the administrator to either 1Gb/s or 2Gb/s, based on the Fibre Channel port that is connected on the other end of each cable. The Fibre Channel ports support either FC-AL or FC-SW.

## SCSI Module



The SCSI module has four ports, each which has an auto-negotiating SCSI bus. The LVD SCSI module supports four independent SCSI buses running at up to 160MB/s each. The HVD SCSI modules supports four independent SCSI buses running at up to 80MB/s each. Both initiator and target modes are supported.



## Network Storage Router N1200



The HP StorageWorks Network Storage Router N1200 is a key component in a complete data protection solution. It is a 1U rackmount router with one Fibre Channel port and two SCSI ports. It allows customers to connect SCSI-based tape libraries to the SAN increasing backup speeds up to 90%. The N1200 offers enterprise class features at a price point appealing to the low to mid range customer. It is targeted towards customers that want to reduce or remove their backup window. It offers support for serverless backup, 2-Gb Fibre Channel and has web-based management allowing for remote management.


Features of the NSR N1200 include:

- Serverless backup support
- 2-Gb or 1-Gb Fibre Channel Support
- Web Based Management — The N1200 and the M2402 share a common web-based interface.
- Ultra2 SCSI LVD support
- Provides connectivity for up to two tape drives per SCSI port
- Exceeding 140 MB/s throughput
- Lower cost per drive connection
- Single fan and single power supply

## Embedded Network Storage Routers

Embedded network storage routers are available for the MSL5000 Library family as well as the ESL9000 Library family. Embedded network storage routers save rack space by implementing the Fibre Channel to SCSI router inside the tape library. In addition, the embedded network storage routers include the same Web-based management interface used in other StorageWorks network storage routers.

### Network Storage Router e1200



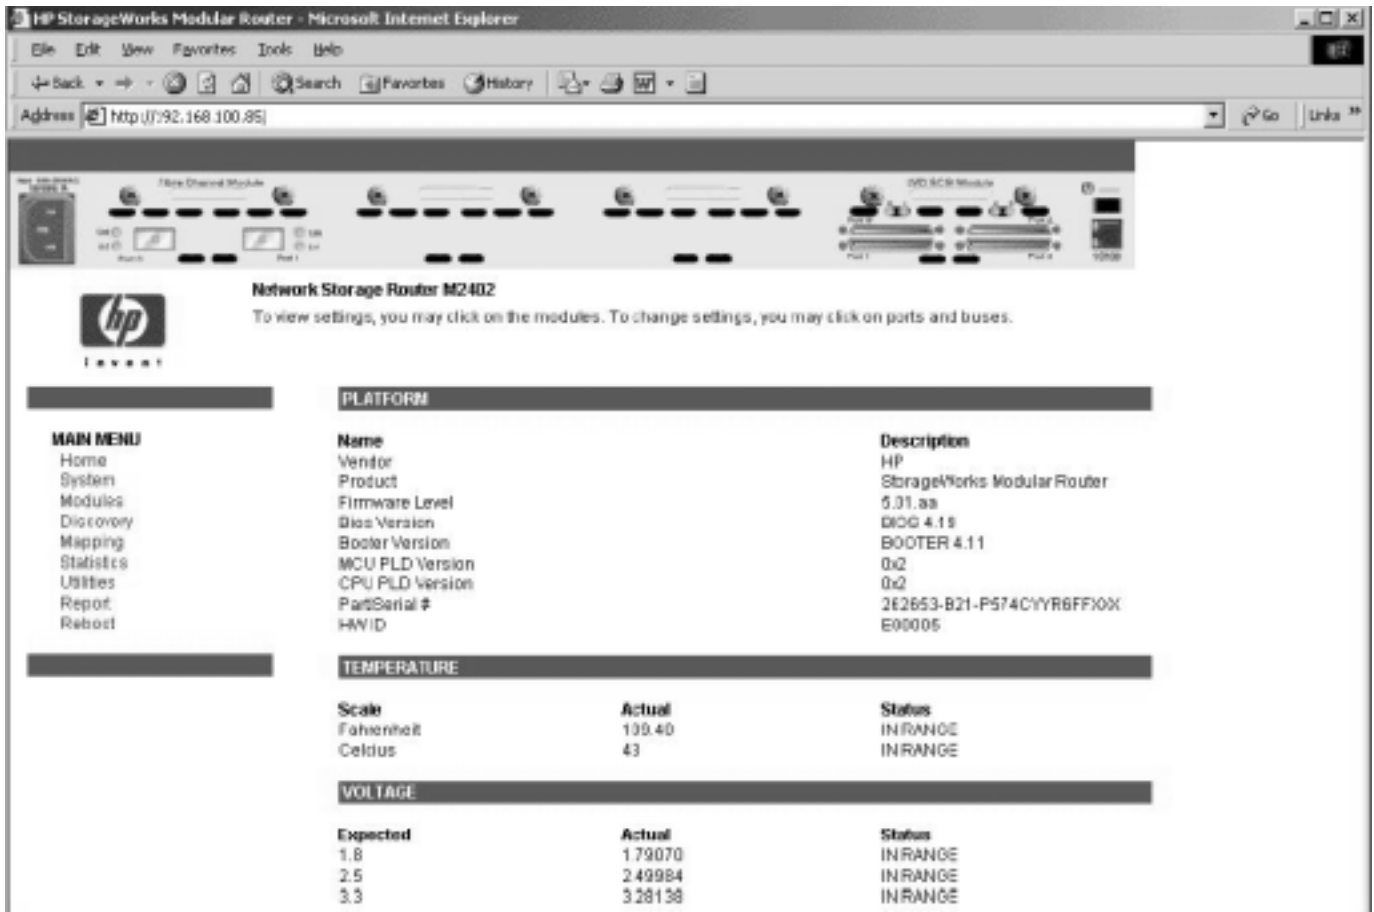
The HP StorageWorks Network Storage Router e1200 consists of a router card with a single Fibre Channel (1Gb or 2Gb) port and two Ultra2 SCSI (LVD/SE) ports. The e1200 is designed to be installed in the cPCI slot of the MSL5000 Library family. Each SCSI port supports up to two SDLT or DLT drives. A maximum of one card is supported in the library.

### Network Storage Router e2400

The HP StorageWorks Network Storage Router e2400 solution consists of a cPCI cage, which must be installed in the rear of the ESL9000 library, and up to four router cards, each with a single Fibre Channel (1Gb or 2Gb) port and two Ultra2 SCSI (LVD/SE) ports. Each SCSI port supports up to two SDLT or DLT drives. A maximum of one card is supported in the library.

The ESL9000 Embedded Fibre Option Kit includes one cPCI cage and two router cards. An upgrade kit which includes two router cards is available. HP highly recommends that the installation CarePak be ordered with the Option Kit.

## Visual Manager User Interface



The Visual Manager user interface provides a graphical format for remotely viewing information and configuring the NSR. The home page contains status information, including a graphical view of the rear. The router image is interactive, allowing quick access to configuration menus for the Fibre Channel ports, SCSI buses, power supply, Ethernet port, and LED beacon.

Router status information in the body of the home page includes

- Platform information about the NSR hardware
- Temperature
- Voltage measurements
- Fan measurements
- Power supply functionality

The Main Menu options bar is located at the left side of the home page. Main Menu options include

- **Home** – Display router status information.
- **System** – Configure standard system components.

- **Modules** – Configure Fibre Channel ports and SCSI buses.
- **Discovery** – View devices and discover new devices.
- **Mapping** – View and configure maps.
- **Statistics** – View router statistics.
- **Utilities** – Configure utility settings.
- **Report** – View and print system information.
- **Reboot** – Restart the router.

## System Menu

		SERIAL		NETWORK	
<b>MAIN MENU</b> Home System Modules Discovery Mapping Statistics Utilities Report Reboot  <b>SYSTEM MENU</b> Serial Network SNMP Active Fabric User Real-Time Clock Power Supply  <b>RESET MENU</b> Factory Settings Reset		<b>Name</b>	<b>Status</b>	<b>Name</b>	<b>Status</b>
		Baud Rate	115200	MAC Address	00:E0:02:E1:F7:25
				IP Address	192.168.100.85 (DHCP)
				Subnet Mask	255.255.255.0 (DHCP)
				IP Gateway	192.168.100.1
				Ethernet Mode	10/100Mbps
				DHCP Client	Enabled
				Hostname	Router
		<b>SNMP</b>		<b>TRAP MANAGER 1</b>	
		<b>Name</b>	<b>Status</b>	<b>Name</b>	<b>Status</b>
		Community Get	public	IP Address	1.1.1.1
		Community Set	private	Port	162
		Traps	Disabled	Filter	Log All Events
				Row State	No Traps (1)
		<b>TRAP MANAGER 2</b>		<b>TRAP MANAGER 3</b>	
		<b>Name</b>	<b>Status</b>	<b>Name</b>	<b>Status</b>
		IP Address	1.1.1.1	IP Address	1.1.1.1
		Port	162	Port	162
		Filter	Log All Events	Filter	Log All Events
		Row State	No Traps (1)	Row State	No Traps (1)
		<b>ACTIVE FABRIC</b>		<b>REAL-TIME CLOCK</b>	
		<b>Name</b>	<b>Status</b>	<b>Name</b>	<b>Status</b>
		Backup Mode	Disabled	Date	Wednesday 10/02/2002
		Controller LUNS	1	Time	09:05:13
		<b>POWER SUPPLY</b>			
		<b>Name</b>	<b>Status</b>		
		Number Installed	1		

The System Menu is used to view and configure serial, network, SNMP, Trap, Active Fabric, clock, and power supply components. Three different trap destinations can be set. Active Fabric settings are used to enable serverless backup capability.

## Modules Menu

**Network Storage Router M2402**  
To view settings, you may click on the modules. To change settings, you may click on ports and buses.

**MAIN MENU**  
Home  
System  
Modules  
Discovery  
Mapping  
Statistics  
Utilities  
Report  
Reboot

**MODULES MENU**  
World Wide Node Name  
**FC MODULE 0**  
Port 0  
Port 1  
**NO MODULE 1**  
**NO MODULE 2**  
**SCSi MODULE 3**  
Bus 0  
Bus 1  
Bus 2

FC MODULE 0		NO MODULE 1	
Name	Status		
Protocol	FCP		
Compatibility	PASSED ALL		
Part/Serial #	262653-B21_q574c0x56YYGGG		
HW ID	P00006		

NO MODULE 2		SCSi MODULE 3	
		Name	Status
		Protocol	SCSi [LVD]
		Compatibility	PASSED ALL
		Part/Serial #	262653-B21_q574c0x56YYGGG
		HW ID	P00005

The Modules Menu allows the administrator to change the World Wide Node Name, Fibre Channel port settings, and SCSI bus settings.



### Note

HP recommends using the default values for World Wide Names.

## Discovery Menu

**MAIN MENU**  
Home  
System  
Modules  
Discovery  
Mapping  
Statistics  
Utilities  
Report  
Reboot

**DISCOVERY MENU**  
**FC MODULE 0**  
Port 0  
Port 1  
**NO MODULE 1**  
**NO MODULE 2**  
**SCSI MODULE 3**  
Bus 0  
Bus 1  
Bus 2  
Bus 3

Perform Discovery on all Ports and Buses:

**SCSI MODULE 3 DISCOVERY**

Protocol	Bus	Device Type	State	Map Count	Target Id	Lun Id
SCSI	3	TAPE	UP	1	2	0
SCSI	3	TAPE	UP	1	3	0

The Discovery Menu provides a view of target devices as well as an option to discover target devices.

## Mapping Menu



The Mapping Menu is used to view and modify host and map information for a Fibre Channel port. Maps and hosts may be added, edited, or deleted.

A map defines the devices accessible through a particular Fibre Channel port. The administrator then assigns each host a map for each Fibre Channel port. Multiple maps may be created for a single Fibre Channel port.



## Statistics Menu

**MAIN MENU**  
Home  
System  
Modules  
Discovery  
Mapping  
Statistics  
Utilities  
Report  
Reboot

**STATISTICS MENU**  
**FC MODULE 0**  
Port 0  
Port 1  
**NO MODULE 1**  
**NO MODULE 2**  
**SCSI MODULE 3**  
Bus 0  
Bus 1  
Bus 2  
Bus 3

Reset System Statistics on all SCSI Buses: **Go**

**FC MODULE 0 STATISTICS**

Name	Port 0	Port 1
In Device Data Sequences	0x000000FB	0x00000000
Out Device Data Sequences	0x00000034	0x00000000
In Link Data Sequences	0x00000057	0x00000000
Out Link Data Sequences	0x000000A5	0x00000000
In PBSY Frames	0x00000000	0x00000000
Out PBSY Frames	0x00000000	0x00000000
In FBSY Frames	0x00000000	0x00000000
In PRJT Frames	0x00000000	0x00000000
Out PRJT Frames	0x00000000	0x00000000
FC Link Down	0x00000003	0x00000000
In Aborts	0x0000004B	0x00000000
Out Aborts	0x00000000	0x00000000
Laser Faults	0x00000000	0x00000000
LOS	0x00000000	0x00000000
Sync	0x00000000	0x00000000
Bad Rx Characters	0x00000000	0x00000000
Link Failures	0x00000001	0x00000001
Bad CRCs	0x00000000	0x00000000
Protocol Errors	0x00000000	0x00000000
Bad SCSI Frames	0x00000000	0x00000000

**NO MODULE 1**

**NO MODULE 2**

**SCSI MODULE 3 STATISTICS**

Name	Bus 0	Bus 1	Bus 2	Bus 3
Mode	LVD	LVD	LVD	SE
Resets	17	1	61	4
Active IDs	0, 1	None	4, 5	2, 3

The Statistics Menu displays Fibre Channel port and SCSI bus information.

## Utilities Menu

MAIN MENU	
Home	
System	
Modules	
Discovery	
Mapping	
Statistics	
Utilities	
Report	
Reboot	
UTILITIES MENU	
Beacon	
FTP Utility	
Trace Settings	
Current Traces	
Previous Traces	
Last Assert Traces	
Clear Current Traces	
Clear Assert Traces	
Event Log Settings	
Event Log	
Clear Event Log	
BIOS INITIALIZATION	
Platform ID Initialization	PASSED
Memory Test Initialization	PASSED
Calendar Power Initialization	PASSED
Calendar Initialization	PASSED
Indicator Initialization	PASSED
IC2 Initialization	PASSED
Resource List Initialization	PASSED
PCI SDRAM Initialization	PASSED
SCSI Configuration	PASSED
Flash Type	PASSED
Instruction SDRAM	PASSED
EBT Error	PASSED
EBT Vendor	PASSED
PCI Module	PASSED
BOOTUP INITIALIZATION	
CPU Program RAM Test	PASSED
PCI Protocol RAM Test	PASSED
SCSI Module 3 Port 0 Script RAM	PASSED
SCSI Module 3 Port 1 Script RAM	PASSED
SCSI Module 3 Port 2 Script RAM	PASSED
SCSI Module 3 Port 3 Script RAM	PASSED
SCSI Module 3 Port 0 POST	PASSED
SCSI Module 3 Port 1 POST	PASSED
SCSI Module 3 Port 2 POST	PASSED
SCSI Module 3 Port 3 POST	PASSED
FC Module 0 Port 0 POST	PASSED
FC Module 0 Port 1 POST	PASSED
Ethernet POST	PASSED
TRACE	
<b>Name</b>	<b>Status</b>
General Errors	ON
FCP Driver	OFF
FCP Transport	OFF
FCP Management	OFF
EVENT LOG	
<b>Name</b>	<b>Status</b>
All Events Logged	ON

The Utilities Menu displays POST results. In addition, the following utilities are available:

- Beacon – Enable the Beacon LED mode.
- FTP – Open an FTP session.
- Trace Settings – Configure trace settings.
- Current Traces – View current trace information.
- Previous Traces – View previous trace information.
- Last Assert Traces – View last assert trace information.
- Clear Current Traces – Clear current trace information.
- Clear Assert Traces – Clear assert trace information.
- Event Log Settings – Configure the Event Log.
- Event Log – Display the Event Log.
- Clear Event Log – Clear the Event Log.

## Report Menu

MAIN MENU

Home

System

Modules

Discovery

Mapping

Statistics

Utilities

Report

Reboot

ENCLOSURE INFORMATION

Platform

Vendor

Product

Firmware Level

Bios Version

Booter Version

MCU PLD Version

CPU PLD Version

Part/Serial #

HW ID

Description

HP

StorageWorks Modular Router

5.01.aa

BIOS 4.19

BOOTER 4.11

0x2

0x2

262653-B21-P574CYYR6FFXXX

E00005

Temperature Scale

Fahrenheit

Celsius

Actual Temperature

109.40

43

Status

IN RANGE

IN RANGE

Expected Voltage

1.8

2.5

3.3

5.0

12.0

1.8

Actual Voltage

1.79070

2.49984

3.28138

4.97364

12.06250

1.76250

Status

IN RANGE

IN RANGE

IN RANGE

IN RANGE

IN RANGE

IN RANGE

Fan

Fan 0

Fan 1

Fan 2

RPM

9246

8544

8881

Status

IN RANGE

IN RANGE

IN RANGE

Power Supply

Power Supply 0

Power Supply 1

Status

FUNCTIONAL

NOT FUNCTIONAL

MODULE INFORMATION

FC Module 0

Protocol

Compatibility

Part/Serial #

HW ID

FCP

PASSED ALL

262653-B21\_q574c0x56YY000

P00006

The Report page displays a consolidated view of all system information.

## Reboot Menu

The Reboot Menu allows the administrator to initiate a reboot of the NSR. When the router is rebooted, current router activities will be disrupted. All submitted configuration changes will be activated during the boot process.

## Optimizing Performance on the Network Storage Router

In order to optimize performance, a maximum of one tape drive per SCSI port is always the best configuration. However, network storage routers will support up to two drives per port for most tape drives except Ultrium 460.

Limitations of the network storage routers include the type of SCSI interface on the router as well as the router's bandwidth. The M2402 supports Ultra3 SCSI, but the e2400, e1200, and N1200 support Ultra2 SCSI. In addition, the M2402 has a maximum bandwidth of 300MB/s.

Below is a chart for supported and optimal configurations.

Router	SDLT 220 (Ultra2)		SDLT 320 (Ultra3)		Ultrium 230 (Ultra2)		Ultrium 460 (Ultra3)	
	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming
N1200	2	2 per router	Not supported	Not supported	2	2 per router	Not supported	Not supported
e1200, e2400	2	2 per card	Not supported	Not supported	2	2 per card	Not supported	Not supported
M2402	2	8 per router, one per SCSI port	2	8 per router, one per SCSI port	2	8 per router, one per SCSI port	1	4 per router, one per SCSI port

## Learning Check

1. What is the purpose of the router Fibre Channel controller?  
.....  
.....  
.....
2. Which Network Storage Router provides the option of a second power supply for redundancy?  
.....
3. Which Network Storage Router provides support for Ultra3 SCSI?  
.....
4. For optimal performance, what should be the maximum number of tape drives per SCSI port on the Network Storage Router?  
.....
5. What is the purpose of the Mapping menu of the Visual Manager User Interface?  
.....  
.....  
.....



---

# Software Components of Enterprise Backup Solutions

## Module 6

### Objectives

After completing this module, you should be able to:

- List the required software components and what each does for the solution.
- List the features of the hp OpenView Data Protector, VERITAS Backup Exec, Computer Associates BrightstorARCserve, VERITAS NetBackup, Legato NetWorker, Tivoli Storage Manager, and CommVault Galaxy software applications.
- Explain the topology of the Data Protector, Backup Exec, ARCserve, NetBackup, NetWorker, Tivoli Storage Manager, and CommVault Galaxy environments.
- Identify the current caveats of the Enterprise Backup Solutions.

## EBS Software Components

A crucial component of hp EBS is the backup management software that enables all the hardware components to work together.

HP has partnered with VERITAS, Computer Associates (CA), Legato, Tivoli, and CommVault to deliver an integrated, centralized, Fibre Channel storage solution. When configured with partner software, EBS enables multiple servers to share libraries.

Each server in the EBS must have certain drivers installed to access the shared tape device over the FC-AL or FC-SW.

The operating system is next, followed by the backup application.

The software component that brings the whole multiserver solution together is the device sharing option. The following table outlines the device sharing software for each solution.

Backup Software	Device Sharing Option
HP Openview Data Protector	Media Agent
VERITAS Backup Exec	Shared Storage Option (SSO)
Computer Associates Brightstor ARCserve	SAN option Tape Library Option (TLO)
VERITAS NetBackup	Multihost License
Legato NetWorker	Smart Media Server (SMS) for NetWorker 6.0.1; Embedded in NetWorker 6.1
Tivoli Storage Manager	Tape Library Sharing
CommVault Galaxy	Media Agent

When using ARCserve 7 and ARCserve 2000, TLO controls the automated loading and unloading of tapes. These functions are handled by SSO within Backup Exec and the Multihost license within NetBackup.



## Maximum Supported Configuration

In general, to configure backup software so that multiple servers are sharing one or more tape libraries, a primary server must be installed. This primary server will manage the database which may track devices, media, ownership, and other indices. The group that is comprised of the primary server, the tape libraries, and the other servers that share the tape libraries is often called a data zone or backup domain.

The members of this data zone may be members of different switch zones in the SAN. However, in order for all servers in the data zone to access the tape library, the tape library must exist in each of these switch zones in the SAN or a separate switch zone must be created that contains the tape library along with all of the servers that are part of the data zone.

In addition, depending on the specific ISV software, members of the data zone can also be running different operating systems.

Data zones have limited supported configurations. For example, with VERITAS Backup Exec, the maximum supported configuration is 32 servers sharing up to 32 tape drives in one data zone. If one were to exceed the maximum configuration, severe performance degradation or loss of communication can occur. In addition, a configuration involving fewer servers and tape drives may be necessary in order to achieve optimal performance. The following table lists the maximum supported configuration with EBS as of December 2002:

Backup Software	Maximum Configuration
HP OpenView Data Protector	16 servers & 16 tape drives
VERITAS Backup Exec	32 servers & 27 tape drives for NetWare 32 servers & 32 tape drives for Windows NT/2000
Computer Associates Brightstor ARCserve	20 servers & 20 tape drives for NetWare 16 servers & 16 tape drives for Windows NT/2000
VERITAS NetBackup	32 servers & 32 tape drives
Legato NetWorker	16 servers & 32 tape drives for Windows 32 servers & 32 tape drives for UNIX
Tivoli Storage Manager	16 servers & 16 tape drives
CommVault Galaxy	16 servers & 16 tape drives
Syncsort Backup Express	16 servers & 16 tape drives
Computer Associates Brightstor Enterprise Backup	16 servers & 16 tape drives
Atempo Time Navigator	16 servers & 16 tape drives
BakBone NetVault	16 servers & 16 tape drives

Although multiple servers can share multiple drives, only one server may control the robot at a given time and only one server may access a particular tape drive at a given time. With some backup management software, a primary server controls the robot while other backup servers send tape mount and dismounts requests to this primary server.

In general, servers in a switch zone should be homogeneous in operating systems if possible. Novell NetWare and Sun Solaris servers are incompatible in the same switch zone. Linux, HP-UX, and IBM AIX servers are each incompatible in zones with all other servers. The following table summarizes switch zone compatibility between different operating systems in a SAN. Platforms in the same column can coexist in the same switch zone.

Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
NetWare	Sun	Linux	HP-UX	IBM AIX
OpenVMS	OpenVMS			
Tru64 UNIX	Windows NT			
Windows NT	Windows			
Windows	2000 Server,			
2000 Server,	Advanced			
Advanced	Server			
Server	Tru64 UNIX			

## Drivers

Drivers are software that allows hardware components to run correctly. The drivers necessary for EBS include:

- Host bus adapter (HBA)
- Tape
- Fibre Channel filter

### HBA Drivers

Each HBA in the server is run by driver software. These drivers are located in the solution kit supplied with each solution. When you install an HBA, you must install the appropriate drivers for that HBA.

### Tape Drivers

The tape drives in the library are run by tape driver software. If applicable, these drivers are located in the solution kit supplied with each solution.

### Fibre Channel Filter Drivers

Some SCSI commands to tape devices do not fit into the Fibre Channel protocol. This driver allows those commands to be executed transparently over Fibre Channel. The filter driver is automatically installed when the host adapter driver is installed.

## Device Persistency

When a host boots, the operating system scans for devices in the SAN. The backup software utilizes the information on devices provided by the operating system. When a user configures the backup software, the backup software saves this configuration based on an association of each tape drive's identity according to the operating system.

Ideally, a tape drive's identity according to the operating system should not change if a new component is added to the SAN switch or if a cable is plugged into a different port on the SAN switch within the same switch zone. The maintenance of identity of a device despite configuration changes to the SAN is called **device persistency**.

In UNIX operating systems, device persistency can be achieved for tape drives in the SAN. Tru64 UNIX binds the World Wide Name of the tape controller to a target identification. Because the World Wide Name does not change, target identification will not change, despite modifications to the SAN. Likewise, SUN Solaris allows a user to map target identification numbers to the tape controller's World Wide Name using the file *st.conf*.

Windows NT/2000 does not offer support for device persistency. The device driver for the host bus adapter assigns targets to World Wide Names in the order of discovery. If a target goes offline or a target is added, the remaining targets are each assigned a different target ID the next time the server reboots. Should a tape drive change device identity, the backup software may no longer be able access that tape drive.

When Windows NT/Windows 2000 boots, targets are assigned sequentially by switch domain number and then by port number. To avoid the lack of device persistency in Windows NT/Windows 2000, the tape controllers should be in the lowest port numbers on the SAN switch with the lowest switch domain number.

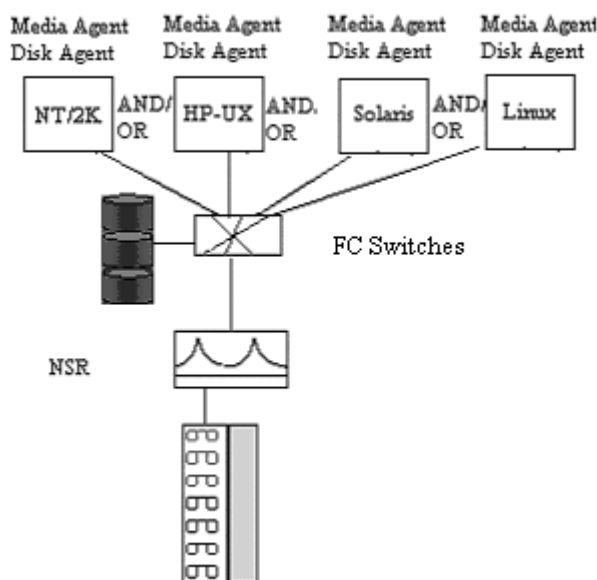
Backup Exec and Computer Associates maintain device persistency in Windows NT/Windows 2000 by tracking the serial number of the tape drive.

## Backup Software

The following backup applications for use with EBS are discussed in this module:

- HP Openview Data Protector
- VERITAS Backup Exec
- Computer Associates Brightstor ARCserve 7 or ARCserve 2000
- VERITAS NetBackup
- Legato NetWorker
- Tivoli Storage Manager
- CommVault Galaxy

## HP OpenView Data Protector Overview



HP OpenView Data Protector provides reliable data protection and comprehensive recovery. Whether backup requirements are for small workgroups or a multi-site heterogeneous data center environment, Data Protector provides centralized or distributed management, depending on the needs of the customer's organization. The modular design of Data Protector provides scalability as storage environments grow. Data Protector supports a comprehensive list of backup clients (disk agents) and backup device servers (media agents) to provide broad compatibility.

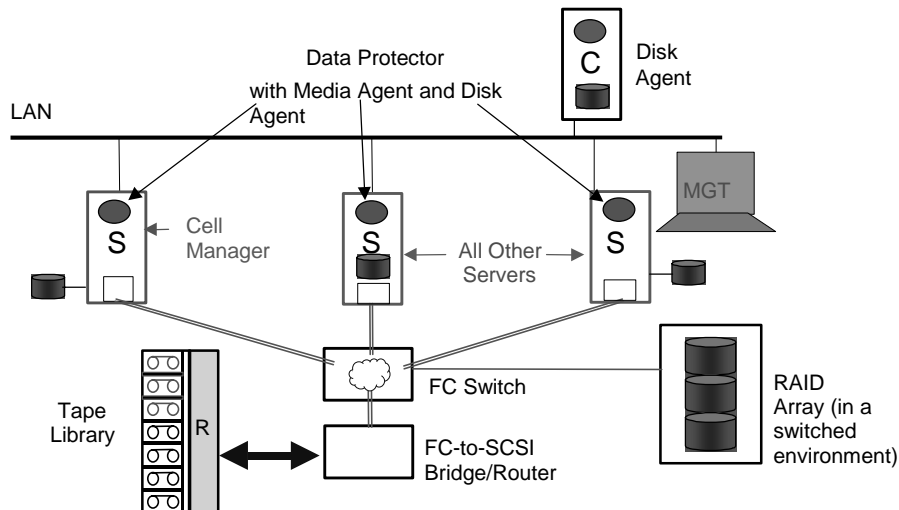
For HP OpenView Data Protector environments, the required components are:

- Windows NT 4.0 with Service Pack 6a, Windows 2000 Server with Service Pack 2, Solaris 8, Red Hat Linux 7.1, and HP-UX 11i
- HP OpenView Data Protector 5.0
- Disk Agent and Media Agent for each server sharing a library across a SAN

## Key Features

- **Scalable and Highly Flexible Architecture** — Data Protector can be used in environments ranging from a single system to thousands of systems on several sites.
- **Easy Central Administration** — Through its easy-to-use graphical user interface (GUI), Data Protector allows the administrator to manage multiple backup environments from a single system. The GUI can be installed on various systems to allow multiple administrators to access Data Protector via locally installed consoles. The Data Protector command-line interface allows you to manage Data Protector using scripts.
- **High Performance Backup** — Data Protector supports parallel data streams. Various types of backups, such as local, network, full, differential, leveled incremental, online, and disk image, provide flexibility in tuning the backups to best fit business requirements.
- **Heterogeneous Support** — Data Protector supports heterogeneous environments. The HP-UX, Windows NT, or Windows 2000 Cell Manager can control all client platforms.
- **Monitoring, Reporting, and Notification** — Reporting and notification capabilities allow the administrator to view the backup status, monitor active backup operations, and customize reports. Reports can be generated using Java-based online web reporting, the Data Protector GUI, or using the omnirpt command on systems running UNIX, Windows NT, or Windows 2000.
- **Integration with Online Database Applications** — Data Protector supports online backup of Microsoft Exchange, Oracle, Informix, SAP, Lotus Domino, and Sybase database objects.

## Data Protector Topology



The Data Protector cell is a network environment that has a Cell Manager, client systems, and devices. The Cell Manager is the central control point where Data Protector software is installed. Client systems are then added to the cell in order to be backed up. Data Protector backs up files to media in the backup devices.

### **Cell Manager**

The Cell Manager:

- Manages the cell from a central point
- Contains the Data Protector database, which contains information about backup details, such as backup durations, media IDs, and session IDs
- Runs core Data Protector software
- Runs Session Managers that start and stop backup and restore sessions and write session information to the Data Protector database

### **Client Systems**

Client systems that need to be backed up must have the Data Protector Disk Agent (also called Backup Agent) installed. The Disk Agent reads or writes data from a disk on the system and sends or receives data from the Media Agent. The Disk Agent is also installed on the Cell Manager, thus allowing you to back up data on the Cell Manager, the Data Protector configuration, and the Data Protector database.

### **Devices**

Client systems with access to backup devices attached directly or through a SAN must have the Media Agent installed. Such client systems are also called Drive Servers. The Media Agent reads or writes data from media in the device and sends or receives data from the Disk Agent.

### **Installation Server**


The Installation Server holds a repository of the Data Protector software packages for a specific architecture. By default, the Cell Manager is also an Installation Server. At least two Installation Servers are needed for mixed environments: one for UNIX systems and one for the Windows NT and Windows 2000 systems.

## **Large Data Protector Enterprise Environments**

A typical enterprise network environment consists of a number of systems from different vendors with different operating systems. The systems may be located in different geographical areas and time zones. All the systems are connected with LAN or WAN networks operating at various communication speeds.

This solution can be used when several geographically separated sites require common backup policies to be used. It can also be used when all departments at the same site want to share the same set of backup devices.






Configuring and managing backups of such a heterogeneous environment is challenging. Data Protector functionality has been designed to highly simplify this task.

### Splitting an Environment into Multiple Cells

Large environments may be split into multiple cells for a number of reasons:

- Geographical grouping of systems
- Logical grouping of systems, for example, departments
- Slow network connection between some systems
- Performance considerations
- Separate administrative control

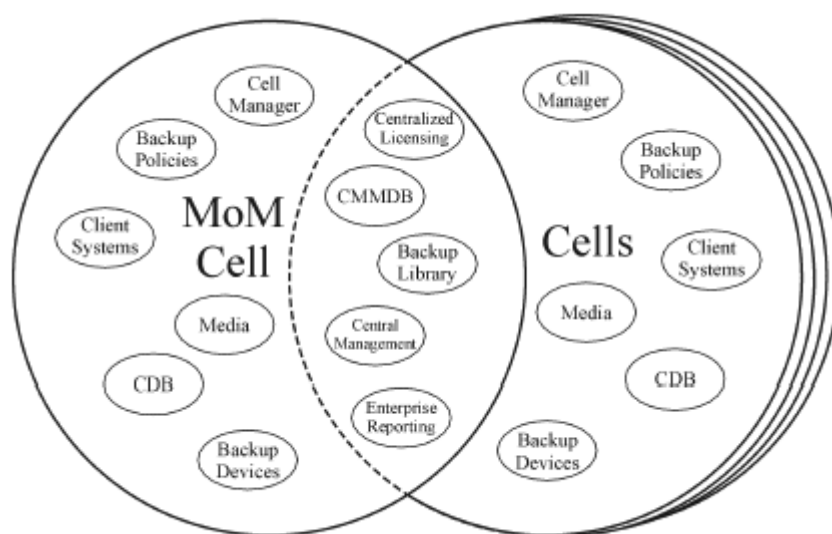
Data Protector allows the administrator to manage multiple cells from a single point.



Data Protector provides the Manager-of-Managers to manage large environments with multiple cells. The MoM allows you to group multiple cells into a larger unit, called a MoM environment that can be managed from a single point, as shown in

Figure 1-9. The MoM enables virtually unlimited growth of your backup environment. New cells can be added or existing ones split.

A MoM environment does not require a reliable network connection from Data Protector cells to the central MoM cell, because only controls are sent over the long distance connections, and backups are performed locally within each Data Protector cell. However, this is based on the assumption that each cell has its own media management database.



Manager-of-Managers provides the following features:

- **Centralized licensing repository** — This enables simplified license management. This is optional but useful for very large environments.
- **Centralized Media Management Database (CMMDB)** — The CMMDB allows you to share devices and media across several cells in a MoM environment. This makes devices of one cell (using the CMMDB) accessible to other cells, which use the CMMDB. The CMMDB, if used, must reside in the MoM cell. In this case, a reliable network connection is required between the MoM cell and the other Data Protector cells. Note that it is optional to centralize the Media Management Database.
- **Sharing libraries** — With the CMMDB, you can share high-end devices between cells in the multi-cell environment. One cell can control the robotics, serving several devices that are connected to systems in different cells. Even the Disk Agent to Media Agent data path can go across cell boundaries.
- **Enterprise reporting** — The Data Protector Manager-of-Managers can generate reports on a single cell basis as well as for the entire enterprise environment.

## EBS with VERITAS Backup Exec Overview

Seagate Software—the original developers of Backup Exec—and HP have worked closely together since 1997 to design a new method for protecting enterprise-wide corporate data. Seagate and HP worked to develop the solution that provides both the performance advantages of locally attached tape storage and the security and management advantages of centralized data protection. After Seagate Software was purchased by VERITAS in 1999, HP continued its collaboration with a new partner.

The Shared Storage Option (SSO) provides robust, centralized backup within a distributed environment. When using the SSO, multiple distributed backup servers can share common, centralized storage devices for better cost efficiency and reliability. At the same time, the SSO enables each backup server to initiate its own backups and restores, regardless of the state of the other backup servers. This functionality eliminates the single point of failure found in solutions that offer dedicated backup servers. VERITAS provides many additional features designed to increase fault tolerance and reliability.

For Backup Exec environments, the required components are:

- Windows NT 4.0 with Service Pack 6, Windows 2000 Server or Advanced Server with Service Pack 1, NetWare 4.2 with Support Pack IWSP8A, NetWare 5.0 with Support Pack 5, or NetWare 5.1 with Support Pack 1
- VERITAS Backup Exec 8.5 or 8.6 for Windows NT/Windows 2000 and VERITAS Backup Exec 9.0 or later for NetWare
- The SSO, which enables backup device sharing across multiple servers

Each server on the SAN must have these components installed to access the shared tape devices.

## Key Features

The Windows NT/Windows 2000 and NetWare versions of Backup Exec share several key features:

- Multidrive concurrent backup using parallel streaming
- Media rotation wizard that easily creates media rotation schemes, implementing either the default Grandfather-Father-Son rotation strategy or custom strategies to meet specified criteria
- Support for centralized multiserver backups over a comprehensive SAN
- Centralized administration
- Data recovery using Intelligent Disaster Recovery
- Local and global monitoring and management
- Pre- and post-scheduling of user-defined commands before or after backup or restore
- Support for removable storage devices
- Data verification options using cyclic redundancy check (CRC)
- Software or device-driven data compression
- Integrated virus protection
- Open and skipped file processing to provide skip, wait, and open file backup capability
- Disk and catalog grooming to free hard disk space of inactive files
- Global Data Manager support for monitoring Backup Exec servers from Global Data Manager's centralized console
- Command line interface to create backup jobs to protect an entire drive, run previously created backup/restore jobs, perform drive utility/operations, and obtain job status
- Support for backup to disk with Backup Exec 8.6
- Support for Intelligent Image Option with Backup Exec 8.6 to provide image-based backup with the ability to restore the full image, directories, or files
- Ability to read ARCserve tapes

In addition, Backup Exec for Windows NT/Windows 2000 has the following feature:

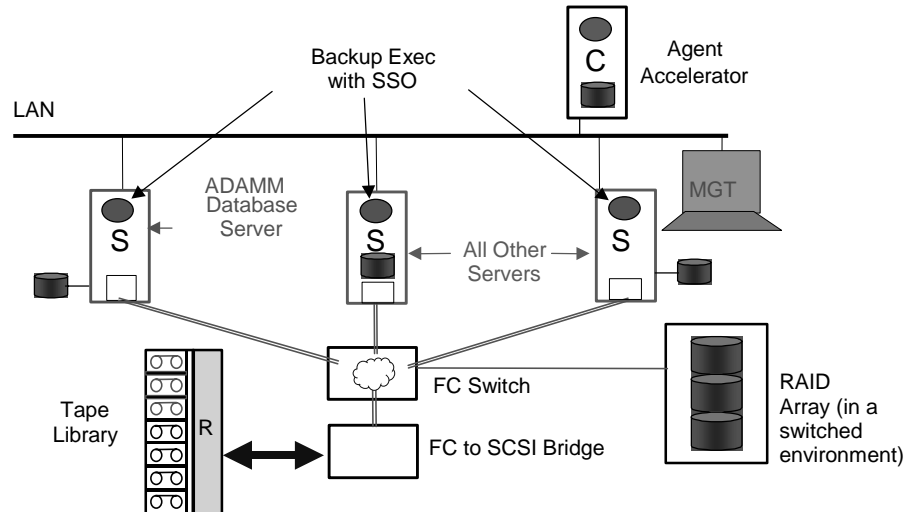
- Seamless compatibility with Windows NT/Windows 2000 native backup applications

In addition, the HP StorageWorks Enterprise Volume Manager(EVM) and Data Replication Manager(DRM) are supported for Windows NT/Windows 2000. EVM is a web-based application software that manages controller-based clone and snapshot operations. DRM provides remote disk mirroring by connecting a primary storage system to a remote primary storage system.

## Backup Exec Topology

A typical EBS solution with Backup Exec consists of the following components:

- Any combination of x86-based servers running Windows NT/Windows 2000 or NetWare (up to 32 servers)
- Host bus adapters
- SAN switches
- NSRs to accommodate tape drives
- Up to 32 tape drives for Windows NT/Windows 2000 and up to 27 tape drives for NetWare
- RAID arrays, including XP storage platforms, Enterprise Virtual Array, MSA1000, and HSG80-based storage platforms
- Backup Exec and appropriate client licenses (one database server)



The first server in the data zone where Backup Exec is installed will be the server with the Advanced Device and Media Management (ADAMM) database. This database contains information such as media sets and retention period. Other servers that access the tape libraries in the data zone also function as backup servers. All servers which access the tape library must have the Shared Storage Option installed.

## EBS with Computer Associates ARCserve Overview

Computer Associates Brightstor ARCserve delivers high performance, easy administration, broad device compatibility, and unsurpassed reliability to protect the computing environment. The suites provide enhanced functionality, including broad client/server support, policy-based data management, online backup of applications, fast disaster recovery, advanced media management, and Fibre Channel SAN support.

ARCserve 9 and ARCserve 7 enable you to fully exploit the HP Fibre Channel technology, offering numerous benefits to the enterprise, including:

- High-speed backup and restore over a fast, fiber optic storage backbone
- Flexibility of extended distances for storage devices
- Total scalability of the storage network
- Support for virtual libraries in Windows NT/Windows 2000 to provide dynamic load balancing, device fault tolerance, and complete statistical information
- Support for centralized multiserver backups over a comprehensive SAN
- Intelligent compression for both client and server, based on device capabilities
- Tape scan and byte-by-byte verification in addition to CRC
- RAID Levels 0, 1, and 5 for maximum fault tolerance and performance in a multidrive backup system for ARCserve 9
- Support for using tape drives in multiple libraries as an array to achieve the performance of the tape array
- Data recovery using the ARCserve Disaster Recovery option
- Online image backup for higher performance
- Flexible device drive selection using SCSI pass-through
- Full or individual value-level backup of local/remote Windows Registries


One important feature of ARCserve 9 is the RAID support for tape drives. By using RAID 1 or 5, it provides full protection even if one of the tape drives in a virtual library should fail (hardware fault tolerance) or if one of the tapes in a set is damaged or lost (media fault tolerance). RAID is not supported with ARCserve 7 for NetWare.

By using RAID 0, ARCserve 9 provides concurrent streams of data to the drives without generating parity information.

For ARCserve 9 or ARCserve 7 environments, the minimum required components are:

- Windows NT 4.0 with Service Pack 6a, Windows 2000 Server or Advanced Server with Service Pack 1, NetWare 4.2 with Support Pack IWSP8A, NetWare 5.0 with Support Pack 5, NetWare 5.1 with Support Pack 1
- ARCserve 9 for Windows NT/Windows 2000 and ARCserve 7 for NetWare
- The SAN option and Tape Library Option (TLO), which enables backup device sharing across multiple servers

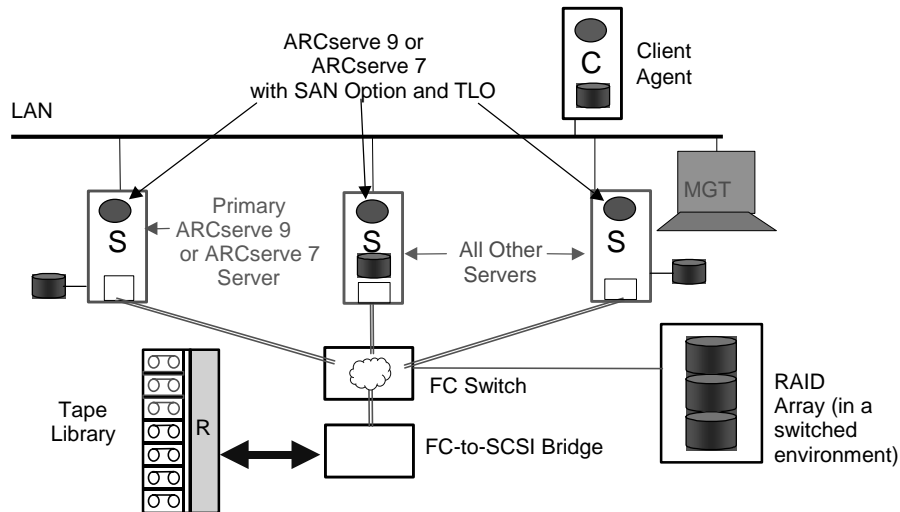
## ARCserve Topology



A typical EBS solution using ARCserve 9 or ARCserve 7 consists of the following components:

- Any combination of x86-based servers running Windows NT/Windows 2000 or NetWare (up to 20 servers)
- Host bus adapters
- SAN switches
- NSRs to accommodate tape drives
- Up to 20 tape drives for NetWare and up to 32 tape drives for Windows NT/Windows 2000
- RAID arrays, including XP storage platforms, Enterprise Virtual Array, MSA1000, and HSG80-based storage platforms
- ARCserve and appropriate client licenses (one primary server)





During a setup process, you can configure any of the storage servers attached to the SAN—typically the most powerful server—as the primary ARCserve 9 or ARCserve 7 server that will coordinate the communication and requests of Fibre Channel storage devices from other servers.

ARCserve 9 and ARCserve 7 also enable multiple servers to work together as a group. This enables Fibre Channel fabric support, which normally involves multiple groups with a fiber connection.

Because the primary ARCserve 9 or ARCserve 7 server monitors and controls access to the Fibre Channel devices, you are provided with some level of load balancing across the storage devices in the data zone. All servers in the data zone detect the same storage devices and are represented as local storage devices.

When a backup or restore request is made, the request is sent through the primary server, which traffics and queues the requests. After control and access is granted, it then treats the storage devices as if they were local devices, exclusive to that server.

The data path for backing up or restoring to an ARCserve 9 or ARCserve 7 server is exclusively through the fiber interconnect, whether it is centrally or remotely managed.

## EBS with VERITAS NetBackup Overview

VERITAS NetBackup 4.5 boasts many industry-first features that continue to provide a Data Center EBS solution. With support for Windows NT/Windows 2000, Sun Solaris, HP Tru64 UNIX, IBM AIX, and HP-UX, NetBackup is used to protect mission-critical enterprise data. NetBackup has a four-tier architecture, advanced media management, and SAN support.

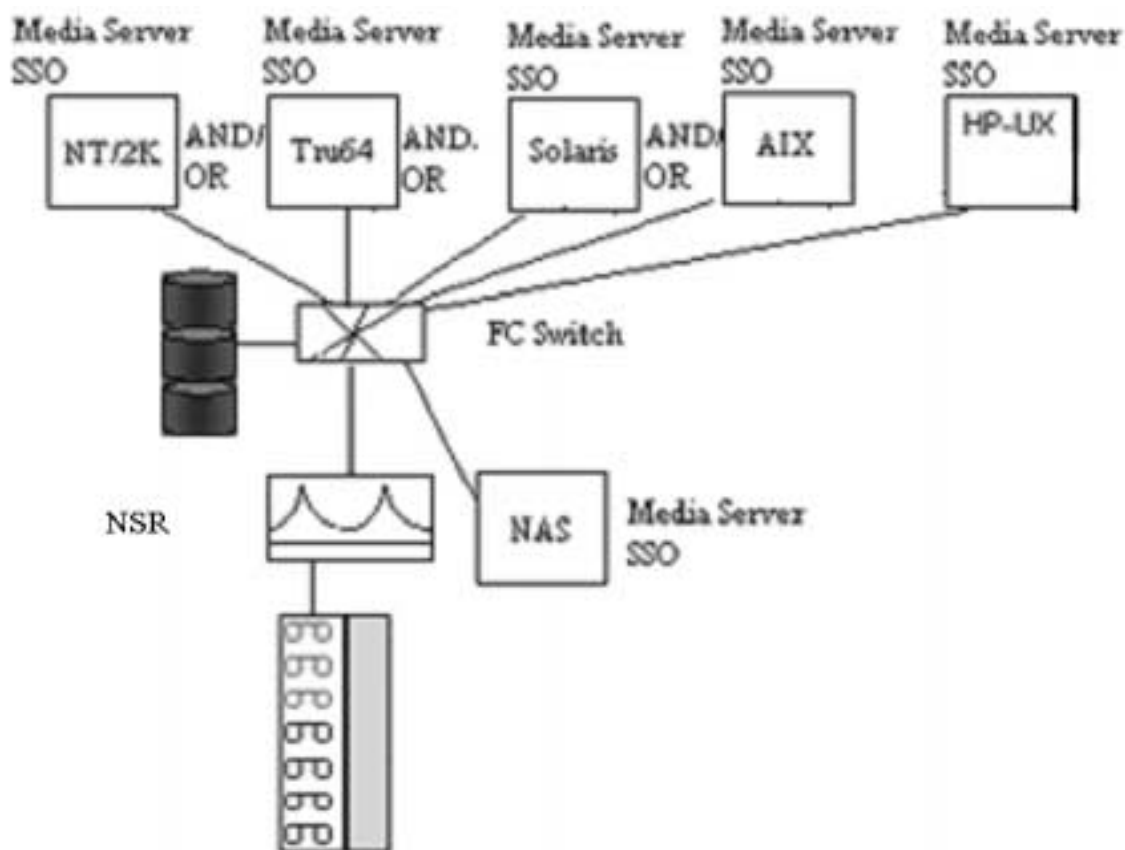
For NetBackup environments, the minimum required components are:

- Windows 2000, Windows NT 4.0 with Service Pack 5, Tru64 UNIX 5.x, Sun Solaris 2.6 ,7, or 8, IBM AIX 4.3.3, or HP-UX 11.x
- NetBackup DataCenter 4.5 for each Operating System with Multihost license

### Key Features

- **Dynamic tape drive sharing** – NetBackup allocates unused tape drives dynamically to servers requesting use and then deallocates tape drives when backup sessions are complete to allow other servers to use the tape drives.
- **Multiplexing backup and restore** – NetBackup can run multiple backups simultaneously and stream the data to one or more devices. Multiplexing can dramatically increase performance.
- **Checkpoint restart of backup sessions** – In the event of a backup failure, NetBackup can restart and execute only the portion of the backup that failed. Checkpoint restart is not supported with restores.
- **Non-proprietary tape format** – NetBackup has the ability to create tar-compatible tapes
- **True image recovery** – NetBackup can be configured to restore only current files. Files that have been deleted can be optionally prevented from being restored, reducing the total number of files to restore in the event of disaster recovery.
- **Intelligent disaster recovery** — Fully integrated disaster recovery for Windows NT environments allows customers to recover remote Windows NT servers quickly, without a lengthy reinstall of the operating system. NetBackup also can copy primary backup tapes and then de-multiplex these tapes so that data is co-located on the second tape, thereby greatly reducing recovery times for critical data or single file restores.
- **Shared Storage Option (SSO)** — Using the SSO, individual tape drives can be shared between multiple NetBackup servers. The SSO enables data centers to leverage their peripheral investments more thoroughly through drive sharing. .
- **Non-disruptive backup** — The Open Transaction Manager provides integrated, open-file backup for non-disruptive backup of Windows NT/Windows 2000 servers.

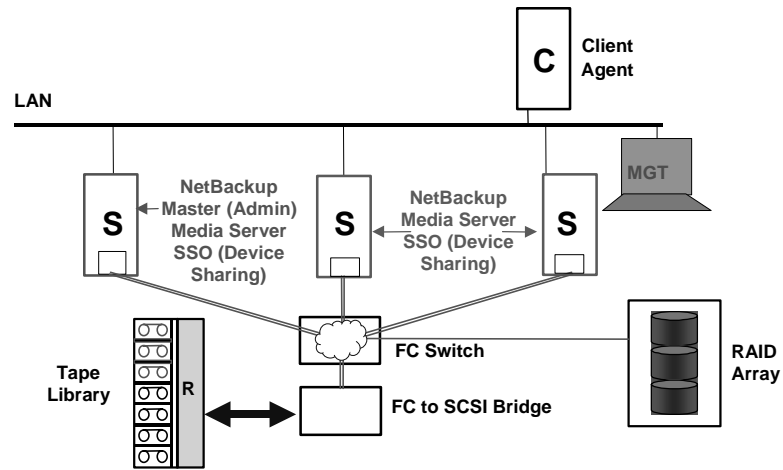
## NetBackup Topology



A typical EBS solution using NetBackup consists of the following components:

- Any combination of x86-based servers running Windows NT/Windows 2000, HP Alpha Servers running Tru64 UNIX, HP-UX servers, Sun UltraSPARC servers running Solaris, or IBM AIX servers (up to 32 servers)
- Host bus adapters
- Fibre Channel switches
- Up to 32 tape drives
- NSRs to accommodate tape drives
- RAID arrays, including XP storage platforms, Enterprise Virtual Array, MSA1000, and HSG80-based storage platforms
- NetBackup DataCenter 3.4 and Multihost licenses

In a switched environment, both primary and secondary storage can be located on the same switched fabric.



VERITAS NetBackup DataCenter v.4.5 provides a four-tier architecture that, combined with advanced media management and speed, can address the largest data center installations.

The first tier consists of the NetBackup Master server, which controls actions such as scheduling and tracking client backups. The Master server can have one or more tape devices or tape libraries attached for backing up data from multiple clients.

Second tier Media servers provide local backup of large applications while backing up other clients over the local area network. A Media server can share a tape library with the Master server or another Media server, or the Media server can operate with its own tape devices or tape library. If a Media server fails, the attached client's backups can be routed to another Media server.

The third tier consists of client agents that back up servers and workstations over local area networks. Both Media servers and clients can be centrally managed from the Master server.

The fourth tier uses Global Data Manager for NetBackup. A NetBackup storage domain consists of one NetBackup Master server and one or more Media servers. Global Data Manager offers centralized management and control of all NetBackup storage domains in the enterprise.

## EBS with NetWorker Overview

Legato NetWorker 6.1 offers heterogeneous platform support, automated media handling, interoperable tape format, data stream parallelism, and remote tape management.

For NetWorker environments, the minimum required components are:

- Windows NT 4.0 with Service Pack 6, Windows 2000, Tru64 UNIX 4.0f or 5.x, Sun Solaris 2.6, 7, or 8.0, IBM AIX 4.3.3, Red Hat Linux 6.2, 7.1, or SuSE Linux 7.0
- NetWorker 6.1.x for each Operating System
- (Optional) AlphaStor, which supports device sharing across multiple data zones

### Key Features

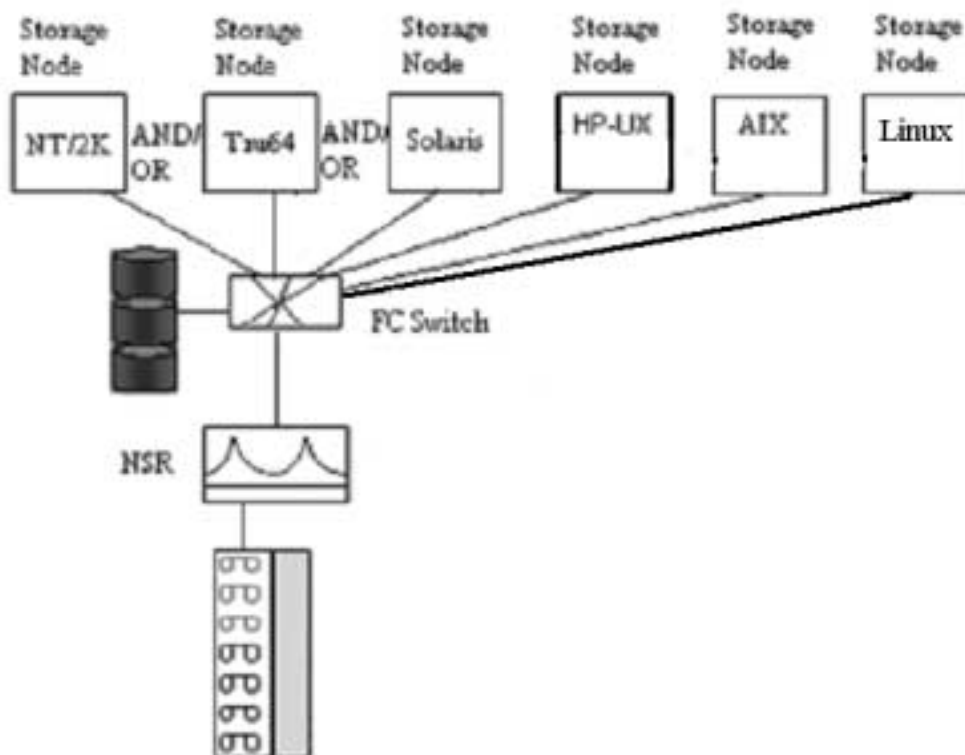
- **Heterogeneous support** — NetWorker protects data running on Windows NT, Windows 2000, Solaris, Tru64 UNIX, Red Hat Linux, and SuSE Linux servers
- **Single administrative view** — NetWorker presents a single administrative view that allows one to configure NetWorker servers running various operating systems
- **Active Client** — The client software responds to backup requests, generates the appropriate data stream, and pushes the data to the server. For example, if an incremental backup is requested, the client selects the modified files. This reduces the amount of communication that occurs between the client and the server.
- **Multiplexing** — NetWorker servers can process multiple client data streams in parallel as well as interleave data streams to a single or multiple devices.
- **Point-in-time recoveries** -- NetWorker provides an index of multiple versions of files and allows the selection of a precise version for recovery.
- **Open file backup** — Add-on modules enable open files and popular databases and applications, including Oracle, SAP R/3 on Oracle, Informix, Sybase, Microsoft SQL Server, Microsoft Exchange Server, and Lotus Notes.
- **Security** — NetWorker provides data encryption and password protection to protect valuable data.
- **Failover support** — Failover capability based on the storage node priority list increases backup reliability and availability.
- **Windows 2000 Cluster Server** — Legato NetWorker and Legato SmartMedia ensures that Microsoft Cluster Server (MSCS) clusters are backed up, while MSCS ensures that online data will be backed up even if the primary server in the cluster fails.
- **NAS switched fabric support** — NAS appliances are specialized file serving devices that provide support for heterogeneous files in a high

capacity package. Customers can deploy their backup of the TaskSmart N2400 server through the SAN.

- **Dynamic Drive Sharing** — Networker 6.1.x allocates unused tape drives dynamically to servers requesting use and then deallocates tape drives when backup sessions are complete to allow other servers to use the tape drives.

## NetWorker Topology

Note: One of the Storage Nodes is the NetWorker server.



A typical EBS solution using NetWorker consists of the following components:

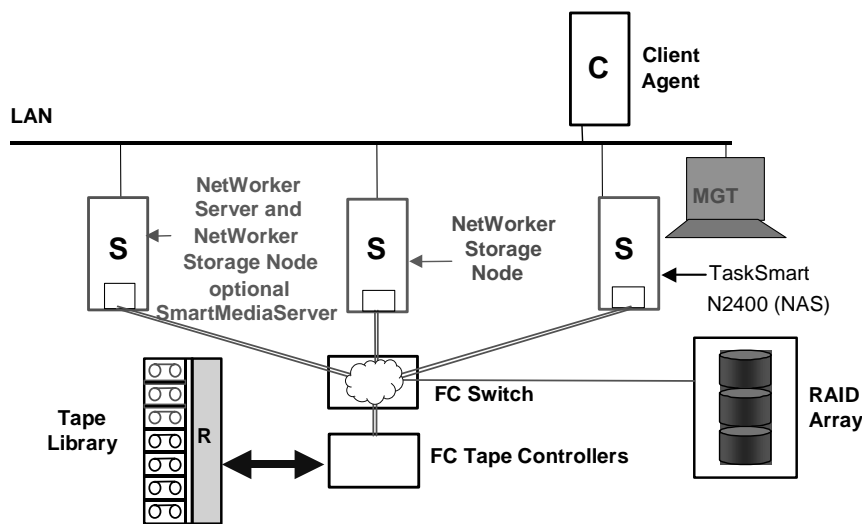
- Any mixture of HP Tru64 UNIX, Sun Solaris, HP-UX, IBM AIX, and x86-based servers running Windows NT/Windows 2000 or Linux
- Host bus adapters
- Fibre Channel switches
- Up to 32 drives
- MDRs or NSRs to accommodate drives
- RAID arrays, including XP storage platforms, Enterprise Virtual Array, MSA1000, and HSG80-based storage platforms

- A NetWorker server, a Legato AlphaStor server (optional), and NetWorker storage nodes

### Note

The AlphaStor server is only required if you want to share libraries between data zones for NetWorker.

In a switched environment, both primary and secondary storage can be located on the same switched fabric.



NetWorker is based on the client/server architecture and comprises three distinct components—client, storage node, and server. This three-tier architecture provides the flexibility and performance required to protect and manage data on the most complex networks. The data protection process is encompassed in a data zone—a collection of clients, storage nodes, and data protected by a single NetWorker server. A NetWorker data zone is collectively administered with a single set of data protection policies and procedures. Legato NetWorker excels as a data zone manager, delivering both centralized and distributed storage management.

The server piece must only be installed on the NetWorker server. The NetWorker server manages the index of backed up files and the index of volumes called the media index. The NetWorker server uses these indices to locate and recover data. Every other machine will be either a storage node or a client. The difference between a storage node and a client is that a storage node has its own dedicated tape device, which can be logical via the SAN, whereas a client can only backup to the NetWorker server or storage node via the LAN. In other words, every server with access to the SAN should be a storage node.

## NetWorker and AlphaStor

Legato NetWorker 6.1.x provides dynamic drive sharing support. The AlphaStor option provides enhancements to NetWorker, including:

- Drive sharing across multiple NetWorker data zones



- Web-based, centralized interface to manage all media, device, and library usage
- Support for multiple tape device types in the same library

## EBS with Tivoli Storage Manager Overview

Tivoli Storage Manager protects and manages data in a client-server environment with multiple operating systems in an enterprise-wide SAN and traditional network environment. Library Sharing allows multiple Tivoli Storage Manager servers to share the same tape library and drives in a SAN, improving backup and recovery performance and tape hardware asset utilization.

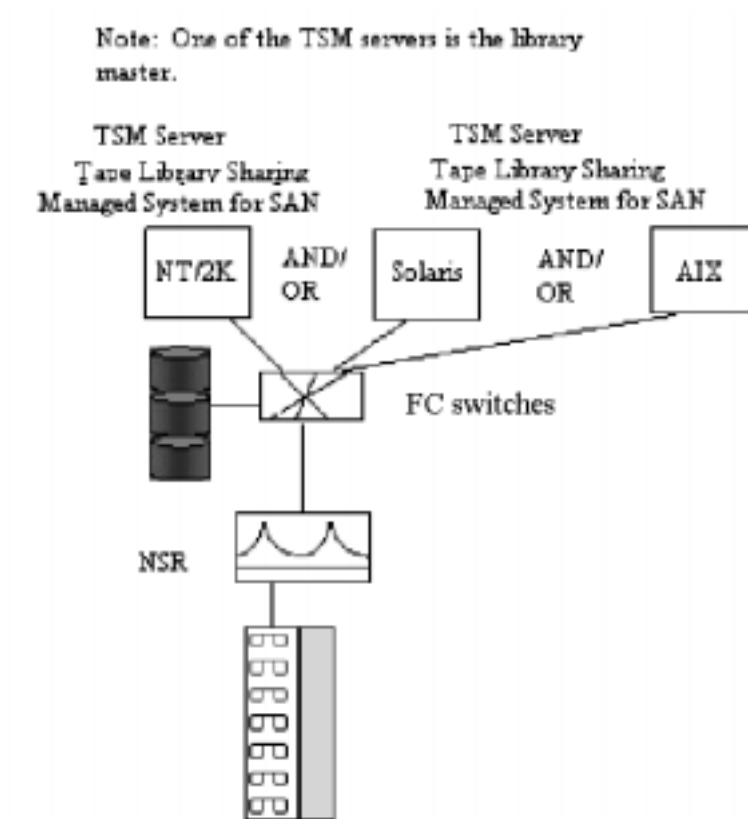
For TSM environments, the minimum required components are:

- Windows NT 4.0 with Service Pack 6a, Windows 2000 with Service Pack 2, Sun Solaris 7 or 8.0, or IBM AIX 4.3.3
- One TSM Managed System for LAN feature is required for each system that will move data across the LAN.
- One TSM Managed System for SAN feature is required for each TSM server that will move data across the SAN.
- One Tape Library Sharing feature is required for each server that accesses a shared library
- One TSM Managed Library feature is required at the library master for each server that the library master will manage

### Key Features

- **Centralized administration** – Tivoli Storage Manager's user-friendly enterprise console enables the administrator to manage any platform in the enterprise via a Web-based interface from anywhere in the enterprise.
- **Progressive backup methodology** – Tivoli Storage Manager saves time and disk space by backing up only new files and modified files. The progress backup feature uses its own database to track data wherever it is stored, delivering direct one-step restore. The progress backup eliminates the need for base plus incremental file restores.
- **Dynamic multi-threaded transfer** – Tivoli Storage Manager permits a client to transfer data using multiple sessions to and from the Tivoli Storage Manager server. This boosts performance over a single-threaded session.
- **Adaptive Differencing Technology** – Tivoli Storage Manager will dynamically transfer data at a byte, block, or file level, based on data size.
- **Tape Library Sharing** – Tivoli Storage Manager allows multiple Tivoli Storage Manager servers to utilize the same tape library and drives.
- **Configuration and policy management** – Tivoli Storage Management configuration and policy information can be defined one time at a configuration server and then propagated to any number of managed Tivoli Storage Manager servers.
- **Data protection for applications** – Using Tivoli Data Protection products, Tivoli Storage Manager offers data protection for most popular groupware, e-mail, databases, and applications.

## Tivoli Storage Manager Topology

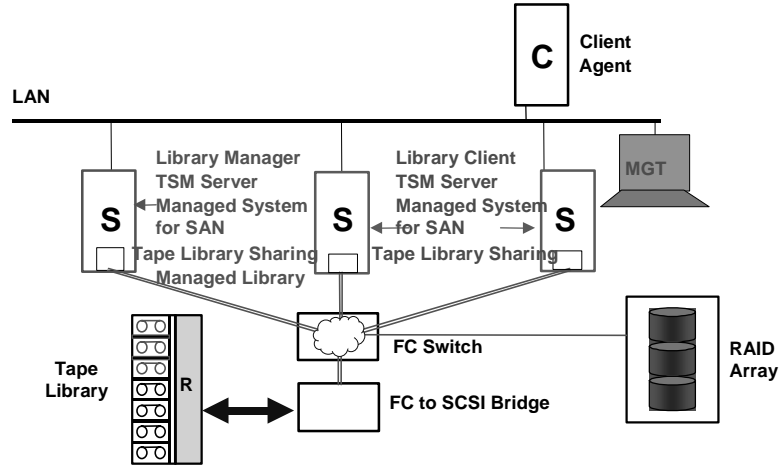


A typical EBS solution using Tivoli Storage Manager consists of the following components:

- Any mixture of Sun Solaris servers, x86-based servers running Windows NT/2000, and IBM AIX servers (up to 16 total)
- Host bus adapters
- Fibre Channel switches
- Up to 16 drives
- MDRs or NSRs to accommodate drives

- TSM servers, one of which is the library master
- HSG80-based RAID arrays

In a switched environment, both primary and secondary storage can be located on the same switched fabric.



When Tivoli Storage Manager servers share a library, one server is the **library manager** and controls the library operations. Other servers, **library clients**, use server-to-server communications to contact the library manager and request library service. The library client can only perform operations to the drives after the library manager grants the library client access to the drives.

## EBS with CommVault Galaxy

CommVault Galaxy serves organizations with data center-sized backup requirements. The solution integrates data protection and archival strategies with disk storage subsystems across multiple hardware platforms and operating systems located on the same SAN. This solution provides for the interconnection of multiple heterogeneous servers to multiple tape backup devices using dynamic device sharing.


For Galaxy environments, the minimum required components are:

- Windows NT 4.0 with Service Pack 5, Windows 2000, or Sun Solaris 8.0
- One server with CommServe StorageManager
- MediaAgent and iDataAgent for each server that shares the library across the SAN

### Key Features

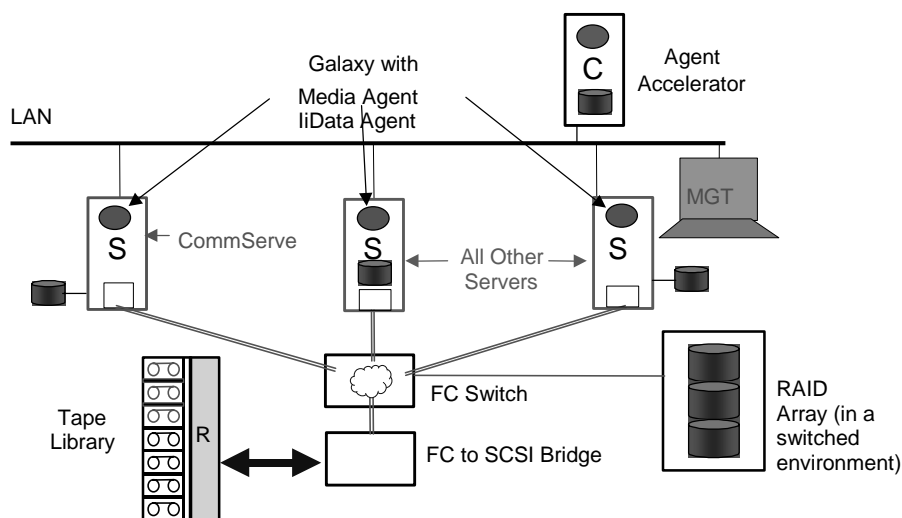
- **Unified management console** – Galaxy utilizes a Java- based, web accessible Graphical User Interface(GUI).
- **Two-part indexing scheme** – Galaxy’s two-part indexing scheme makes fast browses and restores of data possible. The backup index is stored on the backup media as well as each MediaAgent’s local drive. The indexing scheme also shields administrators from having to understand the physical locations of data.
- **Comprehensive storage policies** – Galaxy uses storage policies to map the backup data to a physical media location. Storage policies can include data streams, media groups, and retention periods.
- **Synthetic full backup** – Galaxy has the ability to create a full backup from the last full backup and subsequent incremental or differential backups. This synthetic full backup can be written to the same or different media. The creation of the synthetic full backup does not create traffic with clients.
- **Auxiliary copy** – Galaxy lets users create multiple copies of data for better reliability. If primary copies are in use or unavailable due to media failure, the auxiliary copy is automatically substituted to fulfill the outstanding request. Auxiliary copies can be written to different libraries.
- **Application specific data agents** – Agents are available for Microsoft Exchange, Lotus Notes, Oracle, and Microsoft SQL Server.

## CommVault Galaxy Topology



A typical EBS solution using CommVault Galaxy consists of the following components:

- Any mixture of Sun Solaris, HP-UX, HP Tru64, and x86-based servers running Windows NT/Windows 2000; a Windows NT/Windows 2000 server must be present to manage the configuration
- Host bus adapters
- Fibre Channel switches
- MDRs or NSRs to accommodate tape drives



CommVault Galaxy is built utilizing an advanced multi-tier architecture using combinations of CommServe StorageManager, iDataAgents, and MediaAgents. Each of the combinations forms a CommCell.

The CommServe StorageManager is a central management point for controlling interaction between the iDataAgents and MediaAgents within the CommCell. Only control information, not backup data, passes through the CommServe Storage Manager.



#### Note

The CommServe StorageManager must be installed on a Windows NT/Windows 2000 server.

The MediaAgent manages the movement of data between the backup storage devices and the corresponding iDataAgents at the clients.

The iDataAgent is a software module specific to the type of data being managed. An iDataAgent is required for each managed data type per client computer. Examples of iDataAgent include the Microsoft Exchange iDataAgent and the NTFS iDataAgent.

## Component Summary

For the latest information on qualified hardware and software components , see the hardware and software compatibility matrix at

<http://www.hp.com/products/storageworks/ebs/EBScompatmatrix.html>



## Current EBS Caveats

- Primary and secondary storage cannot be on same loop when using the Storage Hub 7 or Storage Hub 12; they can be supported using an Arbitrated Loop Switch or on a switched fabric.
  - When using the HP Fibre Channel Host Controller with an Arbitrated Loop Switch and Windows NT/Windows 2000 or NetWare, both primary storage (RA4x00) and secondary storage can be on the same loop switch.
  - The HP Fibre Channel Host Controller provides full switched fabric support for Windows NT and NetWare, allowing both primary storage (RA4x00) and secondary storage on the same fabric. The controller also supports fabric support for HSG80-based storage systems and secondary storage on the same fabric.
  - RA4x00 storage systems and HSG80-based storage systems are not supported on the same SAN.
- Each bus of the Modular Data Router and of the Network Storage Router is internally terminated—no external termination is necessary.
- In an EBS configuration, a maximum of sixteen tape drives are supported per a single MDR, with eight SCSI buses, (two quad-port SCSI cards). This maximum configuration is dependent on Windows NT 4.0's and Windows 2000 Advanced Server's limitation of discovering LUNs at location eight and higher.
- When configuring the library for use with a backup application, the software configuration **must** match the actual hardware configuration of the library.

Inconsistency between the hardware and software configurations can result in:

  - Unreadable media errors.
  - Full slots reporting as empty.
  - Lock-up or hanging during tape inventory.
  - Broken tapes.
  - Dropped leaders on the DLT drives.
- Name resolution is used by most ISV software. If DNS or hosts file is not configured correctly, backup software may have problem communicating with servers.

## Computer Associates ARCserve

The solution is homogeneous in both the FC-AL and FC-SW environments.

## **VERITAS NetBackup**

- Purchase and installation of software through VERITAS.
- The SSL2020TL AIT library can only be used in a homogeneous Windows NT/Windows 2000 or homogeneous UNIX environment.
- No device persistency for Window NT/Windows 2000. Device persistency must be maintained through SAN configuration management. A device must remain in the same switch port at the time of software configuration. Device persistency for UNIX is managed by the Operating System through worldwide name binding.
- Hardware installation per Reference Guide and Component User Guides.

## **Legato NetWorker**

- Purchase of enabler keys to activate software through Legato.
- No device persistency for Windows NT/Windows 2000. Device persistency must be maintained through SAN configuration management. A device must remain in the same switch port at the time of software configuration. Device persistency for UNIX is managed by the operating system through worldwide name binding.
- Hardware installation per reference guide and component user guides.

## **Tivoli Storage Manager**

- When using the IBM 3584 LTO tape library, a Windows 2000 server cannot be used as the Library Manager.
- When using the IBM 3584 LTO tape library, the TSM Device Driver must be disabled. The IBM Ultrium device drivers must be installed instead.

## Learning Check

1. Service Pack 6 for Windows NT is a requirement for the Backup Exec solution.  
☐ True  
☐ False
2. In EBS, the Backup Exec servers communicate with the \_\_\_\_\_ server over the LAN to arbitrate for the right to access each of the backup devices.  
.....
3. What is the name of the Backup Exec option that enables servers to share backup devices?
  - a. Enterprise Library Option
  - b. Shared Storage Option
  - c. Autoloader Option
  - d. Fibre Channel Loop Option
4. With the Shared Storage Option in Backup Exec, each backup server on the storage network must be ready to perform a backup before any backups can begin.  
☐ True  
☐ False
5. What is the largest number of servers that can be supported in a maximum configuration of the Enterprise Backup Solution running Legato NetWorker?
  - a. 7
  - b. 8
  - c. 10
  - d. 16
6. What is the largest number of servers that can be supported in a maximum configuration of the Enterprise Backup Solution running Computer Associates ARCserve 9?
  - a. 14
  - b. 16
  - c. 20
  - d. 32

7. Which operating systems offer support for device persistency?
  - a. Windows NT/Windows 2000
  - b. Sun Solaris
  - c. HP Tru64 UNIX
  - d. None of the above



---

# Designing the Multiserver Backup Solution


## Module 7

### Objectives

After completing this module, you should be able to:

- Apply MB/s to GB/hr conversion formulas.
- Perform a customer-needs analysis to collect information necessary for an appropriate design.
- Design a multiserver backup solution based on a customer's requirements.
- Explain the purpose of the hp StorageWorks Backup Sizing Tool and use the tool to design an EBS.
- Explain how various StorageWorks software tools are used to facilitate backups.
- Explain cluster backup and restore strategies.

## MB/s to GB/hr Conversion Formulas



Tape drive capacities and throughput figures are often represented in gigabytes. Primary storage throughput figures are typically represented in megabytes. In the world of enterprise storage, it is important to learn the formulas for converting the megabyte term usage to gigabyte terminology. The preceding graphic demonstrates how to calculate this equation.

### Example

$$5\text{MB/s} \times 3.6 = 18\text{GB/hr}$$

$$18\text{GB/hr} \times 16.66 = 299.88\text{MB/min}$$

$$299.88\text{MB/min} \div 60 = 4.998\text{MB/s}$$



## Performing a Needs Analysis

Before you can begin the Enterprise Backup Solution (EBS) design process, you must have a thorough understanding of the network, the type of data to be backed up, and the backup window parameters.

The best way to familiarize yourself with the network and its data is to perform a needs analysis. From the needs analysis, you will have all the information necessary to design an EBS.

The needs analysis consists of three steps:

1. Create a network topology map.
2. Identify the backup window.
3. Categorize the importance of data.

## Creating a Network Topology Map

The network topology map should contain information for each of the following:

- Network
- Servers
- Volumes

To develop the topology map and determine your customer's needs, ask the following questions.

### Network

**How fast is the LAN?** Backup speed is generally limited to the capabilities of the LAN as well as the I/O capabilities of the server. This speed has a direct correlation with the backup window.

LAN Type	Maximum Speed	Typical Speed
10Base-T	3.6GB/hr	2GB/hr
100Base-T	36GB/hr	15–20GB/hr
FDDI	Similar to 100Base-T	Similar to 100Base-T
Gigabit Ethernet	360GB/hr	36GB/hr-100GB/hr
Fibre Channel	720GB/hr	500GB/hr

### Servers

- **What operating systems are found on the servers?** The version of backup software used depends on the operating system running on the network.
- **Which disk controller are you using?** When determining (or selecting) the disk controller to be used in the backup system, keep future needs in mind. The rate at which a controller can read data from the drives and transfer it to the tape library will determine how many and what type of controllers are required.

## Volumes

- **What is the feed speed of the hard disk system?** This speed will help determine how quickly a volume can be backed up.
- **What is the data compression ratio for the volume?** The type of data contained on the volume greatly impacts compressibility, backup speed, and the number of tapes required to hold the data. Typical compression ratios are shown in the following table and used by the

Data Type	Compression Ratio
Typical file/print server	2:1
CAD	3.8:1
MPEG/JPEG	1:1
Spreadsheet/word processing	2.5:1
Oracle/SAP databases	1.2:1
Microsoft Exchange/SQL databases	1.4:1
Lotus Notes database	1.6:1

---

### Note

These figures are typical and were not tested by HP. Your results might vary, depending on the actual data being compressed.

---

When compressed data is passed through a compression process a second time, it typically occupies more space rather than less because of the overhead of the compression algorithm. This fact can be important when calculating the number of tapes needed.

- **How large are the files?** As file size decreases, so does the tape backup rate because of the overhead of cataloging a higher number of files.

---

### Note

All performance information contained herein is the result of using 1GB files during test procedures.

---

**What is the expected data growth rate?** You must consider the present data volume and anticipate the growth rate to arrive at the best storage solution.

The SSL2020 AIT Library is a single or dual drive library that offers 1TB (native) of storage. This library also offers scalability and flexibility as storage needs grow. The library modules can be stacked five high to give the customer up to 10TB of storage within a 20U space. By stacking the modules the customer can build a 10-drive solution that holds 100 data cassettes, which gives them higher capacity and greater performance.

MSL5000 libraries can be stacked to accommodate further growth.

ESL9000 libraries can be scaled together.

## Identifying the Backup Window

When identifying the backup window, ask these questions:

- **Can the volume to be backed up be taken offline?** Some volumes contain data that must remain available at all times. For example, the database for an online store cannot be taken offline during backup or customers will be unable to make purchases during that time. Plans need to be made to keep this data available during backups. Possibilities include using tools such as the StorageWorks Virtual Replicator or Enterprise Volume Manager (discussed later in this module).
- **How large is the backup window?** With a Fibre Channel system, customers now have access to the information they require, even during backups. Network bandwidth is not reduced during backups because there is a direct path between the server and the tape backup solution. Depending on the backup software being used, however, they might not be able to access files that are currently being backed up.

If immediate and constant access to files is critical, your customers might need to plan for a dedicated backup window. In this case, the configuration of the solution will be based partly on the acceptable backup window.

## Categorizing Data Importance

When categorizing the importance of data, ask these questions:

- **Which data is business-critical?** Business-critical data, such as sales transactions, have precedence over user data. Volumes containing this information should receive preference during the backup window.
- **What types of backups will you want?** The importance of the data also helps determine what types of backups you perform. The type of backup undertaken affects not only the speed of the backup and restore process, but also the amount of data being retained on tape. There are three types of backup, differing in regard to whether they record all the data or only changes made since previous backups:
  - **Full backups only** — All data will be backed up. Backups take longer, but recoveries are faster.
  - **Full and incremental backups** — Incremental backups record the changes since the last backup (full or incremental). Backups are faster, but recoveries take longer because all incremental backups must be restored.
  - **Full and differential backups** — Differential backups record all the changes since the last full backup. Backups take longer than incremental backups, but recoveries are faster (and safer) because all data is restored from only the most recent backup.
- **What is the tape retention schedule?** The importance of the data determines how long backup sets are maintained. Tape retention parameters that must be considered include:
  - Type of backups.
  - How partial backups are retained (if applicable):
    - ◆ On the same tape set as the corresponding full backup.
    - ◆ On a tape set separate from the full backup.
    - ◆ On its own tape set (each partial backup).
  - Whether partial backups will be differential or incremental. Incremental backups are typically smaller because they only back up data that has changed since the last backup job (partial or full).

- How long the partial backups are retained (if applicable):
  - ♦ As long as the full backup (required if partial backups are retained on the same tape set as full backups).
  - ♦ Until x subsequent full backups have been made (where x can range from one to the number of full backups to be retained).
- How many differential backups should be retained between full backups (if differential backups are indicated). Because a differential backup records all the changes since the last full backup, there is redundancy in the data backed up during consecutive partial backups. The value of having this redundancy must be weighed against the cost of the greater storage capacity required. Therefore, the optimal schedule for a particular customer might retain:
  - ♦ All differential backups since the last full backup.
  - ♦ No more than x differential backups (where x can range from two to the maximum number of differential backups between full backups).

## Designing a Solution

After you have completed the needs analysis, you are ready to begin the actual design of an EBS. This consists of five steps:

1. Create backup data domains.
2. Calculate the volume of data to back up.
3. Group the data into backup jobs.
4. Select the appropriate backup device.
5. Document your solution.`

### Creating Backup Data Domains

Create backup data domains or data zones, based on the backup software the customer is using. For example, if the customer is using VERITAS NetBackup, the maximum number of servers allowed in the data zone is 32. If the customer has 50 servers they want in the EBS, you will need at least two data zones.

### Calculating the Data Volume

Calculate the entire amount of data to be backed up for the network, each server, and each volume.

If necessary, you will create backup domains based on the backup window, data importance, amount of data, and capacity of backup devices.

### Grouping Data into Jobs

It is important to group data into different jobs. Doing so provides more data security in case of a tape drive or other failure during the backup job. If one drive fails, not all data is lost, just the data of the job for that particular tape drive.

Grouping data also provides higher backup performance. With one job per tape drive, several drives working at the same time can complete the backup much faster.

Load balancing is also enabled when the data is grouped into jobs. Using drive pools, the backup software selects the drives for the individual backup jobs, ensuring that no one drive is handling the majority of the backup.

## Selecting the Appropriate Device

After you have gathered the information from the needs analysis and grouped the data into jobs, you are ready to select a backup device. This process consists of four steps:

1. Calculate the compressed transfer rate.
2. Adjust the compressed rate for feed speed and lights-out operations.
3. Determine the number of drives needed to meet the backup window.
4. Determine the quantity of retained media.

### Step 1: Calculating the Compressed Rate

Begin with a 1:1 compression ratio. Then apply the compression formula:

#### For DLT Drives

For a 1:1 compression ratio:

Compressed rate = Base rate x compression ratio

For compression ratios greater than 1:1:

Compressed rate = (Base rate x compression ratio) – (compression ratio)<sup>2</sup>

#### Example

$$(15 \times 2) - 2^2 = 26\text{GB/hr}$$



---

**Note**

The maximum compressed rate for a 35/70 DLT drive is 43GB/hr. If the calculated compressed rate comes out to more than 43GB/hr, you must back this number down to 43.

---

#### For LTO and AIT Drives

Base rate x compression ratio = compressed rate

#### Example

$$20 \times 1 = 20\text{GB/hr}$$



---

**Note**

The maximum compressed rate for an AIT 50 drive is 45GB/hr. If the calculated compressed rate comes out to more than 45GB/hr, you must back this number down to 45.

---



**Step 2: Adjusting the Compressed Rate for Feed Speed and Lights-Out Operations**

- An internal SCSI controller can feed approximately 60GB/hr.
- HP SMART Array Controllers can feed approximately 60GB/hr to 200GB/hr.
- The Modular Data Router can feed approximately 100GB/hr (up to 16 drives). The Network Storage Router M2402 can feed approximately 600GB/hr.
- Tape drives can handle a 3:1 feed-to-write ratio. In other words, a DLT drive can write information at 15GB/hr, which requires a 45GB/hr feed rate for maximum performance. A 3:1 feed ratio uses 100% of the write capability of the drive; a 2:1 ratio reduces write speed to 80%; <2:1 to 50%.
  - If the feed speed is 3:1 over the compressed rate calculated in step 1, leave the compressed rate as is.
  - If the feed speed is more than 2:1 but less than 3:1, reduce the compressed rate by 20%.
  - If the feed speed is less than 2:1, reduce the compressed rate by 50%.
- If full backups are more than one tape capacity, reduce the compressed rate by 3%.
- If performing incremental or differential backups, reduce the compressed rate by 14%.

### Step 3: Determining the Number of Drives Needed to Meet the Backup Window

- The total data divided by the adjusted compressed rate equals the total hours of backup for one drive.  
$$\text{Hours of backup for one drive} = (\text{total data}) / (\text{adjusted compressed rate})$$
- The total hours for one drive divided by the backup window equals the number of drives needed.  
$$\text{Number of drives needed} = (\text{hours of backup for one drive}) / (\text{backup window})$$
- If performing full backups, the mechanical process of changing tapes could reduce the base backup rate by up to 7%.
- Incremental backups have slower backup rates than full backups because all the data on the system must be read to determine which files are new or have been modified since the last backup.

### Step 4: Determining the Quantity of Retained Media

1. Calculate the number of full and incremental or differential data sets to be retained.
  - Each tape holds up to 110GB (SDLT), 40GB (DLT), or 50GB (AIT) of data. As the data compression ratio increases, effective tape storage capacity increases. For example, at 1:1 compression, a tape can store 40GB of data; at 2:1 compression, it can store 80GB of data.
  - Total storage capacity is affected by the data compression ratio, type of backups, and so forth.
  - If your retention parameters demand more backups, you will require additional tapes, which could increase the number of libraries required.
2. Assign percentage values to each backup set. Use a full backup as the benchmark and assign it a value of 100%. Partial backup sets will have a smaller value, based on the amount of data that has changed or been added since previous backups.

If the company estimates that 10% of its data files are changed or added to in some way each day, an incremental backup would have a value, on average, of 10%. If they use differential backups, the value for each set will be 10% x the number of days since the last full backup.

3. Calculate the total retained percentage by calculating the value of the backup sets produced during the retention period.

**Example**

- If the company performs weekly full backups (F), no partial backups, and retains backup sets for three weeks before overwriting them, the total retained percentage would be:

$$F \times 3 = 100\% \times 3 = 300\%$$

- If the company performs weekly full backups, incremental backups (I) on the other four days of the week, and retains for three weeks, the total retained percentage would be:

$$(F + I + I + I + I) \times 3 = (100 + 10 + 10 + 10 + 10) \times 3 = 420\%$$

- If the company performs weekly full backups, differential backups (D) on the other four days, and retains sets for three weeks, the total retained percentage would be:

$$(F + D + D + D + D) \times 3 = (100 + 10 + 20 + 30 + 40) \times 3 = 600\%$$

4. Calculate the storage capacity of each tape:

Tape capacity (SDLT) = 110GB x compression ratio

Tape capacity (DLT) = 40GB x compression

Tape capacity (AIT) = 50GB x compression ratio

Calculate the number of tapes required:

$$\# \text{ tapes required} = (\text{total data} \times \text{retained percentage}) / \text{tape capacity}$$

## Documenting Your Solution

As with any network-related project, the most important part of designing and implementing an EBS is solid documentation.

All work done on the project should be documented and kept in a safe place. This documentation can be used by others to recreate the work you have done. It can also be used to troubleshoot problems during or after installation of the EBS.

## EBS Site Survey Form

Use this form to help gather information for the needs analysis and design processes. If the information required for this form differs between servers on the Enterprise Backup Solution, fill out a separate form for each server.

1. Server information:

	Server Type	Capacity (GB)		Server Type	Capacity (GB)
1.	Backup Server		11.		
2.			12.		
3.			13.		
4.			14.		
5.			15.		
6.			16.		
7.			17.		
8.			18.		
9.			19.		
10.			20.		

2. Backup server network topology:

- ☐ 10Mb Ethernet
- ☐ 100Mb Ethernet
- ☐ Gigabit Ethernet

3. Operating system:

- ☐ Novell NetWare
- ☐ Microsoft Windows NT
- ☐ Microsoft Windows 2000
- ☐ HP Tru64 UNIX
- ☐ SUN Solaris
- ☐ IBM AIX

4. Tape management system software:

- ☐ Computer Associates ARCserve
- ☐ VERITAS NetBackup
- ☐ Tivoli Storage Manager
- ☐ VERITAS Backup Exec
- ☐ Legato NetWorker
- ☐ CommVault Galaxy

## 5. Backup scheme:

Day of Week	Backup Window (Hrs)	Start Time	End Time	Backup Type (Full/Incr/Diff)
Sunday				
Monday				
Tuesday				
Wednesday				
Thursday				
Friday				
Saturday				

## 6. What is the expected compression ratio for the backups?

.....

## 7. What is the expected annual growth rate for the data?

.....

## 8. How many years of expected growth should be planned for?

.....

## 9. Granularity of cartridge retention:

Current week: ☐ daily

Current month: ☐ daily ☐ weekly

Current year: ☐ daily ☐ weekly ☐ monthly

## 10. How many weeks or days of historical data is required to be kept “nearline” within the tape library?

.....

11. Will you be using RAIT within the tape library? ☐ yes ☐ no

## 12. List any other important facts about this environment.

.....

.....

## 13. Preferred backup window: from \_\_\_\_\_ to \_\_\_\_\_


## 14. Maximum backup window: from \_\_\_\_\_ to \_\_\_\_\_

15. Can the server be taken offline to perform the backup? ☐ yes ☐ no

## HP StorageWorks Backup Sizing Tool

Sizing (or optimizing) a tape backup system is resource-intensive, complex, and highly dependent on the expertise of the person designing the system. Ever-expanding product lines, options, operating systems, and backup methodologies further complicate the sizing, so recommendations might not always be accurate.

Knowledge of the above necessitated a priority for HP to develop an appropriate tool to assist in weighing the complex array of variables and trade-offs.



Initial Screen in the Backup Sizing Tool

To help you quickly, easily, and accurately select, price, and deploy the optimal EBS for the customer's environment, HP has developed a sizer utility—a robust, interactive tape backup sizing and configuration tool.

The StorageWorks Backup Sizing Tool takes basic parameters about the customer's tape backup environment—such as backup windows, amount of data to be backed up, tape rotation scheme, and number of servers—and outputs a bill of materials for the recommended configuration of the EBS.

You are guided through the sizing process by a series of predetermined questions regarding such things as the applications being used by the customer, the system configuration, and the customer's requirements relating to price, performance, and capacity. Many rules, trade-off decisions, and assumptions must be considered to ensure successful results.



---

The HP StorageWorks Backup Sizing Tool is available as a downloadable executable file at: **<http://www.HP.com/products/storageworks/ebs>**

---

## Other Backup Tools

HP offers a complete range of automated storage management software products that enable you to make the most of the storage devices and servers that HP offers. Many of these products support backup and restore. The result is increased system performance and improved productivity of system administrators and users, as well as data protection.

### StorageWorks Enterprise Volume Manager

StorageWorks Enterprise Volume Manager (EVM) is host-based application software that manages controller-based clone and snapshot operations. Clone and snapshot are point-in-time copy functions that can be used to minimize downtime required for system backups and data migration activities.

EVM addresses the need for business continuance and data protection by allowing for parallel processing. It is designed for customers who do not want to or cannot disrupt their computing operations with management activities such as backup.

With the explosion of data, traditional backup operations cannot be completed in a backup window. EVM virtually eliminates the problem of shrinking backup windows by allowing for non-disruptive backup/recovery operations. Backing up from a clone or snapshot in the background allows virtual 24 x 7 operations.

EVM also provides copies of data for operations such as data mining and warehousing, testing, and work distribution. EVM improves production volume application performance, enhances the use of computing and IT resources, and allows users to perform complex cloning and snapshot tasks through a user-friendly GUI and Script wizard.

EVM operates on the HP RAID Arrays RA/MA8000, ESA/EMA12000, using HSG80 controllers in switch or hub configurations. Currently, EVM is supported with Windows NT/Windows 2000, Tru64, and Sun Solaris.



## StorageWorks Virtual Replicator

HP StorageWorks Virtual Replicator (SWVR) comprises a rich set of innovative capabilities that simplify storage management for Windows NT environments. A key deliverable of the Enterprise Network Storage Architecture (ENSA), SWVR delivers two core capabilities across the ENSA family—*Storage Pooling/Virtualization* and *Snapshots*.

Snapshots can be used to facilitate backups. The Snapshot for Backup wizard schedules a backup operation of a snapshot to tape. The snapshot and backup operations are scheduled using the HP Batch Scheduler.

The Restore from Snapshot wizard schedules the re-creation of a virtual disk from a snapshot. If specified, after successful completion of the schedule, the previous virtual disk and snapshot are both deleted. The restore operation is scheduled using the HP Batch Scheduler.

## Microsoft Exchange Backup/Restore Solution

Restore Accelerator, incorporated with SWVR and EVM, extends the HP leadership position in deploying Microsoft Exchange environments by delivering the highest levels of recoverability for this mission-critical application.

Through the integration of SWVR or EVM and customizable scripts, customers can now perform a quick, complete, and fully automated recovery of their Exchange environment. By using a snapshot of the Exchange information store, customers can resume full operation of their Exchange environment in minutes. This solution also increases data protection by providing an additional recovery alternative to restoring from tape.

Incorporating SWVR or EVM with HP Batch Scheduler or NT Scripting and a scripted procedure that requires minimal customization enables the following:

- Quiese Microsoft Exchange by shutting it down briefly.
- Snapshot the Exchange disks that contain the information store.
- Restart Exchange and resume operations.

## **StorageWorks Data Replication Manager**

Businesses, particularly those that operate in national and world markets, must have business continuance at all times. That means their data must be available continuously with complete data integrity regardless of circumstances. StorageWorks Data Replication Manager (DRM) provides customers the widest choice of performance, distance, and availability to meet their enterprise-level network storage requirements.

Using the DRM software, data replication is performed at the storage system level and in the background to any host activity. The total business continuance solution includes the RA/MA8000 and/or the ESA/EMA12000 external storage systems. The RA/MA8000 and ESA/EMA12000 systems running DRM can either replicate data up to 40km (approximately 25 miles) through an extended SAN over direct Fibre Channel links at full Fibre Channel speeds (100MB/s), or go unlimited distances with new FC-to-ATM gateways.

For maximum protection from wide-area disasters, or where direct Fibre Channel is not available, DRM supports Fibre Channel-to-ATM connectivity into telecommunication networks. Data replication can be performed hundreds or thousands of miles in distance.

## Using EBS to Back Up a Cluster

Server clusters provide high availability of applications and data to users. In a server cluster, two or more servers (called nodes) are linked in a network. The nodes run cluster software that allows each node access to the shared SCSI bus to which any number of disks can be connected. If a node becomes unavailable, cluster resources migrate to an available node (called failover). The disks on the shared SCSI bus and the virtual server are kept available. During failover, users experience only a short interruption in service.

## Backup Exec in a Cluster Environment

Backup Exec protects data on the disks attached to the shared SCSI bus, as well as data on the nodes' local disks. Backup Exec can also protect Microsoft SQL Server and Exchange Server databases that are configured as virtual server applications; that is, the virtual server applications contain an IP address resource and a Network Name resource, and are displayed on the network with a unique server name (the virtual server name).

When installing Backup Exec into a cluster environment, install it as a virtual server application. Assign Backup Exec an IP address resource, a Network Name resource (the virtual server name), and a disk resource.

When a failover occurs, backup jobs that were running are rescheduled. The Backup Exec services are restarted on the node that is designated as the failover node, and the backup jobs are restarted by default. The default setting can be changed so that jobs are not restarted, and a different failover node can be specified if Backup Exec is reinstalled.



### Note

Backup Exec 8.5 supports up to three failover nodes for Microsoft Windows 2000 Datacenter Server.

---

## ARCserve2000 in a Cluster Environment

The ARCserve2000 for Windows NT/Windows 2000 enables the user to back up and restore either the primary or the secondary node of a clustered pair of servers in the EBS environment. The user can create the necessary backups of system resources, recognizing some special requirements of the cluster registry hive/files, the quorum resource, and the disk signature. ARCserve2000 will automatically add the recovered server back into the cluster. However, if the primary node should fail during a backup job, the backup job does not failover to the secondary node nor is the backup job rescheduled on the secondary node.

## **Backup Exec and ARCserve in a Novell NetWare Cluster Environment**

For both Backup Exec and ARCserve in NetWare clusters, backup jobs cannot be configured for failover. If the primary server fails during a backup job, the backup job will fail. This backup job cannot be configured to fail over to a surviving node and be restarted. The system administrator will need to recognize that a failover has occurred and will have to manually configure and submit a backup job for the cluster volume from the surviving node. If a failover occurs prior to a scheduled backup job, the systems administrator must recreate and submit the backup job from a surviving node in the cluster.

## **Clustering Environments in the Enterprise Backup Solution with VERITAS NetBackup DataCenter**

### **Microsoft Cluster Server (Windows 2000)**

#### **Cluster Functionality**

Microsoft Cluster Server is supported in a two-node configuration for EBS. For Windows 2000 Datacenter Server, VERITAS NetBackup supports up to four nodes in a cluster configuration. In cluster configurations, the Backup class should be configured to back up specific drive letters. The cluster server has the ability to fail over the required NetBackup services and the required clustered services that provide the protection for the NetBackup media servers.

#### **Cluster Failover**

The failover of the primary node results in the following events:

- Relevant NetBackup services and cluster services are moved to the secondary node.
- Control of the shared storage is moved to the secondary node.
- The active backup fails and is terminated
- A second backup attempt is initiated on the secondary node if the secondary node is a Media server. If the secondary node is a Master server, the backup job is not restarted. Instead, the secondary node will backup the data at the next scheduled backup.

## **VERITAS Cluster Software for Solaris**

### **Cluster Functionality**

In a two-node cluster configuration, both the primary and secondary nodes have the NetBackup Media Server installed on them with identical backup schedules when the ALL\_LOCAL\_DRIVES directive is selected for each node. The primary node backs up its local disks, as well as the shared storage. The secondary node backs up its local disks, but not the shared storage.

When a primary node failure occurs, VERITAS Cluster Server presents the shared storage to the secondary node. The secondary node performs its regularly scheduled backup, which now includes the shared storage. If a failover were to happen during a backup job of the shared storage, the job would fail and NetBackup would attempt to back it up during the next scheduled backup policy.

### **Cluster Failover**

In the event of a VERITAS Cluster Server failure of the primary node:

- The VERITAS Cluster relocates the shared storage to the secondary node through a disk service.
- NetBackup Media Server services on the secondary node continue to run, and the scheduled backup includes the shared storage specified in the ALL\_LOCAL\_DRIVES directive.

## **Legato NetWorker in a Cluster Environment**

### **Microsoft Cluster Server(Windows NT/Windows 2000)**

EBS with Legato NetWorker in an MSCS environment consists of a NetWorker server, a SmartMedia server, and a two-node cluster.

#### **Cluster Functionality**

EBS with Legato NetWorker in an MSCS environment consists of a NetWorker server, a SmartMedia server, and a two-node cluster. Both the primary and secondary nodes have access to shared storage. Under normal operating The primary node backs up shared storage. In the event of a cluster failover, the ownership of the shared storage is transferred to a new node where it will be found during the next scheduled backup.

#### **Cluster Failover**

The following scenarios occur during a cluster failover:

- Prior to the backup or after backup – Shared storage is failed over from one node to another. The jobs for both physical nodes complete successfully. The jobs for each virtual node complete successfully.
- During backup – The jobs for the physical nodes complete successfully. The saveset will report a successful completion for any virtual node that was in the process of being backed up. However, the backup of data stops at the point of failover. No action is required because the next scheduled backup will backup the data correctly.

## Learning Check

1. List four questions to ask when conducting a backup needs analysis.  
.....  
.....  
.....  
.....
2. What is the name of the tool HP developed to assist in the process of designing the appropriate backup solution for a particular customer's requirements?  
.....
3. With higher data compression ratios, backup performance increases significantly.  
☐ True  
☐ False
4. For maximum performance, a DLT drive should be fed data at about \_\_\_\_\_ the rate that it writes it to the tape.
  - a. One-third
  - b. Equal
  - c. Twice
  - d. Three times

5. Which kind of data is the least compressible?
  - a. Database
  - b. Spreadsheet
  - c. Desktop publishing
  - d. Bitmaps
6. Rank the following controllers by the rate at which they feed data to the backup device. Assign a 1 to the fastest and a 3 to the slowest.
  - a. SMART Array 5300 Controller .....
  - b. Internal SCSI controller .....
  - c. Fibre Channel controller .....
7. Which of the following is **not** part of the information required for sizing the Enterprise Backup Solution?
  - a. The tape controller being used
  - b. The compression ratio of the data being backed up
  - c. The tape retention schedule
  - d. Whether partial backups will be differential or incremental
8. If the feed speed exceeds the compressed transfer rate of the tape drive by less than a 3:1 ratio, but more than a 2:1 ratio, reduce the calculated compressed rate by \_\_\_\_\_ when calculating the number of drives needed to meet the backup window.
  - a. 3%
  - b. 14%
  - c. 20%
  - d. 50%



9. The service department at Weimar Widget Corporation plans to perform one full backup and four incremental backups each week. They estimate that about 5% of their total data is new or modified each day and would require backup.
  - a. If they wish to retain data for four weeks, by what percentage must their total data be multiplied to obtain the total retained data?  
.....
  - b. If the company estimates full backups in the future to measure about 200GB, and they plan to use DLT 35/70 tapes and drives, how many tapes will be required to achieve their retention schedule? Assume that the data is compressible at 2.5:1.  
.....
10. To convert GB/hr to MB/s:
  - a. Multiply by 16.66
  - b. Divide by 60
  - c. Divide by 16.66
  - d. Divide by 3.6
11. The certified backup applications are cluster-aware for backing up Microsoft Cluster Servers, but the Enterprise Backup Solution hardware does not support automatic backup failover.
  - ☐ True
  - ☐ False




### Objectives

After completing this module, you should be able to:

- Describe factors that affect performance of backups and restores
- Describe tools for measuring performance of backups and restores


## System View



Performance will always be limited by one or more bottlenecks in the system, of which the tape drive is just one part. The goal is to make the tape drive the bottleneck so that the system will achieve optimal performance from the tape drive.

This puts the emphasis on the rest of the system. The flow of data throughout your system must be fast enough to provide the tape drive with data at its desired rates. High-speed tape drives such as the Ultrium 460 are so fast this can be a difficult goal to achieve. Industry-quoted speeds are not typical speeds. However, if your system is achieving significantly less performance than the published speeds, there may be one or more bottlenecks in your system.

## Primary Storage Performance



Measuring the performance of the disk subsystem is a good place to start for finding bottlenecks in the system. The performance of the disk subsystem is by far the most critical factor in being able to stream tape drives. In order to maintain streaming in the tape drive and to avoid buffer underruns, a 3:1 source speed to feed speed is needed. Steps for tuning performance on the primary storage include:

- **Controller type** – Select RAID controllers capable of high throughput and utilize multiple channels if possible
- **RAID configuration** – Striping with a high number of spindles generally improves performance
- **Drive speeds** – 15K rpm drives will generally provide faster data access than 10K rpm or 7200 rpm drives
- **RAID cache settings** – Optimal RAID cache settings for applications on the server may not be best for backup performance. In general, it is the cache read parameters that will impact backup performance the most.
- **Stripe size** – Ensure that the stripe size is set correctly at install time to best match the application the server will be running.

## File System Performance

The distribution of file sizes on a system is critical to performance. Large numbers of small files cause a high overhead to the file system. Therefore, the amount of time spent in the file system is high compared to the actual amount of data transferred. Large files on the other hand have a lower overhead for the amount of data transferred. The file size distribution is determined largely by the application. For example, file servers tend to have a higher proportion of small file sizes, which will cause performance issues in file-by-file backup mode.

Deep file structures also can increase the number of directory accesses. Following the file system access paths translates into additional seeks on disk.

File fragmentation whereby a file, instead of being contiguous, is spread out over various random parts of the disk can also severely impact performance. Fragmentation will cause additional seeks during backups. Defragmentation of filesystems prior to backups can improve performance.

If the filesystem consists of a large number of small files, using image backups allows the disk to be read sequentially sector by sector, resulting in no random motion on disk.

Before making the switch to image mode, the following should be considered:

- Image mode is a complete backup of a single partition or LUN. A subset, such as a file or set of files cannot be backed up in this way.
- When restoring, the complete partition (or LUN) must be restored as a single operation unless the application is able to extract the required data itself. This may require the use of a spare partition to restore to so as not to overwrite data in the active partition.

The trade-off is effectively between backup performance and ease of restoring subsets of the partition or LUN.

## Measuring Primary Storage Performance

To measure the data rate from your disk subsystem, use the **Performance Assessment Tool (PAT)** and select the method used by your backup - file-by-file or sequential. PAT allows you to simulate reading data from your disk subsystem. You can select the very files you want to backup and receive a MB/sec rating of how your disk subsystem is performing on your actual data.

In addition, PAT HP-UX allows multiple streams to be specified so you can see the increased performance, which will result if multiple streams are enabled via the backup application (should your backup application support such a configuration).

PAT can be downloaded from <http://www.hp.com> .

## Backup Server Performance

The server is central to the backup process as it runs the backup software and the data is passed into and out of the server's main memory as it is read from disk subsystem and written to tape.

To see if the server is a possible bottleneck use one of the readily available system monitoring tools while performing a backup. Ensure there are still resources free. Watch for:

- CPU bandwidth
- Memory capacity
- I/O bandwidth
- Virus scans
- CPU intensive screen savers

Also check for:

- Compatibility between operating system, backup application, and hardware
- Up-to-date drivers and firmware
- Latest backup application patches

HP recommends:

- Multi-processor server with at least 512MB of system memory
- Not running any other applications during the backup
- Running backup application from fast drives on the server

## Backup Application Setup

These are HP's recommendations for achieving high performance with backup applications:

- For general backup, use an HP recommended backup application. Native applications (e.g. tar, cpio on unix) do not provide high performance.
- For database applications (e.g. SQL 7.0, Exchange 2000), use the backup functionality provided by those applications as they are tuned to make best use of their data structures.
- Use large (SCSI) transfer sizes (>64K).
- Increase the system memory allocated to the backup process if possible
- Use hardware data compression – HP's tape drives do this in hardware at very high speed whereas it is much slower in software on the server. Also, compressing data twice tends to give non-optimal compression ratios.
- Use multi-threading (concurrency) if possible. This allows multiple backups to be interleaved to the tape thus reducing the effect of disk seeks for each one.
- If the system consists of small files (<64KB), consider image/sequential backups.



## Switch Performance



These are HP's recommendations for achieving high performance with backup applications:

- Use switched fabric for highest performance.
- Use 2Gb components from end to end (server to router).
- Use multiple Inter-Switch Links (ISLs) to provide additional bandwidth. An ISL is recommended for every four Ultrium drive attached to the fabric.

## Fibre Channel-to-SCSI Bridge Performance



In order to optimize performance, a maximum of one tape drive per SCSI port is always the best configuration. In addition, the maximum throughput of the router must also be considered.

Below is a chart for supported and optimal configurations.

Router	SDLT 220 (Ultra2)		SDLT 320 (Ultra3)		Ultrium 230 (Ultra2)		Ultrium 460 (Ultra3)	
	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming	Supported drives per SCSI port	Maximum drives for streaming
N1200	2	2 per router	Not supported	Not supported	2	2 per router	Not supported	Not supported
e1200, e2400	2	2 per card	Not supported	Not supported	2	2 per card	Not supported	Not supported
M2402	2	8 per router, one per SCSI port	2	8 per router, one per SCSI port	2	8 per router, one per SCSI port	1	4 per router, one per SCSI port

## Tape Drive Performance



Tape drives are usually not the bottleneck in the system. If other factors have been checked and the tape drive performance is still slow, check the following:

- Cabling from server to tape drive
  - Should be good quality Ultra 160/320 SCSI cabling
  - Should be within SCSI specs – watch for length
  - Should be terminated properly
  - Should be free from any bent pins or badly seated connectors
- Other peripherals on the bus – high-speed tape drives must have exclusive access to the HBA/bus to achieve their potential bandwidth.
- Damaged or worn media – can cause excessive use of tape or physical retries and hence a drop in performance. You can use LTT to check for poor media performance or try a new tape.
- Dirty heads – if this is the case then it is likely the ‘clean’ light will be on. HP tape drives have self-cleaning mechanisms to virtually eliminate dirty heads but you can still use a cleaning cartridge to make sure.

### Measuring Tape Drive Performance

To measure the data rate to your tape drive use the tool **HPTapePerf** (if under Windows environments) or **IDD** if on a UNIX platform. These utilities transfer data directly from the memory of the server to the tape drive and can verify performance independent of the disk subsystem and application. Both utilities can use different compression ratios, block sizes, transfer lengths, and overall size of transfer to be configured. If you have verified that data can be read off the disk subsystem at the expected rates using PAT, then check how fast the system can access the tape drive using HPTapePerf or IDD.

Both HPTapePerf and IDD can be downloaded from <http://www.hp.com>.

## Restore Performance

Most of the use of a tape drive is for the backup of data. Restore tends to be reserved for when something has gone wrong on the disk (corruption, loss) and you want some or all of your data back.

The tape drive writes and reads at the same speed but, importantly, your disk system doesn't. When writing data back to disk, the file structures need to be created and this takes extra head movements. Also, you can't take advantage of the read cache.

You can measure the performance of your disk subsystem using the **MkData** tool. This creates a virtual directory structure based on file size, directory depth, directory width and compression ratio and measures the data rates it achieves.

If you have an issue with restore performance, then check your backup performance first. If that is OK then it is almost certainly the disk subsystem that is the bottleneck. There are three main options:

- Use even more serious hardware
- Use image backup/restore – the file structures are copied at the same time
- Accept that restore performance is not as fast as backup performance. If you only restore occasionally this may be acceptable.

## Learning Check

1. What is usually the most critical factor in streaming tape drives?

.....

.....

2. What does the Performance Assessment Tool assess?

.....

.....

3. Which tool can be used to assess restore performance?

.....

4. Which tool can be used to assess tape drive performance?

.....

### Objectives

After completing this module, you should be able to:

- Describe the Library and Tape Tools and its features.
- Describe the HP management support for the Enterprise Backup Solutions.
- Identify the limitations of Fibre Channel on secondary storage compared with primary storage.
- Diagnose problems and perform repairs to the system using the troubleshooting flowchart method and tools provided by HP.

## HP StorageWorks Library and Tape Tools

HP StorageWorks Library and Tape Tools (L&TT) is a collection of storage hardware management and diagnostic tools for tape, tape automation, and archival products. L&TT assembles these tools into a single, convenient program.

### Software Features

L&TT offers the following features:

- **Installation Check**—L&TT guides you through a basic installation check of your product. The software assists the user in choosing an appropriate HBA and SCSI ID(s), ensuring that the device is detected by the system, and verifying key device functionality. The installation check feature is essentially HTML documentation that helps with the most common generic installation issues while also describing how to use L&TT to verify the device installation.
- **Device Identification**—L&TT clearly identifies the storage products connected to the system, along with key information on product configuration and status.
- **Troubleshooting Tests**—L&TT provides various tests to verify product functionality or to isolate product issues. Tests include device self-tests, read/write tests on drives, exerciser tests for autoloaders and libraries, and specific device utilities.
- **Firmware Upgrades**—L&TT provides a convenient way of updating product firmware, enabling users with an Internet connection to take advantage of ongoing enhancements. The software can be configured to check the Web automatically for firmware updates for connected devices, or users can manually check the Web for updates if the automatic update feature is not desired. If updated firmware is available, the program notifies the user, and the updates can easily be copied to the system. With libraries, users can upgrade the library and the embedded drive firmware in the same operation. Wherever possible, the embedded drives are updated in parallel to reduce time.
- **Support Ticket Generation**—If you experience a problem with a storage product, L&TT can generate a support ticket that includes essential information for troubleshooting the problem. As an alternative to phone support, you can e-mail the support ticket to a support center for assistance. This information streamlines the support process and enables the support staff to better serve you if a support call is made later. When a support ticket for a device is generated, L&TT performs a Device Analysis test on the device. The support ticket contains generic information about a device, as well as the results of the Device Analysis test. The Device Analysis test can be performed by itself, but HP recommends generating a support ticket because the resulting data is presented in a more useful format.
- **Automatic Notification of Web Updates**—If a connection to the Internet is present and web updates are enabled in the tool preferences, L&TT automatically informs you of the following updates, if available, each time the program is started:
  - New versions of L&TT

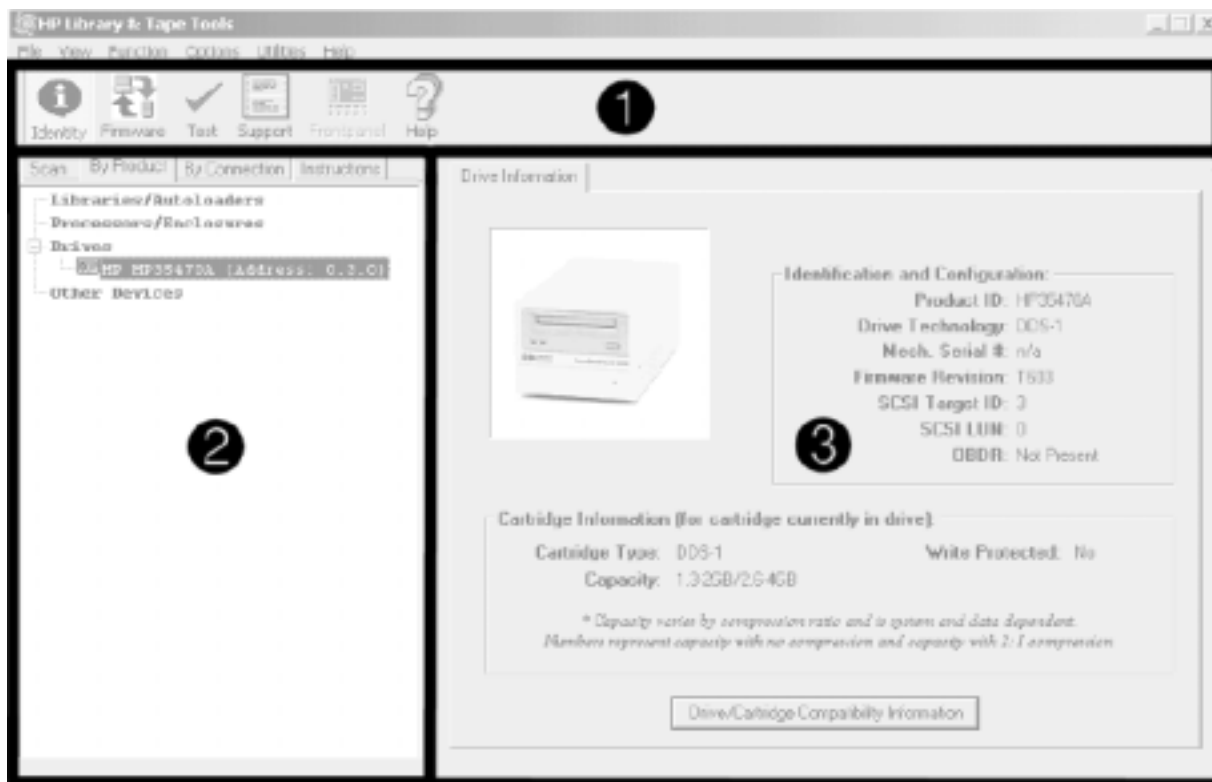
- New firmware files for connected devices
- New device-specific functionality (such as new or updated tests) for connected devices

## Supported Products and Operating Systems

For a complete listing of compatible products, refer to the L&TT Compatibility Matrix found at <http://www.hp.com/support/lttcompatibility>. The level of functionality that L&TT offers for each device varies depending on features of the device, and the degree of device integration into L&TT.

The Windows version of L&TT uses a graphical user interface (GUI), whereas the NetWare, HP-UX, and Tru64 versions of the program use a command screen interface (CSI).

## Application Window Layout



The three sections of the L&TT main screen are:

1. **Taskbar**—This section contains buttons that provide quick access to the main functions of L&TT and to the online help system.
2. **Device List**—This is a multi-function window that offers several options on the following tabs:
  - **Scan**—This option provides either a summary status or detailed information (depending on whether **Show Details** or **Hide Details** is selected) about the bus scanning process. If a problem is encountered during the scan, this information may help in determining the cause. When the scan completes successfully, the device list automatically

switches to the **By Product** tab. The **Scan** tab also lets you rescan the bus. If any devices have been hot-swapped or powered on after the OS has booted, in most cases, the scan feature can discover those devices without requiring a reboot of the system.

- **By Product**—This option shows a list of all the products connected to the system. The list is grouped into the following four categories:
  - ◆ Libraries and autoloaders
  - ◆ Drives
  - ◆ Enclosures and processors
  - ◆ Other devices

The three number fields listed after the device represent the device address. Each field in the address is separated by a period: the first field represents the HBA channel, the second field represents the SCSI ID, and the third field represents the LUN.

- **By Connection**—This option shows all products connected to the system, grouped by the HBA they are connected to. This view makes it easy to see which devices are connected to the same bus as the device in question, and may help in understanding system I/O performance issues.
  - **Instructions**—This window contains brief instructions on how to use the selected screen. This view can be disabled in the software preferences.
3. **Device Information screen**—All the main functions of the program are displayed in this window. The content of this window depends on the device and tool function selected.

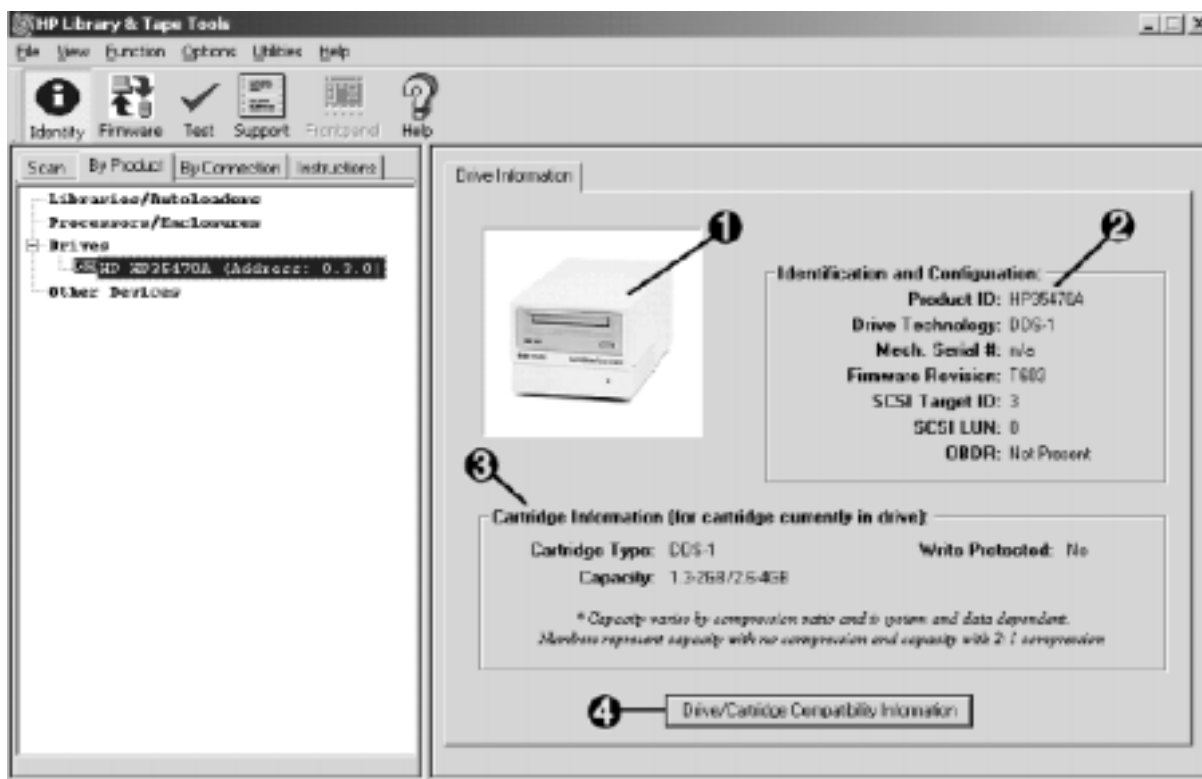
## Using the Device Information Screen

When you select a product from the **Device List**, the **Device Information** screen displays information relevant to the device. The **Device List** also changes to display instructions on this page (if the instructions preference is enabled). If another tool function is currently active, then clicking the **Identity** button on the toolbar opens the **Device Information** screen. The **Device Information** screen provides an overview of the selected hardware device and its current configuration and status.



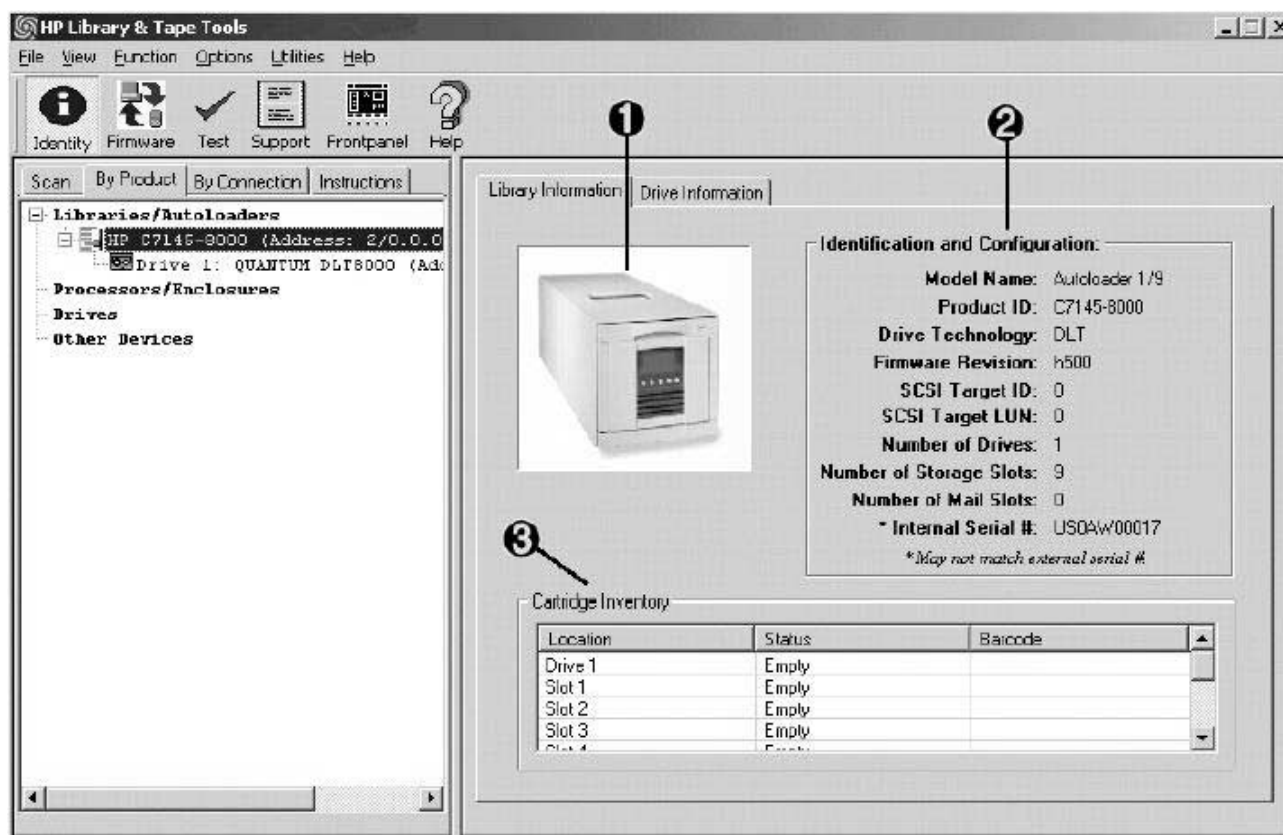
## Standalone Devices

The **Device Information** screen, showing a standalone device, contains a single tab labeled **Drive Information**. The **Device Information** screen has four main sections, as shown.



1	This provides a visual representation of the selected product (the images are grayscale)
2	<p>This provides the following general information on the product:</p> <ul style="list-style-type: none"> <li>• Product ID (Inquiry String)</li> <li>• Drive Technology (DDS/DL/LTO)</li> <li>• Mech Serial # (Drive serial number)</li> <li>• Firmware Revision</li> <li>• SCSI Target ID</li> <li>• SCSI LUN</li> </ul>
3	If a data cartridge was inserted in the drive when the identity tool was selected (or when the current screen was "refreshed"), information regarding cartridge type, capacity, and write protect status is shown here.
4	Click this button to display a list of media that is compatible with the product.

## Library and Autoloader Products




When a library or autoloader is selected in the **Device List**, the **Device Information** screen displays two tabs: one for library information and one for the drives contained within the library.

1	This provides a visual representation of the selected product.
2	This area provides the relevant information on the selected library product.
3	This provides a cartridge (media) inventory for all the drives and slots. The current status (full or empty) of all the storage elements and the barcode number (if barcode reading is available) is displayed here.

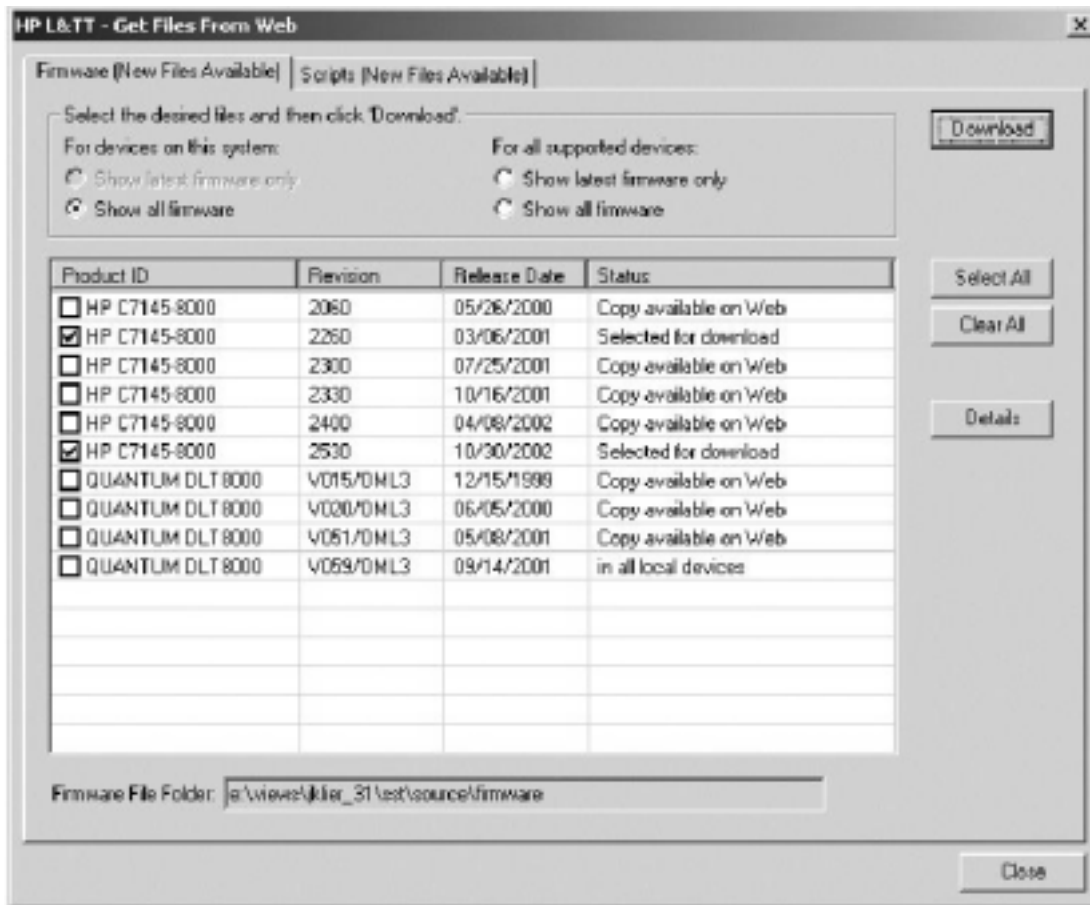
To see information for the drives within the library, select the **Drive Information** tab. This tab displays a screen very similar to the one for standalone products, with the addition of the **Selected Drive** dropdown menu. Because a library can contain multiple drives, you can select and view information for each drive using the **Selected Drive** dropdown menu. All of the drives in the library are available in the dropdown menu.

## Using the Firmware Management Screen

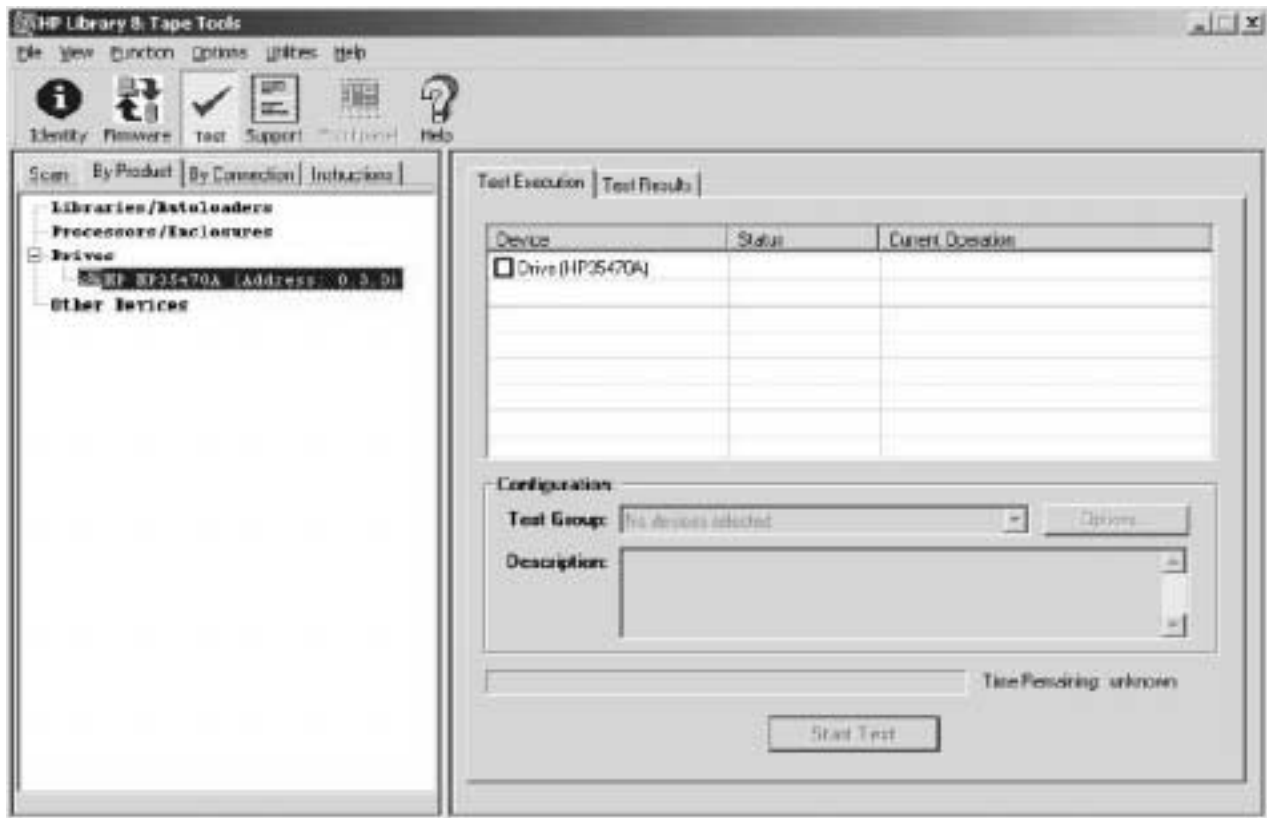


The Firmware Management functionality of L&TT provides a convenient way to easily upgrade the firmware of connected devices. Library firmware and firmware of drives within the libraries can be upgraded as well. In many cases, multiple drives within the library can be updated simultaneously to save time. L&TT uses special firmware files that associate the firmware with a particular product. This feature prevents the user from downloading the wrong image and ensures that the downloaded firmware is the same or a newer revision than the current device firmware.

Under the **Local Firmware Files** tab, a **Get Files from Web** button allows the administrator to download firmware files to the local management host.



## Using the Tests Screen



L&TT provides numerous tests for analyzing storage hardware. L&TT tests are SCSI script-based. A SCSI script is a sequence of SCSI commands sent to a device, usually for the purpose of testing the device's functionality. SCSI scripts can also be used to perform non-diagnostic utility functions, such as rewinding a tape, recalibrating a tape head, or “exercising” a tape-moving robot arm. L&TT uses the following three types of SCSI scripts:

- **Built-in tests** are hardcoded into the program and are installed during the L&TT installation process.
- **Script-based tests** are additional tests based on scripts that are stored in the *scripts* directory in the L&TT installation directory. The script-based test files have the \*.*tst* extension. Users can download new and updated scripts. With the Windows version of L&TT, the management of script files is accomplished through the **Get files from Web** item on the **File** menu. HP is constantly adding scripts for new devices and updating scripts to enhance the testing functionality of older devices. As new scripts are developed, they are published to the external website in the scripts section.
- **SCSI utility scripts** provide access to generic SCSI operations. These scripts are accessed by selecting **Run Script...** from the **Utilities** menu and have the \*.*scc* extension.

## Built-In Tests

- **Connectivity Test** examines the ability of the host to communicate with the device. This is a test of the SCSI communication path only; no media is required. This option is useful when trying to establish that the SCSI cables and adapters are working correctly. This test is available for DLT, LTO, and DDS drives only.
- **Device Self-Test** executes the internal self-test of the device, and any failure is reported. This test is useful in confirming that the device itself is operating properly. Not all devices support an internal self-test.
- **Device Analysis Test** runs a device-specific script that carries out a detailed examination of the device. The test examines all available logs from the device to look for any known issues, customer use issues, and common error patterns. The test also provides summary information and presents suggestions as to the cause of certain errors and any remedial actions the user should take. This option is particularly useful if you suspect that there may be a problem with the device.
- **Read/Write Test** writes and reads data from tape. In addition to verifying that the data read-back matches the data written, the read/write error rates are also checked, and the test fails if the error rates exceed allowable limits. This test can be used to verify the ability of the drive to read and write data independent of the storage application. Use a blank tape when performing this test because this test is destructive to data on the tape. This test is available for all tape drives.
- **Media Validation Test** determines if the drive can read the data on any media currently in the drive. The test is used to determine the quality of the media. The test reads data from the media and calculates the error rate. If the error rate exceeds the maximum allowable threshold, the test will require a cleaning cartridge and a second tape to determine if the problem is with the drive or the tape. The amount of data read depends on the drive type and coverage level of the test, but it generally varies between a few megabytes and one gigabyte. No data is written to tape and the test is not destructive to data. This test is useful if the user suspects that a backup failed because of a bad tape. This test is available for all tape drives.

---

### Note

This test requires that a minimum amount of data is already written on the tape. Otherwise, the test will fail.

---

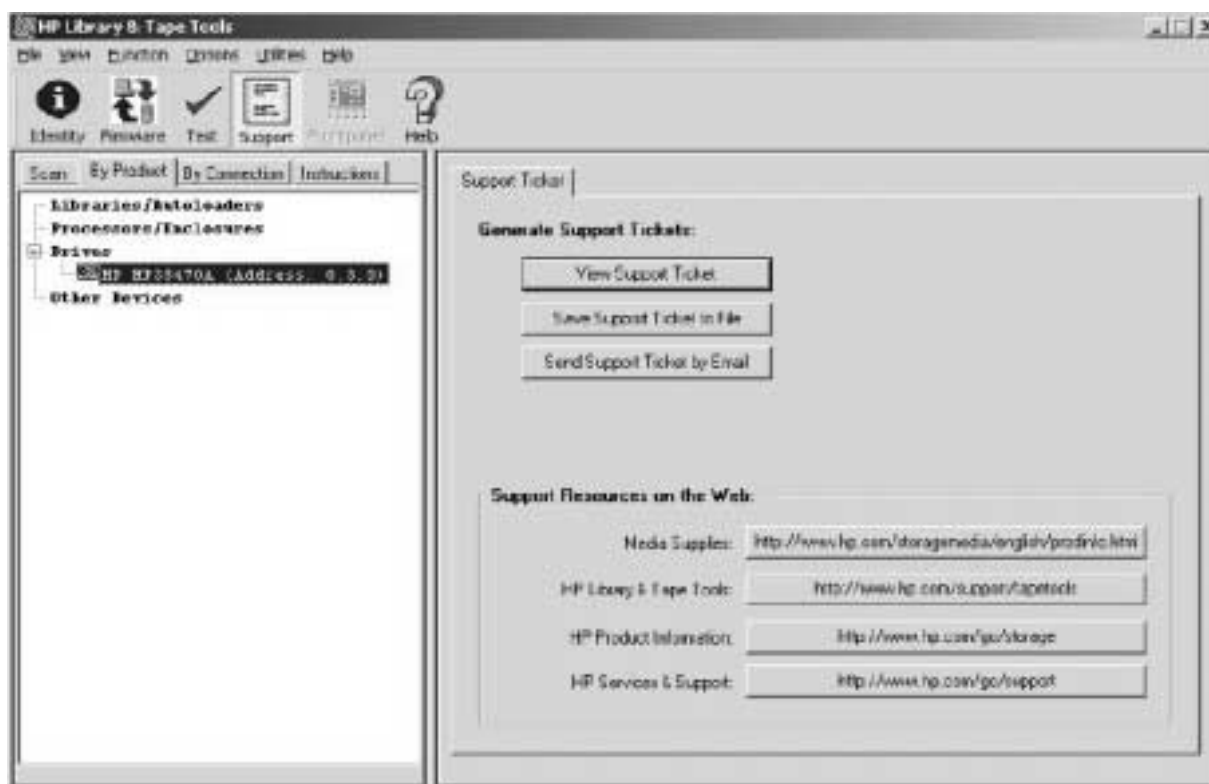
- **Media Analysis Test** reads all data on a cartridge from beginning to end. The test determines if the media is readable, and it also provides statistical data such as block size, compression ratio, and speed. This test is available for DLT, LTO, and DDS drives only.
- **Library Exerciser Test** takes a test cartridge from a mail slot and performs random slot-to-slot moves. The test then moves the test cartridge into every

available drive and performs a short drive test. This test is intended to exercise the robotics and robot drive hardware.

### Script-Based Tests

- **Compression Test** writes a data pattern of known compressibility to tape. The test then rewinds the tape and reads the file just written. The read and write hardware compression ratios are calculated to ascertain that the hardware compression is enabled and working. If the script finds that hardware compression has been disabled by the host or by switch configuration, it reports this to the user. The compression ratios are displayed to the user and should be approximately 2:1. This test is useful if tape capacity or hardware compression is in doubt. This test is available for DLT, LTO, and DDS drives only.
- **DLT FUP Creation Test** takes a firmware image and writes it to a tape cartridge for use as a firmware update (FUP) tape. This test is only available for DLT drives.
- **LTO FUP Creation Test** takes a firmware image and writes it to a tape cartridge for use as a FUP tape. This test is only available for LTO drives. DLT SRAM Check Tests (short and long versions) check drive logs for existing evidence of SRAM chip corruption, and they perform a read/write test to check the drive for susceptibility to SRAM corruption problems. The test is destructive to data on the cartridge in the drive. These tests are only available for DLT drives.
- **Initialize Element Status Test** performs an *initialize element status* command on the library. This test is available for all libraries.
- **Lateral Motion Test** performs lateral motions of the robotics to enhance performance of brush tracks. This test is available for C7200 library models only.
- **MO Configuration Test + Reset** checks the configuration parameters of the MO drive and identifies configurations that differ from the factory default settings. If settings differ, the test can reset the drive to the factory defaults. This test is only available for MO drives.
- **LTO Thermal Test** monitors the temperature sensors of an LTO drive to ensure that the drive is cooling properly when installed in a server.

## Using the Support Ticket Generation Screen



L&TT can automatically generate a support ticket based on information that the program determines about the system. The information in this support ticket can then be used to diagnose storage hardware issues. L&TT can automatically e-mail the support ticket to the HP help desk to assist the help desk technician in resolving issues. To access the **Support Ticket Generation** screen, click the device you want to generate the support ticket for in the device list, and then click the **Support** button on the main toolbar.

When a support ticket is generated, L&TT collects configuration information and executes a Device Analysis test on the selected device. This information can then be viewed, saved, printed, or e-mailed. A support ticket is saved as a log file having the \*.ltx extension. By default, log files are saved to the *log* directory. The support ticket viewer is built into the L&TT application. The viewer enables viewing of detailed device information and general information on the system it is connected to. This information is displayed in a standard tree-view format. The user (or support technician) can change the detail level of the support ticket, browse through the support ticket information, and expand or collapse specific sections of the support ticket as required.

---

### Note

Regardless of the selected detail level, the log file always contains all of the information collected from the product. The support ticket viewer filters this information, depending on the detail level.

---



L&TT support tickets are saved in a proprietary format and therefore require the support ticket viewer. However, with the Windows version of L&TT a support ticket can be converted to HTML format using the currently selected detail level. The HTML version can easily be printed or viewed by other users without the support ticket viewer.

---

**Note**

After an HTML support ticket has been generated, the detail level within the HTML file cannot be changed. Make sure that the detail level is set to the desired level before generating the HTML report.

---

## Management Support

HP offers a complete range of automated storage management software products that enable you to make the most of HP storage devices and servers. These products can support the full range of storage management operations—from backup and restore to subsystem management and more. The result is increased system performance and improved productivity of system administrators and users, as well as data protection.

These tools include the capabilities for automated policy-driven operations, media and device management, and cataloging, as well as easy-to-use administration that allows the concept of “lights-out computing” to be implemented for storage management.

The management support tools discussed in this section are:

- HP Insight Manager 7
- Media Robot Utility

## HP Insight Manager 7

HP Insight Manager 7 provides web-based management for HP servers and any device compliant with HTTP, SNMP MIB-2, or DMI v2. The strength of HP Insight Manager 7 lies in its ability to provide system administrators real control through an easy-to-use industry standards-based web interface. It is feature-rich, intuitive, extensible, and designed to unlock the built-in manageability of HP hardware.

Its user interface is intuitive and packed with functionality. Accessed from anywhere on the intranet, systems administrators can manage devices, manage events, and administer HP Insight Manager 7. Preconfigured common system views and customizable queries and events provide out-of-box productivity.

### Features

- Accessibility to management information anywhere on the intranet, more securely than SNMP
- Automatic discovery and identification of industry-standard SNMP MIB-2, DMI-2, and HTTP-based web servers and compliant devices (including third-party devices)
- Comprehensive event management that enables proactive management of distributed devices by alerting in advance of system failures
- Predictive fault management and pre-failure warranty for storage, memory, and processor to protect your hardware investment
- SQL-compatible repository for discovered devices, alarms, asset, and status reporting (supports Microsoft SQL Server 6.5 and 7.0 and Microsoft Data Engine)
- Integrated Cluster Monitor that provides a one-step management control point for cluster management
- Highly scalable architecture designed for extensibility

### Backup Solutions Support

HP Insight Manager 7 enables you to monitor the Fibre Channel-to-SCSI bridge and the libraries. You can see the worldwide name, firmware revision, and status of each component, including the tape drives located in the library.

## Media Robot Utility

Media Robot Utility (MRU) software enables you to monitor and control libraries and mini-libraries conveniently from your desk. MRU provides an application programming interface (API), a GUI on OpenVMS and Tru64 UNIX, a GUI on Windows NT, and a command line interface (CLI).

### Features

You can manipulate and monitor libraries using simple commands:

- **Load** — Moves a cartridge from a slot to a drive
- **Unload** — Moves a cartridge from a drive to a slot
- **Inject** — Moves a cartridge from a port to a slot
- **Eject** — Moves a cartridge from a slot to a port
- **Show** — Displays information
- **Home** — Returns the cartridge to the location it came from
- **Find** — Finds a specific cartridge within the library
- **Help** — Displays information on the commands

### Benefits

- **Eases storage management** — Verifies library/mini-library installation and configuration.
- **Facilitates diagnostics** — Verifies that hardware and control communications paths are functioning properly. If operations such as backup fail, MRU isolates the library and verifies its operations.
- **Customizes storage management** — Enables you to write storage management solutions for your libraries.
- **Provides easy-to-use interfaces** — Choose between CLI, GUI, or API to move cartridges.

## Limitations of Fibre Channel

Certain aspects of Fibre Channel place constraints on how the solution operates.

Fibre Channel class-3, the protocol that HP uses for all Fibre Channel storage products, is considered a “lightweight” protocol. It is designed specifically for use with storage subsystems due to low latency performance. Storage subsystems have extremely low latency requirements, much lower than typical IP protocol stacks can provide. To keep the Fibre Channel class-3 protocol lightweight, error recovery mechanisms are left for upper layers to resolve.

In a primary storage environment, the blocks are relatively small and randomly accessible. Error events caused by Fibre Channel are detected and communicated to a higher level protocol, and an entire block can be retransmitted without impact to performance, provided the error events are infrequent.

In a secondary storage environment, applications do not handle error events as gracefully. Any Fibre Channel error event could manifest itself as a backup or restore job failure. If the job is a scheduled job, it will be rescheduled.

Fortunately, error events in Fibre Channel do not occur frequently in a static environment. The most common error event occurs as a result of loop disruptions. In general, these loop disruptions are a controllable occurrence.

## Common Types of Loop Disruption

The most common type of loop disruption is caused by a loop initialization procedure (LIP). An LIP occurs when a loop component is added or removed from the loop or when a loop component, such as a server, is restarted. If an LIP occurs when a data frame is on the loop, that frame will be lost.

In a primary storage environment, a higher-level protocol is responsible for re-transmitting the data and the incident is unnoticed by the users of the data. In a secondary storage environment, the application vendors have not yet dealt with this characteristic. The error is noticed by the user and shows up as a job failure.

Therefore, certain activities, such as hot-plugging loop components and resetting loop components, should be performed only when there are no backup or restore jobs executing on the loop.

You must restart each Windows NT server whenever you add backup devices to the loop, just as you must restart for the new devices to be seen by Windows NT.

## Diagnosis and Repair Using Troubleshooting Flowcharts

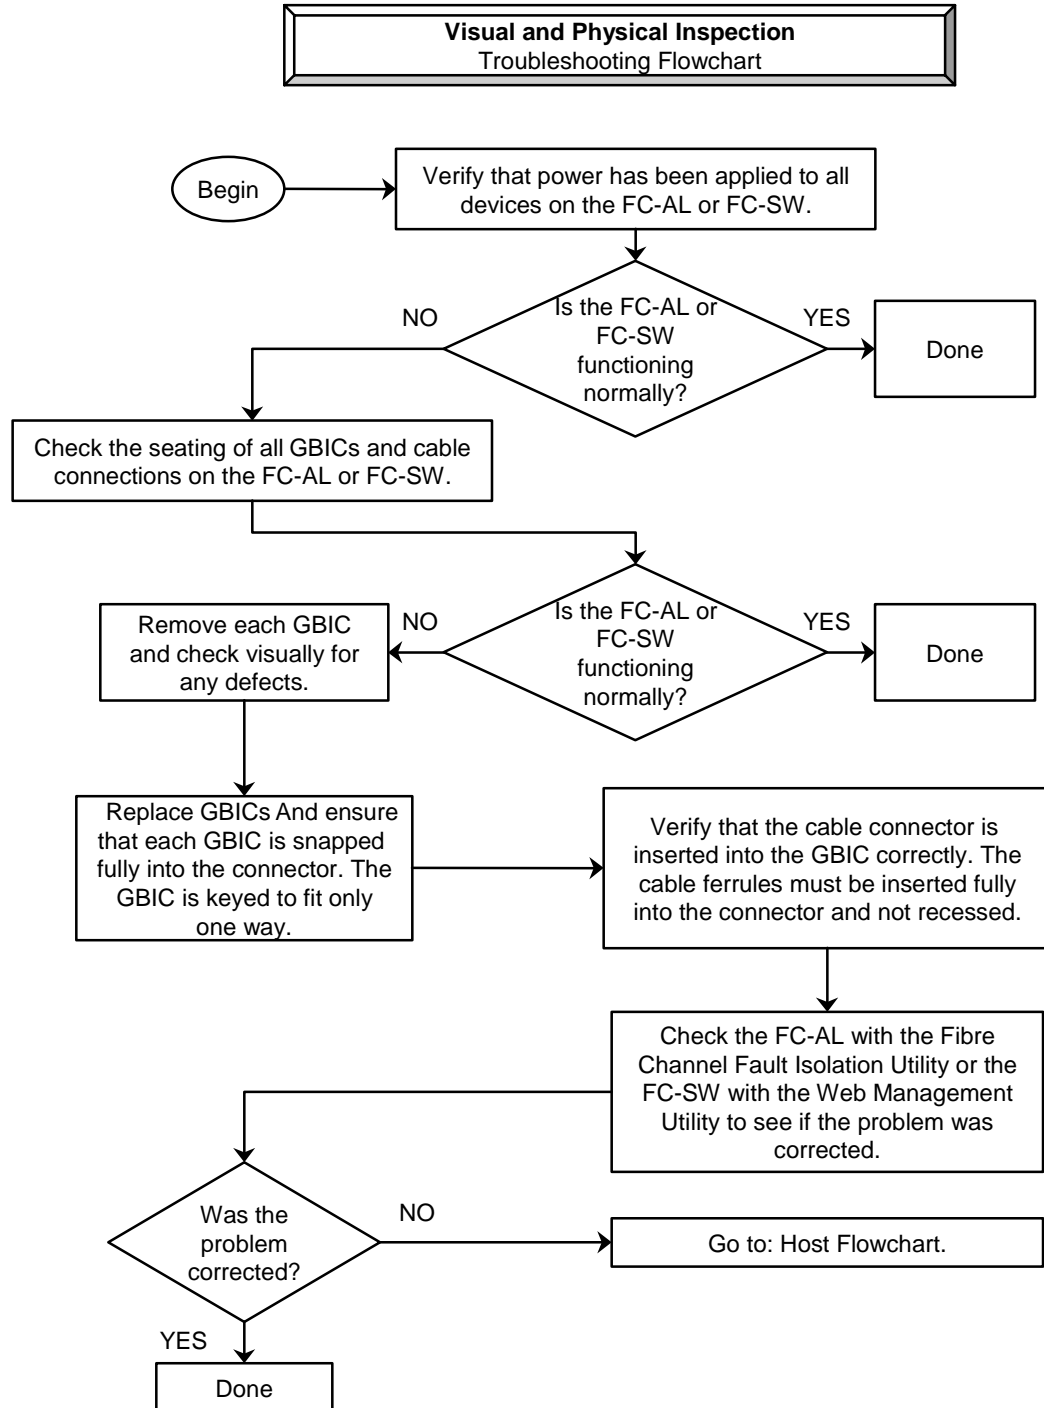
This section contains troubleshooting flowcharts that can be used to help diagnose problems in an Enterprise Backup Solution (EBS). The flowcharts and software utilities provided by HP can be used together to isolate problems and verify repairs.

When diagnosing problems with EBS, follow this order:

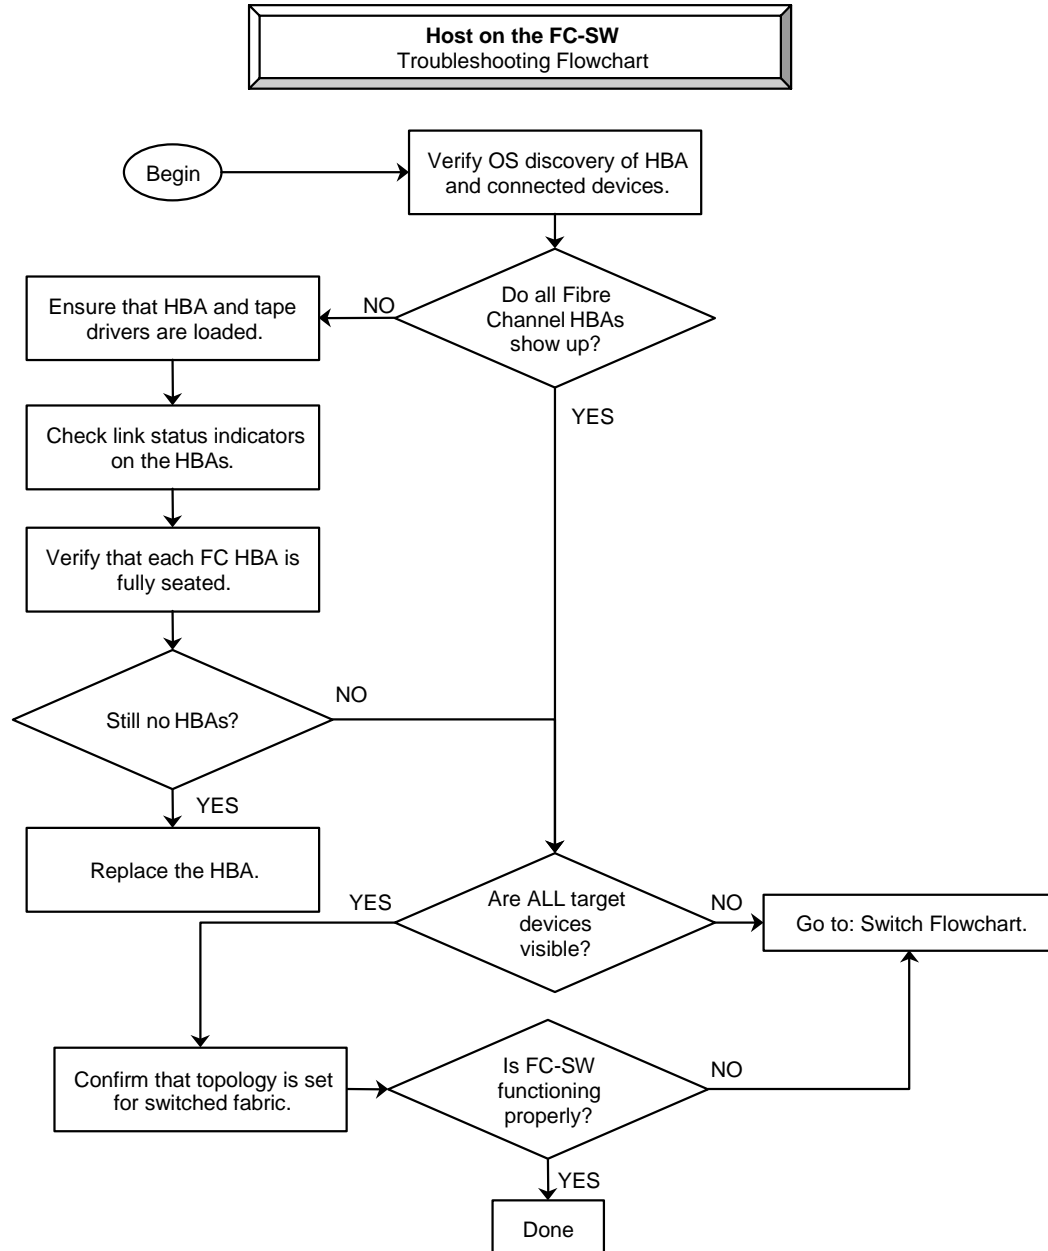
1. Perform a visual and physical inspection of the entire system.
2. Determine if there are problems with the host.
3. Determine if there are problems between the host and the hub or switch. Verify that the host is a member of the switch zone.
4. Determine if there are problems with the hub (if applicable).
5. Determine if there are problems with the switch (if applicable).
6. Determine if there are problems between the hub or switch and the Fibre Channel-to-SCSI bridge. Verify that the Fibre Channel-to-SCSI bridge is a member of the switch zone.
7. Determine if there are problems with the Fibre Channel-to-SCSI bridge.
8. Determine if there are problems with the library.

The troubleshooting flowcharts provided here help you move from one component of the EBS to another to diagnose problems.

## Visual and Physical Inspection



## Checking the Host on the FC-SW






## Device Confirmation

To verify operating system discovery of the HBA and connected devices, use the following methods.

For Windows NT, use the SCSI Adapters control panel.



For Novell NetWare, use the `list storage adapters` command.

For Tru64 UNIX, use the SCSI CAM utility.

```
# scu sc edt
Scanning all available buses, please be patient...
# scu sh edt

CAM Equipment Device Table (EDT) Information:

Device: RRD47    Bus: 0, Target: 5, Lun: 0, Type: Read-Only Direct Access
Device: RZ1DB-CA Bus: 1, Target: 0, Lun: 0, Type: Direct Access
Device: TL895    Bus: 4, Target: 2, Lun: 1, Type: Medium Changer
Device: TZ89     Bus: 4, Target: 2, Lun: 2, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 3, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 4, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 5, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 1, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 2, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 4, Type: Sequential Access
Device: HSG80    Bus: 4, Target: 5, Lun: 1, Type: Direct Access

#
```

In addition, check for detected tape devices by typing *hwmgr -view dev| grep tape*  
.

For Sun Solaris, check the `var/adm/messages` directory.

```
sd6 at SUNW,fas0: target 6 lun 0
sd6 is /sbus@1f,0/SUNW,fas@e,8800000/sd@6,0
fcaw0: JNI Fibre Channel Adapter model FCW
fcaw0: 64-bit (32bit mode) SBus 1: IRQ 3: FCODE Version 11.0.11 [18237e]: SCSI
ID 125: AL PA 01
fcaw0: Fibre Channel WWN: 100000e069402eb7 WWP: 200000e069402eb7
fcaw0: FCA SCSI/IP Driver Version 2.3, Oct 19, 1999 for Solaris 2.5,2.6
fcaw0: All Rights Reserved.
fcaw0: < Total IOPB space used: 1145024 bytes >
fcaw0: < Total DMA space used: 8458269 bytes >
NOTICE: fcaw0 NPORT Initialization Complete, SID=402eb7
fcaw0: New Fabric Parameters Received. Resetting...
NOTICE: fcaw0 NPORT Initialization Complete, SID=402eb7
fcaw0: Host: Port 220d13 (100000e069402eb7:200000e069402eb7)
fcaw0: Port 220013 (50001fe100014720:50001fe100014723) available.
fcaw0: Port 220713 (100000e00202148a:100000e00202148a) available.
fcaw0: Port 210013 (50001fe100014790:50001fe100014792) available.
fcaw0: Port 210213 (100000e002021474:100000e002021474) available.
fcaw0: Port 210a13 (10000000c920e159:10000000c920e159) available.
fcaw0: Port 210c13 (10000000c9213b0b:10000000c9213b0b) available.
fcaw0: Target 64: Port 220013 (50001fe100014720:50001fe100014723) online.
fcaw0: Target 64 Lun 1: Port 220013 (50001fe100014720:50001fe100014723) present.
sd16 at fcaw0: target 40 lun 1
sd16 is /sbus@1f,0/fcaw@1,0/sd@40,1
<DEC-HSG80-V84F cyl 26277 alt 2 hd 20 sec 169>
root on /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0:a fstype ufs
zs0 at sbus0: SBus0 slot 0xf offset 0x1100000 Onboard device sparc9 ipl 12
zs0 is /sbus@1f,0/zs@f,1100000
```

In addition, in Solaris, check for the detected tape drives by typing `ls /dev/rmt`.

## Topology Confirmation

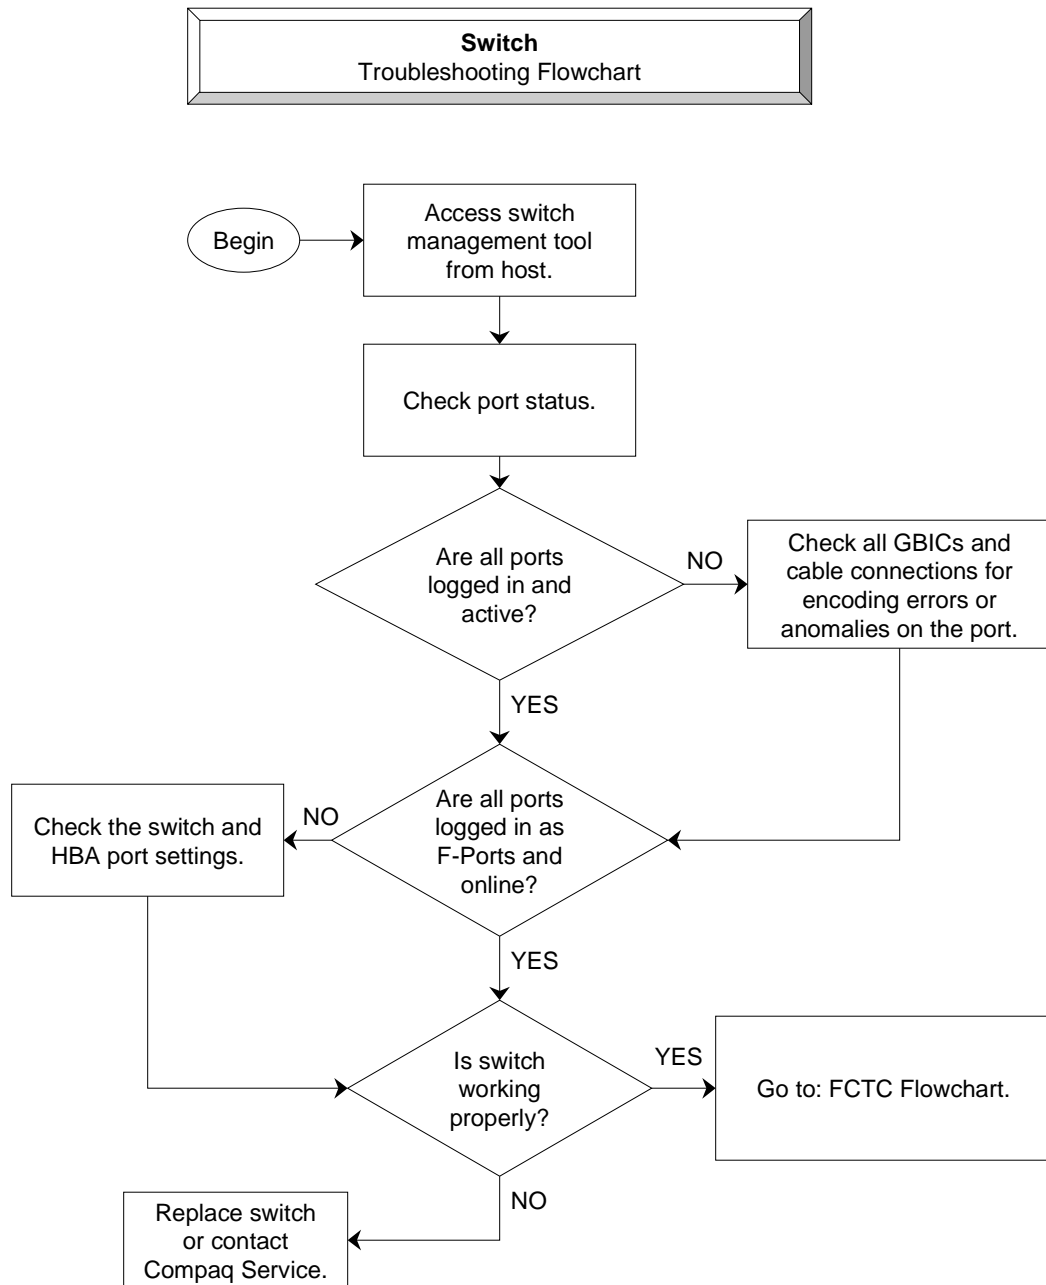
To confirm that the topology is set for switched fabric, use the following methods.

For Windows NT/2000, ensure that the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lp6nds35\Parameters\Device\DriverParameter` registry entry includes the `Topology=1` value.

When using the KGPSA HBAs in a Tru64 UNIX or Sun Solaris server, the default topology is set for switched fabric. You can use the Emulex utility that ships with the HBA drivers to confirm the topology.

NetWare logs into the switch as an FL-port. You do not have to confirm the topology setting for NetWare.

## Checking the Switch



## Switch Management

Web Management Tools let you interact dynamically with any switch in the SAN to monitor status and performance. To manage a switch over the web, launch SWCC or log on to a switch from a host with a Java-enabled web browser by means of the Internet or an intranet.

Web Management Tools display information you can use to make overall topology decisions and provide a means to make administrative changes. Security is enforced by a user name and encrypted password.

Web Management Tools feature:

- Switch identification in the network
- Fabric topology and routing information
- Switch and port configuration information
- Real-time graphical switch and port status and statistics report
- Graphical representation of port performance (polled up to every 2.5 seconds depending on the operating system and browser used)
- Four minutes of aggregate bandwidth throughput shown graphically, scaled dynamically, and based on activity (polled up to every 2.5 seconds depending on the operating system and browser used)
- Secured management through either the GUI or a Telnet session
- Five screens for monitoring information
- Two screens for administrative interfaces
- Help functions including glossary help, online HTML help, and pop-up help for error conditions
- Out-of-band interface through a 10/100BaseT Ethernet connection
- In-band interface over a Fibre Channel link

The General Switch View page is used most often in troubleshooting a switch. Access this page for a switch by double-clicking that switch's icon on the Fabric View page.

The General Switch View page shows a graphic representation of the front panel of each connected switch. Normal long-term monitoring is conducted from this page, which provides a real-time view of each switch's overall health and status in the fabric. Switch status is updated every 1 to 2.5 seconds, depending on the operating system and web browser used.

The upper half of the screen shows port and LED indicator status, and the lower half displays general switch information. The following table summarizes the components of the General Switch View page. Right-click a component on this page to display a corresponding Help dialog box.

### General Switch View Page Components

Component	Description
Port modules	Link to the Port Detail View page
Admin button	Links to the system administration pages
Perform button	Links to the Performance View page
Telnet button	Launches a Telnet session
Thermometer	Indicates the highest temperature from the last data sample. Click the thermometer to display the temperature readings from five switch sensors.
Fans	Indicate the status of the switch's fans
WWN	Unique numeric identifier for each switch
Domain ID	A number from 0 to 239 that uniquely identifies each switch in a fabric
Role	The switch's role: <ul style="list-style-type: none"><li>■ Principal — The switch is the principal switch.</li><li>■ Subordinate — The switch is enabled, but it is not the principal switch.</li><li>■ Disabled — The switch is disabled.</li></ul>
State	The switch state: Online, Offline, Testing, or Faulty
Firmware	The firmware version
EtherIP	The switch's Ethernet IP address
Ether NM	The Ethernet subnetmask
FC IP	The Fibre Channel IP address
FC NM	The Fibre Channel subnetmask
Gateway	The gateway address

Verify that the host, storage system, and tape controller are in the same zone.



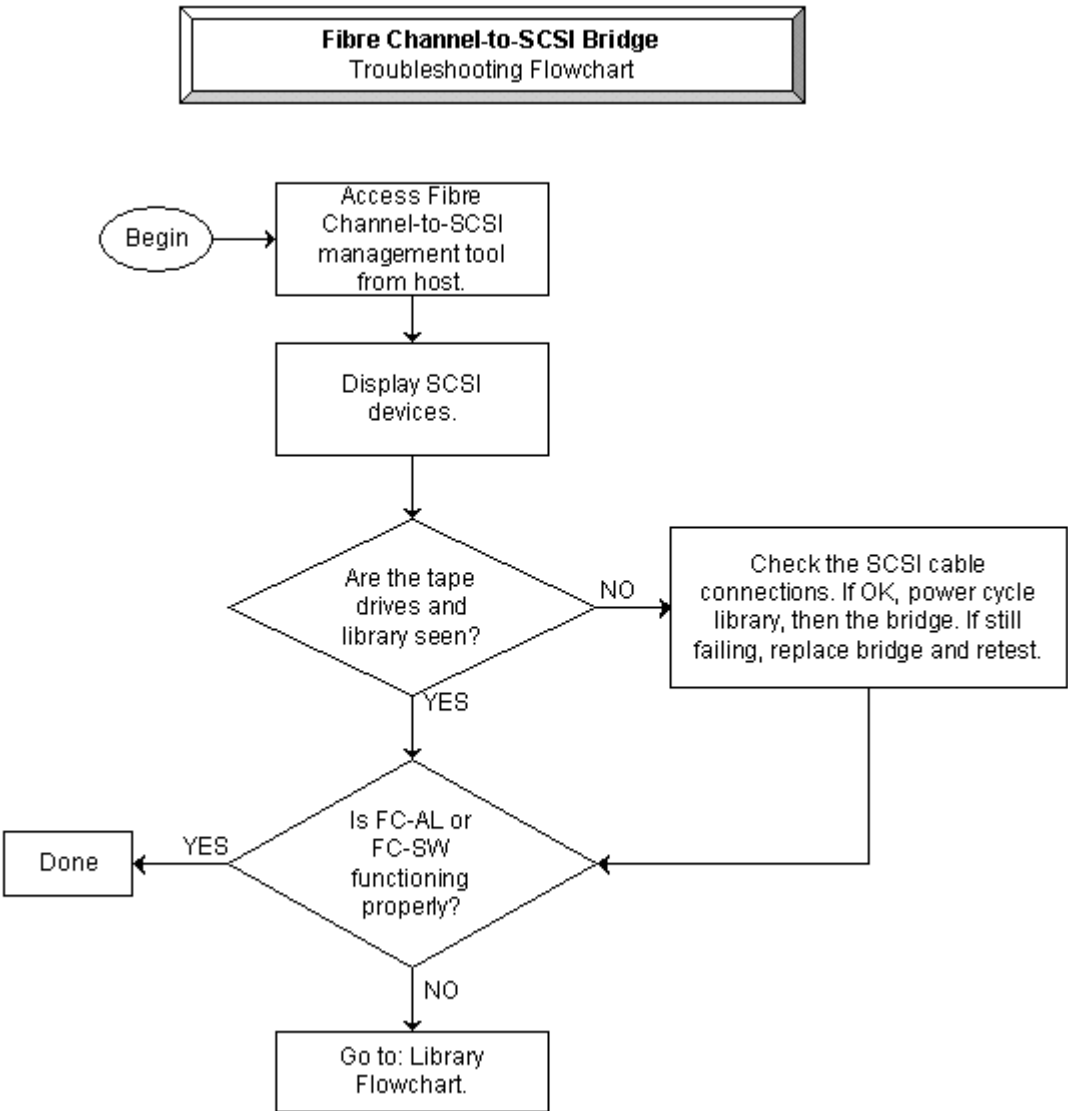
---

**Note**

For more information on troubleshooting switches, see the appropriate installation and hardware guide.

---

# Checking the Fibre Channel-to-SCSI Bridge



## Troubleshooting the Modular Data Router

Perform the following steps to troubleshoot the Modular Data Router (MDR):

1. Verify AC power.
2. Inspect cables.
3. Perform POST.

### Verify AC Power

The MDR has an LED on the rear panel that illuminates when it is connected to an AC power source and the power switch has been turned on. If this LED does not illuminate, verify the following:

- The MDR has been plugged into an active AC power source.
- All peripheral devices are connected to AC power.

### Inspect Cables

There are two LEDs on the MDR Fibre Channel module. Normally, these indicate the condition of transmitted and received signals.

When a cable attached to an operating loop or fabric is plugged into the GBIC in the module, both LEDs should illuminate. If they do not, the likely causes are a bad GBIC or Fibre Channel cable.

If a pattern of alternating periods of rapid flashes followed by a short period of non-illumination is displayed, there may be an internal problem in the Fibre Channel module.

Inspect the SCSI cabling between the MDR and all attached SCSI devices. Check for the following:

- Loose, dirty, broken, or bent cabling and connectors
- Cables resting on the SCSI module's plastic handle



---

**Note**

The VHDCI SCSI connectors used on the MDR cannot support the weight of the attached SCSI cables. Always make sure the plastic handles are used to support the cable.

---

The MDR SCSI modules are marked with icons indicating whether they contain Low-Voltage Differential/Single-Ended (LVD/SE) or High-Voltage Differential (HVD) transceivers. Verify that the target SCSI devices have the same type of transceivers. The MDR will not be able to communicate with the SCSI devices if the transceiver types do not match.



Each SCSI module has one LED indicator for each SCSI bus. The LED is an indicator of activity on the SCSI bus. Depending on the level of SCSI activity, the LED could be off (indicating no SCSI activity) or it may appear to be solidly illuminated during high rates of SCSI activity.

If the LED appears on all the time but there is no apparent I/O activity happening:

- The SCSI bus may be hung.
- This could be an indication of a faulty SCSI module in the MDR.
- It could be an indication of a faulty component outside the MDR such as a SCSI cable or tape drive.

### **Power-On Self-Test (POST)**

Immediately after the MDR is turned on, it performs a power-on self-test to verify its internal hardware is working properly.

- The entire process should last approximately 90 seconds.
- When POST starts, all LEDs on the Management Module will illuminate and then go through a series of indicator patterns as various portions of the internal hardware are tested.
- After POST is complete:
  - The test is completed properly if the Management Module LEDs display a *single* LED moving back and forth.
  - If the LEDs illumination pattern *does not* change for more than 90 seconds, there could be a hardware failure in the MDR.

If this happens:

- ◆ Power off the MDR.
- ◆ Reseat all of the option modules, making certain to tighten the modules in with their thumbscrews.

If POST stops with a static LED pattern again, connect a terminal or terminal emulator to the MDR's DB-9 serial port to monitor POST progress and error messages.

The DB-9 serial port provides an advanced configuration and management interface. This serial port reports POST progress and system error messages, as well as providing command line interfaces for boot management and application management.

To monitor POST progress:

1. Connect one of your servers to the MDR by attaching the provided DB9 serial NULL MODEM cable to the serial port on the MDR.
4. Select *Start* → *Programs* → *Accessories* → *Hyperterminal* → *Hyperterminal* and enter a name for your hyperterminal in the New Connection dialog box.
5. Select an *icon* to represent the tape controller and click *OK*.
6. The Phone Number dialog box displays. From the *Connect using:* pull-down menu, select the communications port assigned to your serial port; click *OK*. (You might have to enter an area code.)
7. Enter the following information in the Properties dialog box; select *OK*.

a. Bits per second (baud rate)	<b>9600</b>
b. Data bits	<b>8</b>
c. Parity	<b>None</b>
d. Stop bits	<b>1</b>
e. Flow control	<b>None</b>
8. Select *File* → *Properties* → *Settings* and from the Emulation pull-down menu, select *VT100* → *OK*. The Power-On menu should display.
9. Press *Enter* if the Power-On menu does not display on its own.

## Troubleshooting the Network Storage Router

### LED Indicators

The LED indicators on the router are useful for diagnosing various problems:

- The Fibre Channel LEDs indicate Fibre Channel activity and link status. If one of these indicators does not light or stays continually lit without any corresponding SCSI bus activity, it may indicate a problem with the Fibre Channel link. Verify the Fibre Channel configuration.
- The SCSI Bus LEDs indicate SCSI activity. These indicators are lit during power up, configuration, and when the unit is transferring data. If the SCSI indicator stays continually lit without any corresponding Fibre Channel LED activity, it may indicate a problem with the SCSI bus configuration. Verify the SCSI bus configuration.
- The Ethernet LEDs indicate activity and link status. If one of these indicators does not light or stays continuously lit, it may indicate a problem with the network connection. Verify the network connection. The port must be connected to a 10/100BaseT Ethernet network to function properly.

### Verifying SCSI Bus Configuration

Items to check include:

- **Termination**—Problems with termination can cause intermittent or hard failures. A SCSI bus must be terminated on both ends. Termination problems are common when both narrow and wide devices are on the same bus.
- **Bus Type**—On an LVD SCSI module, SE and LVD devices can be connected to the same bus. However, if one SE device is detected during power up, communication to all devices will convert to SE mode.
- **Device ID**—Each device on a SCSI bus must have a unique ID. Verify configured IDs are not in use by other devices on the same SCSI bus.
- **Cabling**—Check SCSI cables to verify they are functional. SCSI rules for total length, distance between devices, and stub length must be followed. Connections should also be checked and reseated if necessary.
- **SCSI Devices**—Verify that the SCSI devices on a particular SCSI bus can be seen in the Configuration Menu of the router. If the router cannot see the devices, verify SCSI configuration, cabling, and termination.

### Verifying Fibre Channel Configuration

If SCSI devices are recognized on the SCSI buses, but do not appear to the Fibre Channel host, it may be that the Fibre Channel link is not properly established. Most hubs and switches have link indicators showing link status. When the router is connected and powered on, this link indicator should show a good link. If it does not, check the cabling or connections.

One method of verifying link integrity when connected to a functional host, involves disconnecting and then reconnecting the Fibre Channel cable. This

procedure should cause momentary activity of this indicator as the link reinitializes.

Additionally, verify that the media type of the router and the attached hub, HBA, or switch is of corresponding types.

---

**Note**

By default, the Fibre Channel port speed is set to 1 Gb/s. Changes to the Fibre Channel port speed must be manually set, such as for 2 Gb/s. If set incorrectly and plugged into a Loop or Fabric, the unit may receive Framing errors due to the incorrect Fibre Channel link speed.

---

## **Verifying the Router Configuration**

If you are in doubt about the configuration or about the location of the error, restore the router to the factory default configuration and proceed to configure the unit one step at a time, verifying the functionality of the configuration after each change is made.

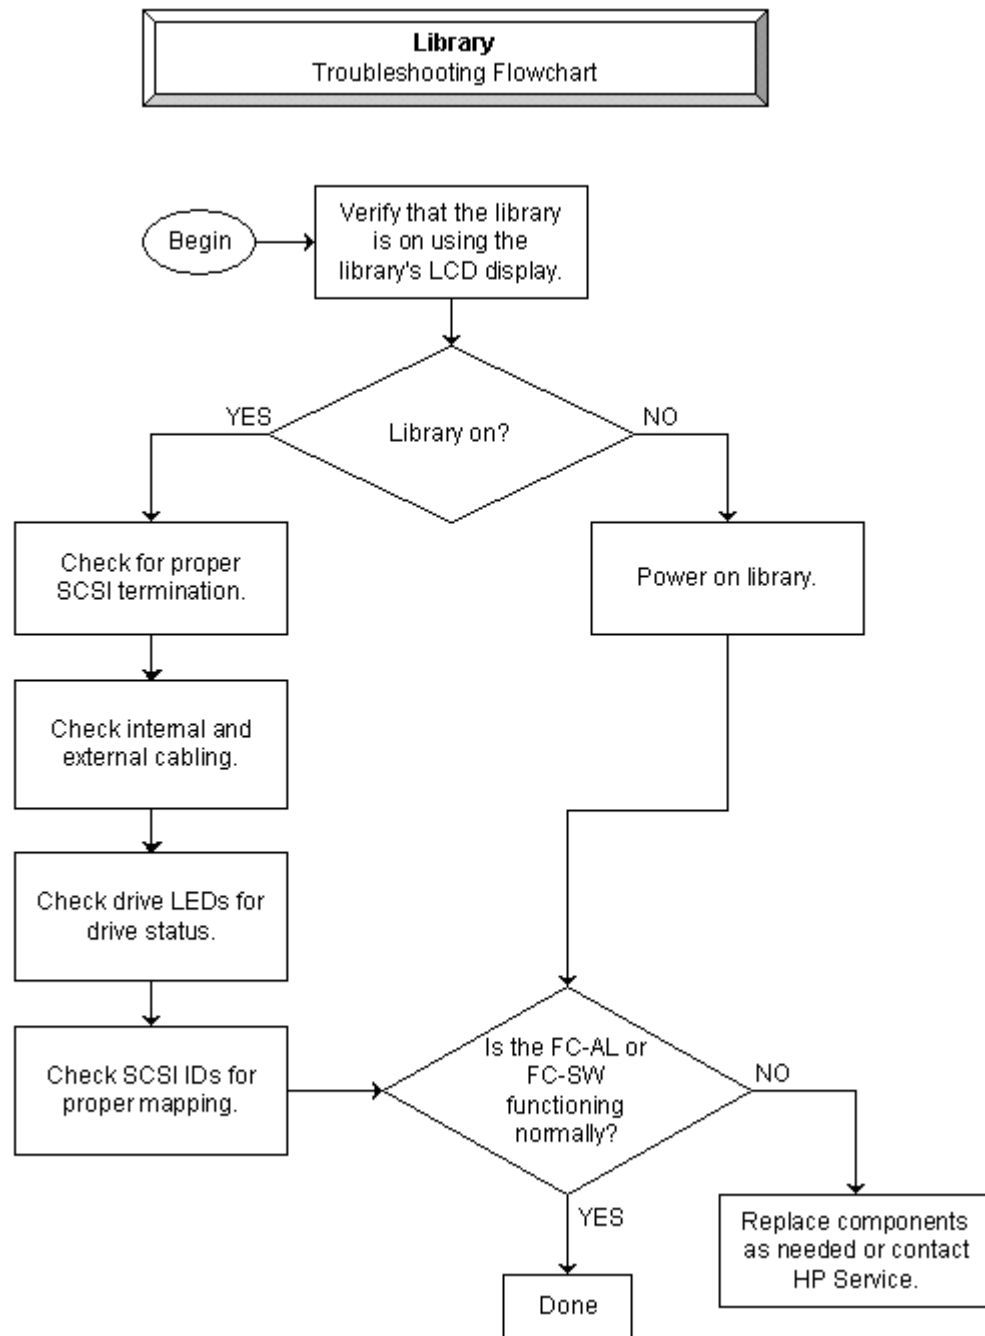
## **Verifying Devices**

It may be useful to connect the SCSI target devices directly to a SCSI interface (for example, a host SCSI bus) to verify that the devices are functional.

## **Verifying the Host Configuration**

In some cases, it may be that the Fibre Channel HBA or host device driver may not be working properly. Check the configuration of these elements. It may be useful to check the release notes for the device driver to see if there are any specific issues or required configuration. It may also be useful to ensure that the current version of the HBA driver is being used.

## Checking the Library



## Typical Causes of Library Failure

When diagnosing hardware failures, observe the physical characteristics of the library with power off and power on. You can often locate failures during a visual inspection.

If, after all the troubleshooting, you find that the library is the source of the hardware failure, consider these common causes:

- Mechanical failure of tape drive
- Tape drive booted with tape inside
- One or more cables disconnected
- Cartridge lock assembly in wrong position
- Object blocking sensor path
- SCSI IDs not configured properly

## SCSI IDs

The order of the SCSI ID assignments for the robotics controller and drives is important when using the FCTC-II. The robotics controller must be set to SCSI ID 1. Drives must be set to SCSI IDs 2 and 3 at the hardware level. For example, the following table lists the IDs for a TL895 with seven drives.

Device	FCTC #1		FCTC #2	
	Bus 0	Bus 1	Bus 0	Bus 1
Robot	1			
Drive 0	2			
Drive 1	3			
Drive 2		2		
Drive 3		3		
Drive 4			2	
Drive 5			3	
Drive 6				2

SCSI IDs for TL895 with Seven Drives

## Library Troubleshooting Strategy

Conduct repeated tests and visual inspections of the library, and check cables and other connections to ensure that they are joined correctly. Do not simply rely on the diagnostics tools.

When you are not sure about how to proceed, remember to **run, watch, and record**.

1. **Run** one or more tests.
2. **Watch** the mechanisms.
3. **Record** the error messages.

## POST

POST is often the key to troubleshooting. During POST, watch the mechanisms of the tape library and note where they stop.

Observe what the unit does during the test. If an error is reported, what action was the unit attempting when the failure occurred? Visually inspect the mechanisms. Are all the cables connected? Does anything seem abnormal?



---

### Note

Watch the unit perform normally as much as possible so that you can detect problems visually.

---

## Troubleshooting the SSL2020TL and TL891 Libraries

The following problems can cause the SSL2020TL and TL891 Libraries to malfunction:

- Platform problems
- General drive errors

### Platform Problems

Incorrect installation and configuration cause platform problems. In this case, the Library appears to be operating normally, but no data can be interchanged. You might or might not get an error code on the Control Panel. To identify an error caused by this type of problem, check your installation and configuration setup.

### General Drive Errors

General drive errors usually result from a miscommunication between a module's processor and a drive processor or a mechanical malfunction within the Library.

Both platform problems and general drive errors display an error message and a Fault System Code (FSC) on the Control Panel. Use the FSC to report errors to your Technical Support representative, or in some cases, to determine a recovery procedure.

A simple error recovery procedure displays at the bottom line of the Control Panel display. You can clear some errors by pressing the *Enter* button on the Control Panel or by cycling the power to the Library.



## Troubleshooting the MSL5000 Library Family

There are two main types of problems that can cause the MSL5026 Library to malfunction or fail to perform correctly: platform problems and general drive errors. Some errors cause Fault Symptom Codes (FSCs) to be displayed on the GUI touch screen, along with a description of the fault. Error Recovery Procedures (ERP) are available for each FSC. FSCs are described in the MSL5000 User Guide.

### Platform Problems

These errors arise out of incorrect installation and configuration. The most common characteristic of this type of error is that the drive seems to operate normally, except that no data can be interchanged. You might not get an error code on the GUI touchscreen. To identify an error as this type of problem, check your installation and configuration setup.

### General Drive Errors

These errors usually result from a miscommunication between the drive processors or involve a mechanical malfunction. In most cases, both platform and general drive errors report an error message and an FSC on the GUI touch screen. The FSC can also be used to determine a recovery procedure.

### Error Recovery Procedures

ERP Code	Corrective Action
C	Cycle power the library using the Power option on the GUI touch screen. Wait 30 seconds to power up again.
D	Turn off the power to the library and inspect the connectors and cables.
F	Invalid operation. Set parameters correctly and try again.
G	Call Technical Support.

### Fault Symptom Codes

Fault symptom codes, along with the suggested error recovery procedures, are listed in the MSL5000 User Guide.

## Troubleshooting the TL895 Library and ESL9000 Library Family

The following sections outline many of the messages and problems you might encounter during the initial configuration and operation of the ESL9000 Library:

- Messages
  - Status
  - Sense Data Value
- Problems
  - Start-up
  - Control Panel
  - Robotics
  - Operating
  - Tape drive

### Messages

Some events produce status messages on the control panel; others produce error messages called *sense data values*. Sense data value messages display on the Control Panel or host monitor, and include a number and description of the error that occurred. Information on messages can be found in the user guide.

### Problems

In addition to status messages, other problems might occur during library operation. The following tables list the most common problems and the steps you can take to resolve them. For problems that cannot be resolved using the information in this module, notify your system administrator or field service engineer.

## Start-Up Problems

Problem	Corrective Action
The library does not power on.	Ensure that the library power switch is set to the I (on) position and the power cord is connected to a grounded electrical outlet.
The library or tape drives do not respond on the SCSI bus.	Verify that each SCSI device on the same SCSI bus has a unique address and that the last device is properly terminated.
The library reports “not ready” during initialization.	Determine the failure type by checking any previous error codes returned to the host computer. Correct the cause of the error.
One or more tape drives fail to spin up during start-up.	Check all SCSI cabling and termination on the back panel of the library. If necessary, contact your field service representative about replacing the drives.
There is a cartridge in the gripper, preventing system calibration and inventory.	Manually remove the cartridge from the gripper. Then restart the library.
The library starts up in offline mode.	Press the Standby button. Verify that the library switches to online. You can use the Control Panel to select either online or offline mode at power-up.

## Control Panel Problems

Problem	Corrective Action
The Control Panel is blank or does not respond to touch.	Contact an authorized field service engineer.
An error or information message displays.	Write down the details of the message, including the SK, ASC, and ASCQ numeric values. Then press <i>OK</i> to clear the message. For instructions on resolving a possible error, refer to the user guide.

## Robotics Problems

Problem	Corrective Action
The robot does not move at power-up.	Verify that all internal packing materials (foam pads and tie wraps) were removed during the installation procedure. Check the Stop and Standby buttons; ensure the library is online and the Stop button is released. Ensure that all library doors are closed.
The gripper partially grips a tape cartridge.	Issue a <i>Move Cartridges</i> command to move the cartridge from the gripper to an empty storage element.
The bar code reader on the gripper fails.	Verify that nothing obstructs the reader. Then restart the library. If the problem continues, contact an authorized field service engineer.
The robot times-out or fails during an operation.	Confirm that the tape cartridge involved in the operation is properly positioned in the bin or drive and ready to be picked. Verify that the robot is not obstructed in any way. Retry the operation. If it still fails, contact an authorized field service engineer.
The robot drops a cartridge.	Open the storage array door safely. (To do so, press the <i>Standby</i> button first to place the library offline, and then unlock and open the door.) Retrieve the cartridge, orient it properly, and place the cartridge in an empty storage bin. (Do not try to place the cartridge in the gripper.) Afterwards, close and lock the storage array door and press the <i>Standby</i> button again to return the library to online mode.
A cartridge is in the gripper at start-up, when a move command is requested, or after a place command is executed.	Press the <i>Standby</i> button to place the library offline; then open the storage array door. Manually remove the cartridge from the gripper and place it in an empty bin. Then close the door and press the <i>Standby</i> button to return the library to online mode.
The gripper does not have a cartridge after completing a pick command.	Verify that a cartridge is in the source location. Retry the command. If the pick operation fails again, contact an authorized field service engineer.

## Operating Problems

Problem	Corrective Action
The host computer cannot communicate with the library.	This might be a SCSI bus time-out or a premature disconnect problem. Check cable connections, cable length, SCSI addresses, and termination. Restart the host and the library. If the host and library still are not communicating, contact an authorized field service engineer.
A cable or terminator is disconnected from the library back panel.	Reconnect the cable or terminator.
A tape cartridge (medium) is reported not present.	The gripper cannot sense a tape cartridge in a particular storage element even though the inventory reports that it is present. Verify that the designated cartridge is present. If it is, ensure that it is properly seated. (For a tape drive, ensure that the cartridge is completely unloaded.) Then retry the command. If the error persists, contact an authorized field service engineer.
A move command failed.	Check the source and destination elements. The source element should hold the cartridge to be moved; the destination element should be empty. Ensure that the gripper is empty and all actuators are free of obstruction. Also, verify that the library is online and the <i>Stop</i> button is released. Retry the command.
A flash memory error is reported.	Contact an authorized field service engineer.
A <i>maximum temperature exceeded</i> warning displays.	Power off the library and allow it to cool down. Lower the room temperature, if possible, and increase ventilation around the library. (If the operating temperature is too high, the library automatically shuts down until the temperature drops.)

## Tape Drive Problems

Problem	Corrective Action
The library is unable to communicate with a drive, indicated by a <i>Drive Communication Time-out</i> error.	Contact an authorized field service engineer.
Failure to mount or read/write with new or used cartridge.	The cartridge is bad. Retry the same operation with a different cartridge. OR the drive heads are dirty. Use a Cleaning Tape III head-cleaning cartridge.
Application software-generated initialize command fails with parity error.	The tape calibration failed. Use a Cleaning Tape III head-cleaning cartridge or try a different data cartridge.
Green <i>Operate Handle</i> indicator stays lit and tape does not move; yellow <i>Tape in Use</i> indicator is lit, but does not blink.	There was a cartridge loading error. Dismount the cartridge and inspect it for a mispositioned leader. If the drive leader is not in the correct location, call your HP authorized reseller or service provider. OR inspect the drive for a damaged, misplaced, or unhooked leader. If the drive leader is not in the correct location, call your HP authorized reseller or service provider.
All right-side or left-side indicators are blinking.	The drive failed the self-test or detected an error during operation. Ensure that the drive handle is in the down position. Try to clear the error by pressing the <i>Unload</i> button, or power off and then back on. If the error does not clear (the tape does not rewind and unload, and the lights blink), you have a hardware failure; call your HP authorized service provider.
The tape drive does not eject a cartridge.	Reset the library and retry the unload command. If the tape still does not unload, press the <i>Standby</i> button to place the library offline, open the front door, and manually unload and eject the cartridge.
A drive handle error occurs.	Contact an authorized field service engineer.
Unable to back up or copy data to tape.	The cartridge is write-protected. Set the tab on the cartridge to write-enabled or, if software is protected, modify that protection. OR no tape is in the drive. Insert a tape.
<i>Write Protected</i> indicator flashes orange.	Perform the head cleaning procedure. If the error repeats, try another tape.
<i>Tape in Use</i> and <i>Write Protected</i> LEDs flash rapidly in unison.	Eject the tape. Perform the head cleaning procedure. If the error repeats, replace the tape. Power off and power on the drive. If the error repeats, replace the drive.
After applying power, nothing happens. All indicators are off, except green <i>Operate Handle</i> indicator.	No tape is loaded. Load a tape.
Drive is not available to the system.	Check the AC power source to be sure power is getting to the drive. OR check the SCSI ID switch settings. OR ensure that the power cable is plugged in and that the SCSI cable is seated properly.

## Learning Check

1. The Media Robot Utility (MRU) can be used to perform which of the following?
  - a. Move a cartridge in a library from a drive to a slot in the library.
  - b. Load firmware into a tape drive.
  - c. Notify the administrator of faults with the library through email.
  - d. View errors on the Network Storage Router.
2. Error events caused by Fibre Channel are detected and communicated to a higher-level protocol without impact to performance in the secondary storage environment.  
☐ True  
☐ False
3. A LIP occurs when a:
  - a. Data frame is sent.
  - b. Job is restored.
  - c. Server is restarted.
  - d. Data frame is lost.
4. If you suspect that the tape drive is exhibiting faulty behavior, but it passes POST without errors, what should you try next to diagnose the problem?  
.....  
.....
5. A typical cause of a tape library failure is: (Select all that apply.)
  - a. One or more cables disconnected
  - b. An object blocking the sensor path
  - c. A GBIC removed from the library
  - d. A bad host adapter

### Objectives

After completing this module, you should be able to:

- Describe the Library and Tape Tools and its features.
- Describe the HP management support for the Enterprise Backup Solutions.
- Identify the limitations of Fibre Channel on secondary storage compared with primary storage.
- Diagnose problems and perform repairs to the system using the troubleshooting flowchart method and tools provided by HP.



## HP StorageWorks Library and Tape Tools

HP StorageWorks Library and Tape Tools (L&TT) is a collection of storage hardware management and diagnostic tools for tape, tape automation, and archival products. L&TT assembles these tools into a single, convenient program.

### Software Features

L&TT offers the following features:

- **Installation Check**—L&TT guides you through a basic installation check of your product. The software assists the user in choosing an appropriate HBA and SCSI ID(s), ensuring that the device is detected by the system, and verifying key device functionality. The installation check feature is essentially HTML documentation that helps with the most common generic installation issues while also describing how to use L&TT to verify the device installation.
- **Device Identification**—L&TT clearly identifies the storage products connected to the system, along with key information on product configuration and status.
- **Troubleshooting Tests**—L&TT provides various tests to verify product functionality or to isolate product issues. Tests include device self-tests, read/write tests on drives, exerciser tests for autoloaders and libraries, and specific device utilities.
- **Firmware Upgrades**—L&TT provides a convenient way of updating product firmware, enabling users with an Internet connection to take advantage of ongoing enhancements. The software can be configured to check the Web automatically for firmware updates for connected devices, or users can manually check the Web for updates if the automatic update feature is not desired. If updated firmware is available, the program notifies the user, and the updates can easily be copied to the system. With libraries, users can upgrade the library and the embedded drive firmware in the same operation. Wherever possible, the embedded drives are updated in parallel to reduce time.
- **Support Ticket Generation**—If you experience a problem with a storage product, L&TT can generate a support ticket that includes essential information for troubleshooting the problem. As an alternative to phone support, you can e-mail the support ticket to a support center for assistance. This information streamlines the support process and enables the support staff to better serve you if a support call is made later. When a support ticket for a device is generated, L&TT performs a Device Analysis test on the device. The support ticket contains generic information about a device, as well as the results of the Device Analysis test. The Device Analysis test can be performed by itself, but HP recommends generating a support ticket because the resulting data is presented in a more useful format.
- **Automatic Notification of Web Updates**—If a connection to the Internet is present and web updates are enabled in the tool preferences, L&TT automatically informs you of the following updates, if available, each time the program is started:
  - New versions of L&TT

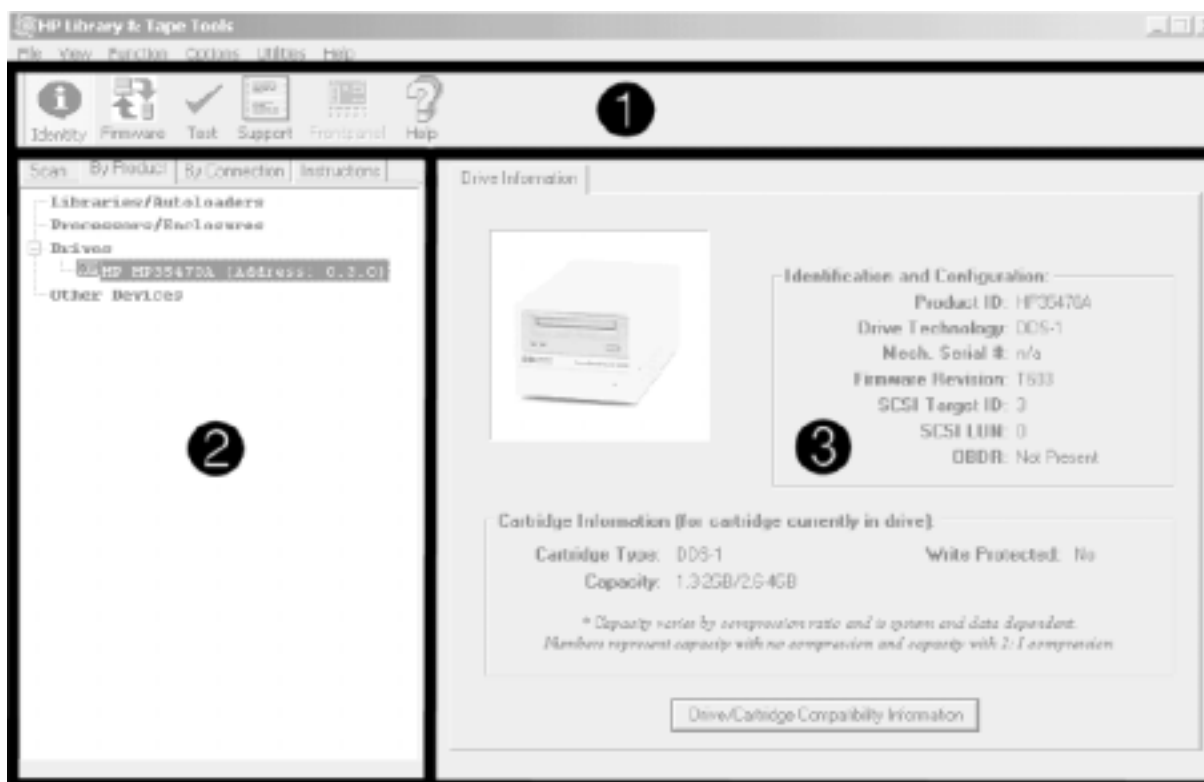
- New firmware files for connected devices
- New device-specific functionality (such as new or updated tests) for connected devices

## Supported Products and Operating Systems

For a complete listing of compatible products, refer to the L&TT Compatibility Matrix found at <http://www.hp.com/support/lttcompatibility>. The level of functionality that L&TT offers for each device varies depending on features of the device, and the degree of device integration into L&TT.

The Windows version of L&TT uses a graphical user interface (GUI), whereas the NetWare, HP-UX, and Tru64 versions of the program use a command screen interface (CSI).

## Application Window Layout



The three sections of the L&TT main screen are:

1. **Taskbar**—This section contains buttons that provide quick access to the main functions of L&TT and to the online help system.
2. **Device List**—This is a multi-function window that offers several options on the following tabs:
  - **Scan**—This option provides either a summary status or detailed information (depending on whether **Show Details** or **Hide Details** is selected) about the bus scanning process. If a problem is encountered during the scan, this information may help in determining the cause. When the scan completes successfully, the device list automatically

switches to the **By Product** tab. The **Scan** tab also lets you rescan the bus. If any devices have been hot-swapped or powered on after the OS has booted, in most cases, the scan feature can discover those devices without requiring a reboot of the system.

- **By Product**—This option shows a list of all the products connected to the system. The list is grouped into the following four categories:
  - ◆ Libraries and autoloaders
  - ◆ Drives
  - ◆ Enclosures and processors
  - ◆ Other devices

The three number fields listed after the device represent the device address. Each field in the address is separated by a period: the first field represents the HBA channel, the second field represents the SCSI ID, and the third field represents the LUN.


- **By Connection**—This option shows all products connected to the system, grouped by the HBA they are connected to. This view makes it easy to see which devices are connected to the same bus as the device in question, and may help in understanding system I/O performance issues.
  - **Instructions**—This window contains brief instructions on how to use the selected screen. This view can be disabled in the software preferences.
3. **Device Information screen**—All the main functions of the program are displayed in this window. The content of this window depends on the device and tool function selected.

## Using the Device Information Screen

When you select a product from the **Device List**, the **Device Information** screen displays information relevant to the device. The **Device List** also changes to display instructions on this page (if the instructions preference is enabled). If another tool function is currently active, then clicking the **Identity** button on the toolbar opens the **Device Information** screen. The **Device Information** screen provides an overview of the selected hardware device and its current configuration and status.

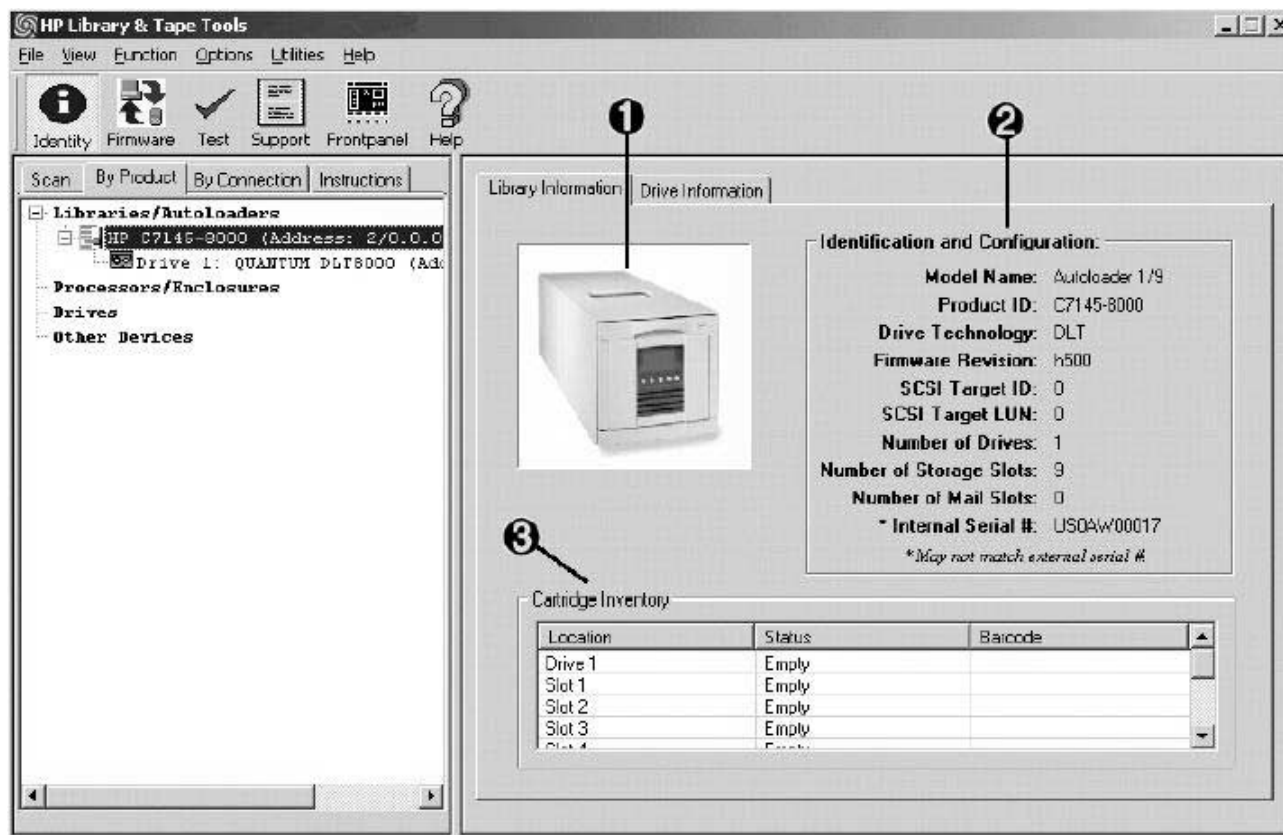
## Standalone Devices

The **Device Information** screen, showing a standalone device, contains a single tab labeled **Drive Information**. The **Device Information** screen has four main sections, as shown.



1	This provides a visual representation of the selected product (the images are grayscale)
2	This provides the following general information on the product: <ul style="list-style-type: none"> <li>• Product ID (Inquiry String)</li> <li>• Drive Technology (DDS/DL/LTO)</li> <li>• Mech Serial # (Drive serial number)</li> <li>• Firmware Revision</li> <li>• SCSI Target ID</li> <li>• SCSI LUN</li> </ul>
3	If a data cartridge was inserted in the drive when the identity tool was selected (or when the current screen was "refreshed"), information regarding cartridge type, capacity, and write protect status is shown here.
4	Click this button to display a list of media that is compatible with the product.

## Library and Autoloader Products

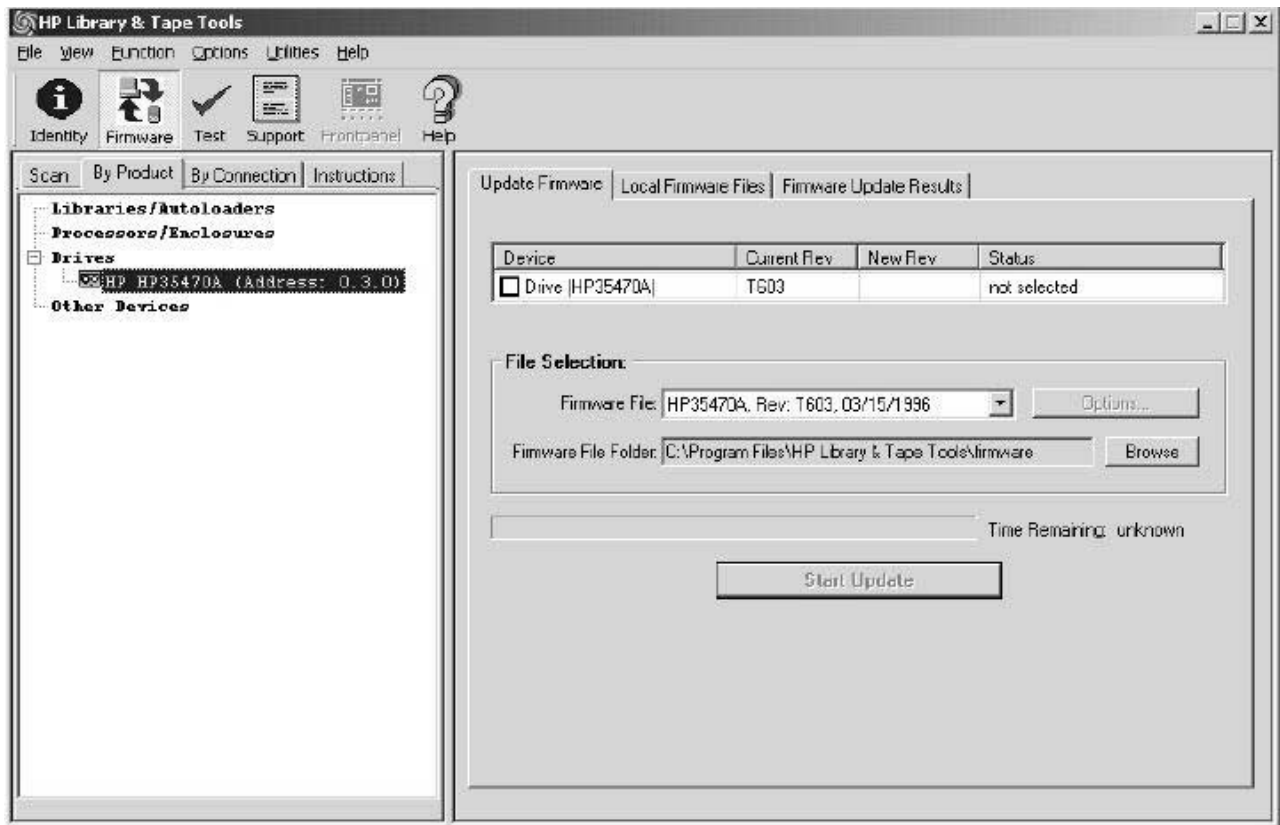


When a library or autoloader is selected in the **Device List**, the **Device Information** screen displays two tabs: one for library information and one for the drives contained within the library.

1	This provides a visual representation of the selected product.
2	This area provides the relevant information on the selected library product.
3	This provides a cartridge (media) inventory for all the drives and slots. The current status (full or empty) of all the storage elements and the barcode number (if barcode reading is available) is displayed here.


To see information for the drives within the library, select the **Drive Information** tab. This tab displays a screen very similar to the one for standalone products, with the addition of the **Selected Drive** dropdown menu. Because a library can contain multiple drives, you can select and view information for each drive using the **Selected Drive** dropdown menu. All of the drives in the library are available in the dropdown menu.

## Using the Firmware Management Screen




The Firmware Management functionality of L&TT provides a convenient way to easily upgrade the firmware of connected devices. Library firmware and firmware of drives within the libraries can be upgraded as well. In many cases, multiple drives within the library can be updated simultaneously to save time. L&TT uses special firmware files that associate the firmware with a particular product. This feature prevents the user from downloading the wrong image and ensures that the downloaded firmware is the same or a newer revision than the current device firmware.

Under the **Local Firmware Files** tab, a **Get Files from Web** button allows the administrator to download firmware files to the local management host.



## Using the Tests Screen



L&TT provides numerous tests for analyzing storage hardware. L&TT tests are SCSI script-based. A SCSI script is a sequence of SCSI commands sent to a device, usually for the purpose of testing the device's functionality. SCSI scripts can also be used to perform non-diagnostic utility functions, such as rewinding a tape, recalibrating a tape head, or “exercising” a tape-moving robot arm. L&TT uses the following three types of SCSI scripts:

- **Built-in tests** are hardcoded into the program and are installed during the L&TT installation process.
- **Script-based tests** are additional tests based on scripts that are stored in the *scripts* directory in the L&TT installation directory. The script-based test files have the \*.*tst* extension. Users can download new and updated scripts. With the Windows version of L&TT, the management of script files is accomplished through the **Get files from Web** item on the **File** menu. HP is constantly adding scripts for new devices and updating scripts to enhance the testing functionality of older devices. As new scripts are developed, they are published to the external website in the scripts section.
- **SCSI utility scripts** provide access to generic SCSI operations. These scripts are accessed by selecting **Run Script...** from the **Utilities** menu and have the \*.*scc* extension.



## Built-In Tests

- **Connectivity Test** examines the ability of the host to communicate with the device. This is a test of the SCSI communication path only; no media is required. This option is useful when trying to establish that the SCSI cables and adapters are working correctly. This test is available for DLT, LTO, and DDS drives only.
- **Device Self-Test** executes the internal self-test of the device, and any failure is reported. This test is useful in confirming that the device itself is operating properly. Not all devices support an internal self-test.
- **Device Analysis Test** runs a device-specific script that carries out a detailed examination of the device. The test examines all available logs from the device to look for any known issues, customer use issues, and common error patterns. The test also provides summary information and presents suggestions as to the cause of certain errors and any remedial actions the user should take. This option is particularly useful if you suspect that there may be a problem with the device.
- **Read/Write Test** writes and reads data from tape. In addition to verifying that the data read-back matches the data written, the read/write error rates are also checked, and the test fails if the error rates exceed allowable limits. This test can be used to verify the ability of the drive to read and write data independent of the storage application. Use a blank tape when performing this test because this test is destructive to data on the tape. This test is available for all tape drives.
- **Media Validation Test** determines if the drive can read the data on any media currently in the drive. The test is used to determine the quality of the media. The test reads data from the media and calculates the error rate. If the error rate exceeds the maximum allowable threshold, the test will require a cleaning cartridge and a second tape to determine if the problem is with the drive or the tape. The amount of data read depends on the drive type and coverage level of the test, but it generally varies between a few megabytes and one gigabyte. No data is written to tape and the test is not destructive to data. This test is useful if the user suspects that a backup failed because of a bad tape. This test is available for all tape drives.

---

### Note

This test requires that a minimum amount of data is already written on the tape. Otherwise, the test will fail.

---


- **Media Analysis Test** reads all data on a cartridge from beginning to end. The test determines if the media is readable, and it also provides statistical data such as block size, compression ratio, and speed. This test is available for DLT, LTO, and DDS drives only.
- **Library Exerciser Test** takes a test cartridge from a mail slot and performs random slot-to-slot moves. The test then moves the test cartridge into every

available drive and performs a short drive test. This test is intended to exercise the robotics and robot drive hardware.

### Script-Based Tests

- **Compression Test** writes a data pattern of known compressibility to tape. The test then rewinds the tape and reads the file just written. The read and write hardware compression ratios are calculated to ascertain that the hardware compression is enabled and working. If the script finds that hardware compression has been disabled by the host or by switch configuration, it reports this to the user. The compression ratios are displayed to the user and should be approximately 2:1. This test is useful if tape capacity or hardware compression is in doubt. This test is available for DLT, LTO, and DDS drives only.
- **DLT FUP Creation Test** takes a firmware image and writes it to a tape cartridge for use as a firmware update (FUP) tape. This test is only available for DLT drives.
- **LTO FUP Creation Test** takes a firmware image and writes it to a tape cartridge for use as a FUP tape. This test is only available for LTO drives. DLT SRAM Check Tests (short and long versions) check drive logs for existing evidence of SRAM chip corruption, and they perform a read/write test to check the drive for susceptibility to SRAM corruption problems. The test is destructive to data on the cartridge in the drive. These tests are only available for DLT drives.
- **Initialize Element Status Test** performs an *initialize element status* command on the library. This test is available for all libraries.
- **Lateral Motion Test** performs lateral motions of the robotics to enhance performance of brush tracks. This test is available for C7200 library models only.
- **MO Configuration Test + Reset** checks the configuration parameters of the MO drive and identifies configurations that differ from the factory default settings. If settings differ, the test can reset the drive to the factory defaults. This test is only available for MO drives.
- **LTO Thermal Test** monitors the temperature sensors of an LTO drive to ensure that the drive is cooling properly when installed in a server.

## Using the Support Ticket Generation Screen



L&TT can automatically generate a support ticket based on information that the program determines about the system. The information in this support ticket can then be used to diagnose storage hardware issues. L&TT can automatically e-mail the support ticket to the HP help desk to assist the help desk technician in resolving issues. To access the **Support Ticket Generation** screen, click the device you want to generate the support ticket for in the device list, and then click the **Support** button on the main toolbar.

When a support ticket is generated, L&TT collects configuration information and executes a Device Analysis test on the selected device. This information can then be viewed, saved, printed, or e-mailed. A support ticket is saved as a log file having the \*.ltx extension. By default, log files are saved to the *log* directory. The support ticket viewer is built into the L&TT application. The viewer enables viewing of detailed device information and general information on the system it is connected to. This information is displayed in a standard tree-view format. The user (or support technician) can change the detail level of the support ticket, browse through the support ticket information, and expand or collapse specific sections of the support ticket as required.

---

### Note

Regardless of the selected detail level, the log file always contains all of the information collected from the product. The support ticket viewer filters this information, depending on the detail level.

---

L&TT support tickets are saved in a proprietary format and therefore require the support ticket viewer. However, with the Windows version of L&TT a support ticket can be converted to HTML format using the currently selected detail level. The HTML version can easily be printed or viewed by other users without the support ticket viewer.

---

**Note**

After an HTML support ticket has been generated, the detail level within the HTML file cannot be changed. Make sure that the detail level is set to the desired level before generating the HTML report.

---

## Management Support

HP offers a complete range of automated storage management software products that enable you to make the most of HP storage devices and servers. These products can support the full range of storage management operations—from backup and restore to subsystem management and more. The result is increased system performance and improved productivity of system administrators and users, as well as data protection.

These tools include the capabilities for automated policy-driven operations, media and device management, and cataloging, as well as easy-to-use administration that allows the concept of “lights-out computing” to be implemented for storage management.

The management support tools discussed in this section are:

- HP Insight Manager 7
- Media Robot Utility

## HP Insight Manager 7

HP Insight Manager 7 provides web-based management for HP servers and any device compliant with HTTP, SNMP MIB-2, or DMI v2. The strength of HP Insight Manager 7 lies in its ability to provide system administrators real control through an easy-to-use industry standards-based web interface. It is feature-rich, intuitive, extensible, and designed to unlock the built-in manageability of HP hardware.

Its user interface is intuitive and packed with functionality. Accessed from anywhere on the intranet, systems administrators can manage devices, manage events, and administer HP Insight Manager 7. Preconfigured common system views and customizable queries and events provide out-of-box productivity.

### Features

- Accessibility to management information anywhere on the intranet, more securely than SNMP
- Automatic discovery and identification of industry-standard SNMP MIB-2, DMI-2, and HTTP-based web servers and compliant devices (including third-party devices)
- Comprehensive event management that enables proactive management of distributed devices by alerting in advance of system failures
- Predictive fault management and pre-failure warranty for storage, memory, and processor to protect your hardware investment
- SQL-compatible repository for discovered devices, alarms, asset, and status reporting (supports Microsoft SQL Server 6.5 and 7.0 and Microsoft Data Engine)
- Integrated Cluster Monitor that provides a one-step management control point for cluster management
- Highly scalable architecture designed for extensibility

### Backup Solutions Support

HP Insight Manager 7 enables you to monitor the Fibre Channel-to-SCSI bridge and the libraries. You can see the worldwide name, firmware revision, and status of each component, including the tape drives located in the library.

## Media Robot Utility

Media Robot Utility (MRU) software enables you to monitor and control libraries and mini-libraries conveniently from your desk. MRU provides an application programming interface (API), a GUI on OpenVMS and Tru64 UNIX, a GUI on Windows NT, and a command line interface (CLI).

### Features

You can manipulate and monitor libraries using simple commands:

- **Load** — Moves a cartridge from a slot to a drive
- **Unload** — Moves a cartridge from a drive to a slot
- **Inject** — Moves a cartridge from a port to a slot
- **Eject** — Moves a cartridge from a slot to a port
- **Show** — Displays information
- **Home** — Returns the cartridge to the location it came from
- **Find** — Finds a specific cartridge within the library
- **Help** — Displays information on the commands

### Benefits

- **Eases storage management** — Verifies library/mini-library installation and configuration.
- **Facilitates diagnostics** — Verifies that hardware and control communications paths are functioning properly. If operations such as backup fail, MRU isolates the library and verifies its operations.
- **Customizes storage management** — Enables you to write storage management solutions for your libraries.
- **Provides easy-to-use interfaces** — Choose between CLI, GUI, or API to move cartridges.

## Limitations of Fibre Channel

Certain aspects of Fibre Channel place constraints on how the solution operates.

Fibre Channel class-3, the protocol that HP uses for all Fibre Channel storage products, is considered a “lightweight” protocol. It is designed specifically for use with storage subsystems due to low latency performance. Storage subsystems have extremely low latency requirements, much lower than typical IP protocol stacks can provide. To keep the Fibre Channel class-3 protocol lightweight, error recovery mechanisms are left for upper layers to resolve.

In a primary storage environment, the blocks are relatively small and randomly accessible. Error events caused by Fibre Channel are detected and communicated to a higher level protocol, and an entire block can be retransmitted without impact to performance, provided the error events are infrequent.

In a secondary storage environment, applications do not handle error events as gracefully. Any Fibre Channel error event could manifest itself as a backup or restore job failure. If the job is a scheduled job, it will be rescheduled.

Fortunately, error events in Fibre Channel do not occur frequently in a static environment. The most common error event occurs as a result of loop disruptions. In general, these loop disruptions are a controllable occurrence.

## Common Types of Loop Disruption

The most common type of loop disruption is caused by a loop initialization procedure (LIP). An LIP occurs when a loop component is added or removed from the loop or when a loop component, such as a server, is restarted. If an LIP occurs when a data frame is on the loop, that frame will be lost.

In a primary storage environment, a higher-level protocol is responsible for re-transmitting the data and the incident is unnoticed by the users of the data. In a secondary storage environment, the application vendors have not yet dealt with this characteristic. The error is noticed by the user and shows up as a job failure.

Therefore, certain activities, such as hot-plugging loop components and resetting loop components, should be performed only when there are no backup or restore jobs executing on the loop.

You must restart each Windows NT server whenever you add backup devices to the loop, just as you must restart for the new devices to be seen by Windows NT.



## Diagnosis and Repair Using Troubleshooting Flowcharts


This section contains troubleshooting flowcharts that can be used to help diagnose problems in an Enterprise Backup Solution (EBS). The flowcharts and software utilities provided by HP can be used together to isolate problems and verify repairs.

When diagnosing problems with EBS, follow this order:

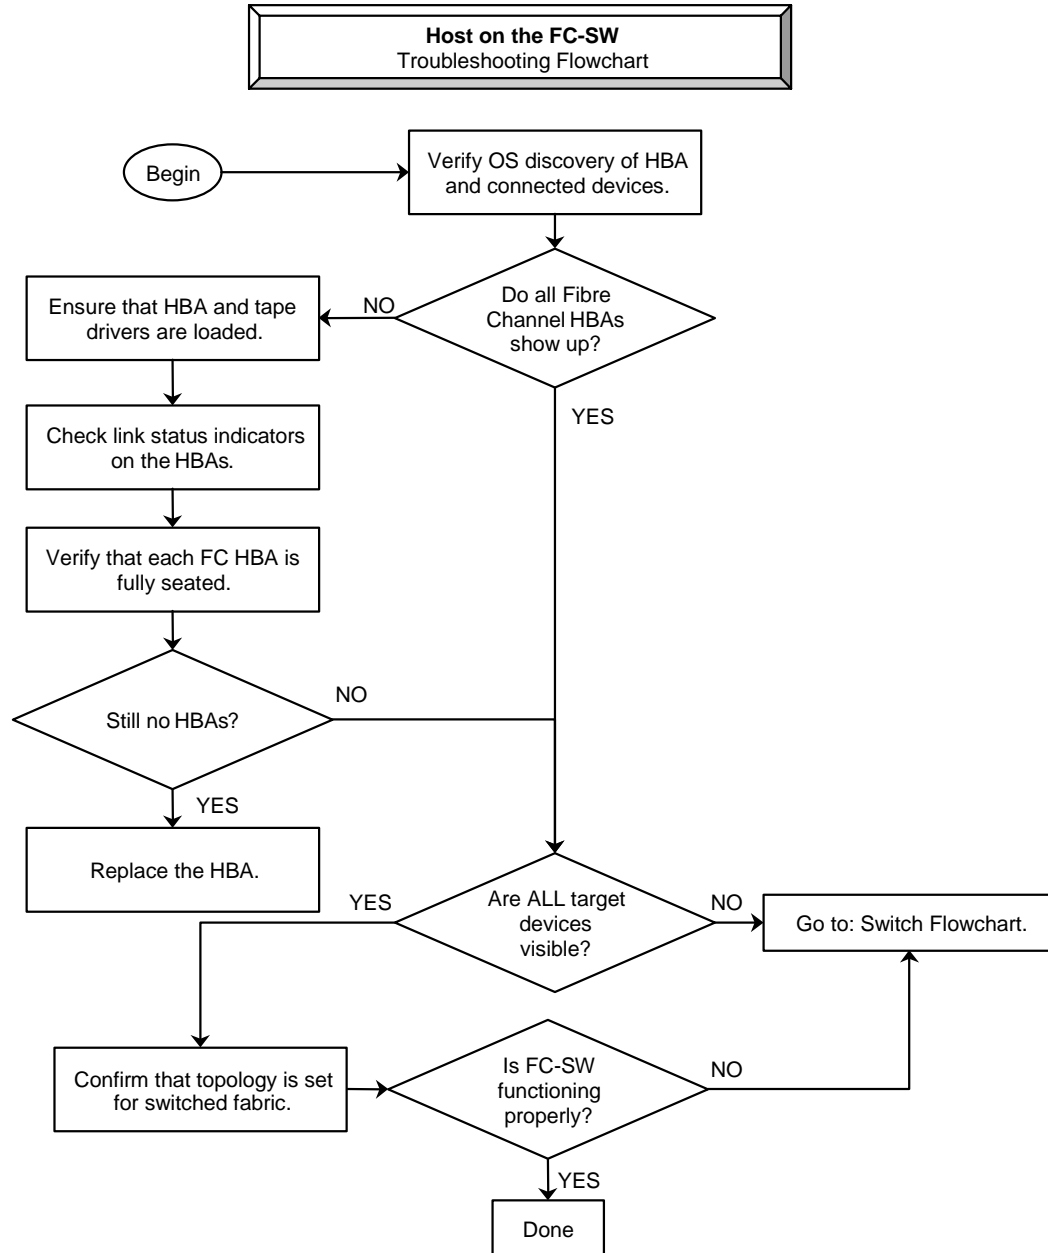
1. Perform a visual and physical inspection of the entire system.
2. Determine if there are problems with the host.
3. Determine if there are problems between the host and the hub or switch. Verify that the host is a member of the switch zone.
4. Determine if there are problems with the hub (if applicable).
5. Determine if there are problems with the switch (if applicable).
6. Determine if there are problems between the hub or switch and the Fibre Channel-to-SCSI bridge. Verify that the Fibre Channel-to-SCSI bridge is a member of the switch zone.
7. Determine if there are problems with the Fibre Channel-to-SCSI bridge.
8. Determine if there are problems with the library.

The troubleshooting flowcharts provided here help you move from one component of the EBS to another to diagnose problems.

## Visual and Physical Inspection



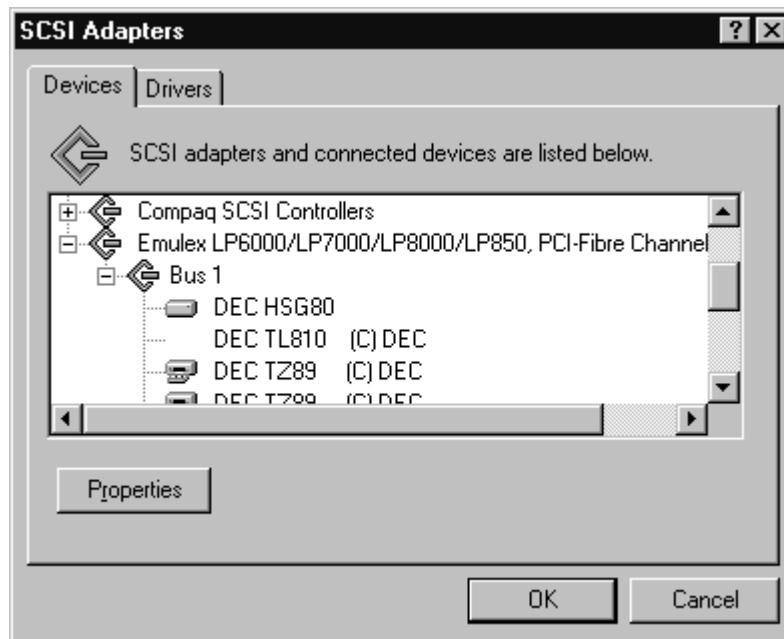
## Checking the Host on the FC-SW



## Device Confirmation

To verify operating system discovery of the HBA and connected devices, use the following methods.

For Windows NT, use the SCSI Adapters control panel.



For Novell NetWare, use the `list storage adapters` command.

For Tru64 UNIX, use the SCSI CAM utility.

```
# scu sc edt
Scanning all available buses, please be patient...
# scu sh edt

CAM Equipment Device Table (EDT) Information:

Device: RRD47    Bus: 0, Target: 5, Lun: 0, Type: Read-Only Direct Access
Device: RZ1DB-CA Bus: 1, Target: 0, Lun: 0, Type: Direct Access
Device: TL895    Bus: 4, Target: 2, Lun: 1, Type: Medium Changer
Device: TZ89     Bus: 4, Target: 2, Lun: 2, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 3, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 4, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 2, Lun: 5, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 1, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 2, Type: Sequential Access
Device: TZ89     Bus: 4, Target: 4, Lun: 4, Type: Sequential Access
Device: HSG80    Bus: 4, Target: 5, Lun: 1, Type: Direct Access

#
```

In addition, check for detected tape devices by typing *hwmgr -view dev| grep tape*

.

For Sun Solaris, check the `var/adm/messages` directory.

```
sd6 at SUNW,fas0: target 6 lun 0
sd6 is /sbus@1f,0/SUNW,fas@e,8800000/sd@6,0
fcaw0: JNI Fibre Channel Adapter model FCW
fcaw0: 64-bit (32bit mode) SBus 1: IRQ 3: FCODE Version 11.0.11 [18237e]: SCSI
ID 125: AL PA 01
fcaw0: Fibre Channel WWNN: 100000e069402eb7 WWPN: 200000e069402eb7
fcaw0: FCA SCSI/IP Driver Version 2.3, Oct 19, 1999 for Solaris 2.5,2.6
fcaw0: All Rights Reserved.
fcaw0: < Total IOPB space used: 1145024 bytes >
fcaw0: < Total DMA space used: 8458269 bytes >
NOTICE: fcaw0 NPORT Initialization Complete, SID=402eb7
fcaw0: New Fabric Parameters Received. Resetting...
NOTICE: fcaw0 NPORT Initialization Complete, SID=402eb7
fcaw0: Host: Port 220d13 (100000e069402eb7:200000e069402eb7)
fcaw0: Port 220013 (50001fe100014720:50001fe100014723) available.
fcaw0: Port 220713 (100000e00202148a:100000e00202148a) available.
fcaw0: Port 210013 (50001fe100014790:50001fe100014792) available.
fcaw0: Port 210213 (100000e002021474:100000e002021474) available.
fcaw0: Port 210a13 (10000000c920e159:10000000c920e159) available.
fcaw0: Port 210c13 (10000000c9213b0b:10000000c9213b0b) available.
fcaw0: Target 64: Port 220013 (50001fe100014720:50001fe100014723) online.
fcaw0: Target 64 Lun 1: Port 220013 (50001fe100014720:50001fe100014723) present.
sd16 at fcaw0: target 40 lun 1
sd16 is /sbus@1f,0/fcaw@1,0/sd@40,1
<DEC-HSG80-V84F cyl 26277 alt 2 hd 20 sec 169>
root on /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0:a fstype ufs
zs0 at sbus0: SBus0 slot 0xf offset 0x1100000 Onboard device sparc9 ipl 12
zs0 is /sbus@1f,0/zs@f,1100000
```

In addition, in Solaris, check for the detected tape drives by typing `ls /dev/rmt`.

## Topology Confirmation


To confirm that the topology is set for switched fabric, use the following methods.

For Windows NT/2000, ensure that the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lp6nds35\Parameters\Device\DriverParameter` registry entry includes the `Topology=1` value.

When using the KGPSA HBAs in a Tru64 UNIX or Sun Solaris server, the default topology is set for switched fabric. You can use the Emulex utility that ships with the HBA drivers to confirm the topology.

NetWare logs into the switch as an FL-port. You do not have to confirm the topology setting for NetWare.

## Checking the Switch



## Switch Management

Web Management Tools let you interact dynamically with any switch in the SAN to monitor status and performance. To manage a switch over the web, launch SWCC or log on to a switch from a host with a Java-enabled web browser by means of the Internet or an intranet.

Web Management Tools display information you can use to make overall topology decisions and provide a means to make administrative changes. Security is enforced by a user name and encrypted password.

Web Management Tools feature:

- Switch identification in the network
- Fabric topology and routing information
- Switch and port configuration information
- Real-time graphical switch and port status and statistics report
- Graphical representation of port performance (polled up to every 2.5 seconds depending on the operating system and browser used)
- Four minutes of aggregate bandwidth throughput shown graphically, scaled dynamically, and based on activity (polled up to every 2.5 seconds depending on the operating system and browser used)
- Secured management through either the GUI or a Telnet session
- Five screens for monitoring information
- Two screens for administrative interfaces
- Help functions including glossary help, online HTML help, and pop-up help for error conditions
- Out-of-band interface through a 10/100BaseT Ethernet connection
- In-band interface over a Fibre Channel link

The General Switch View page is used most often in troubleshooting a switch. Access this page for a switch by double-clicking that switch's icon on the Fabric View page.

The General Switch View page shows a graphic representation of the front panel of each connected switch. Normal long-term monitoring is conducted from this page, which provides a real-time view of each switch's overall health and status in the fabric. Switch status is updated every 1 to 2.5 seconds, depending on the operating system and web browser used.



The upper half of the screen shows port and LED indicator status, and the lower half displays general switch information. The following table summarizes the components of the General Switch View page. Right-click a component on this page to display a corresponding Help dialog box.

### General Switch View Page Components

Component	Description
Port modules	Link to the Port Detail View page
Admin button	Links to the system administration pages
Perform button	Links to the Performance View page
Telnet button	Launches a Telnet session
Thermometer	Indicates the highest temperature from the last data sample. Click the thermometer to display the temperature readings from five switch sensors.
Fans	Indicate the status of the switch's fans
WWN	Unique numeric identifier for each switch
Domain ID	A number from 0 to 239 that uniquely identifies each switch in a fabric
Role	The switch's role: <ul style="list-style-type: none"><li>■ Principal — The switch is the principal switch.</li><li>■ Subordinate — The switch is enabled, but it is not the principal switch.</li><li>■ Disabled — The switch is disabled.</li></ul>
State	The switch state: Online, Offline, Testing, or Faulty
Firmware	The firmware version
EtherIP	The switch's Ethernet IP address
Ether NM	The Ethernet subnetmask
FC IP	The Fibre Channel IP address
FC NM	The Fibre Channel subnetmask
Gateway	The gateway address

Verify that the host, storage system, and tape controller are in the same zone.



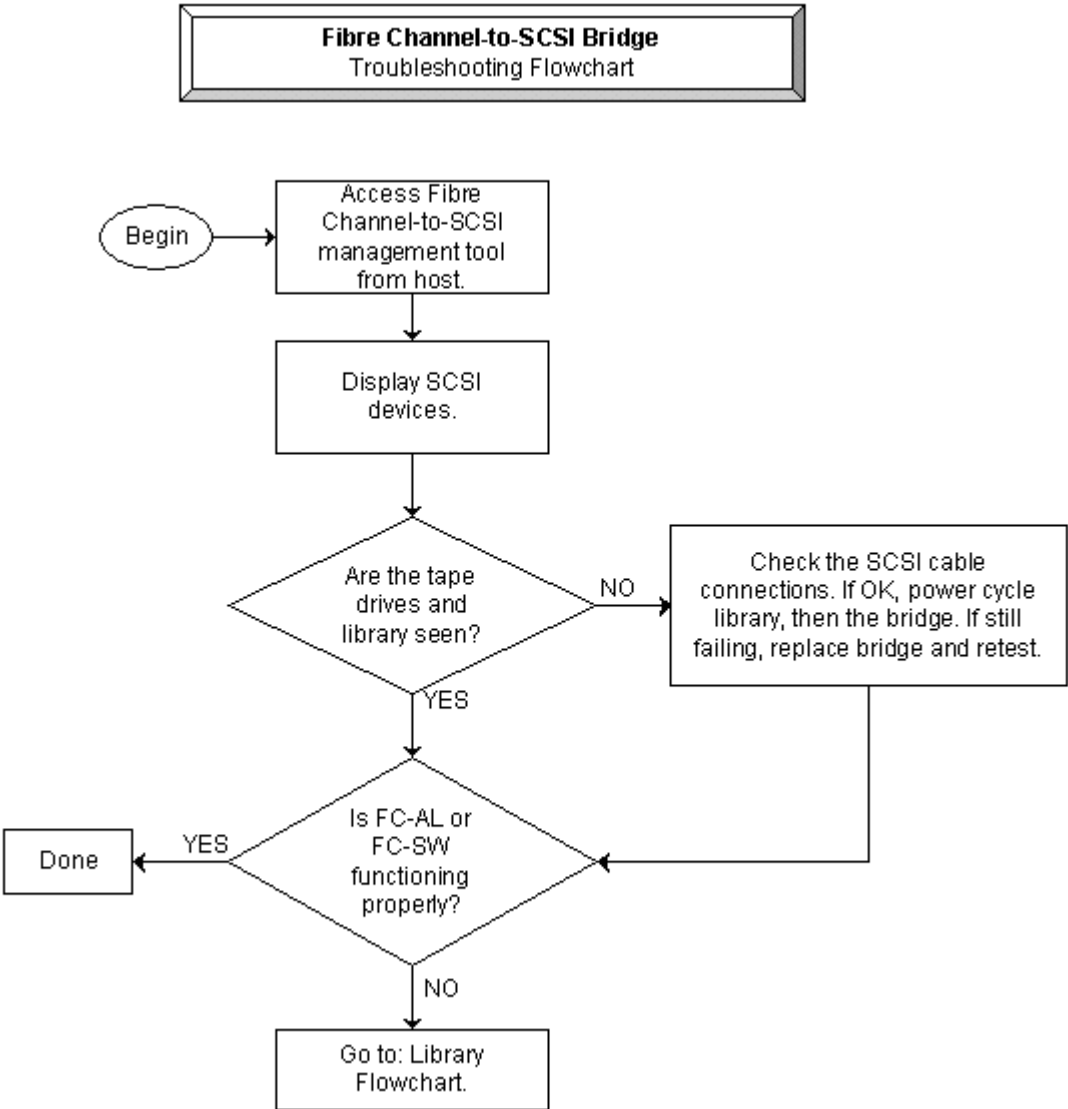
---

**Note**

For more information on troubleshooting switches, see the appropriate installation and hardware guide.

---

# Checking the Fibre Channel-to-SCSI Bridge



## Troubleshooting the Modular Data Router

Perform the following steps to troubleshoot the Modular Data Router (MDR):

1. Verify AC power.
2. Inspect cables.
3. Perform POST.

### Verify AC Power

The MDR has an LED on the rear panel that illuminates when it is connected to an AC power source and the power switch has been turned on. If this LED does not illuminate, verify the following:

- The MDR has been plugged into an active AC power source.
- All peripheral devices are connected to AC power.

### Inspect Cables

There are two LEDs on the MDR Fibre Channel module. Normally, these indicate the condition of transmitted and received signals.

When a cable attached to an operating loop or fabric is plugged into the GBIC in the module, both LEDs should illuminate. If they do not, the likely causes are a bad GBIC or Fibre Channel cable.

If a pattern of alternating periods of rapid flashes followed by a short period of non-illumination is displayed, there may be an internal problem in the Fibre Channel module.

Inspect the SCSI cabling between the MDR and all attached SCSI devices. Check for the following:

- Loose, dirty, broken, or bent cabling and connectors
- Cables resting on the SCSI module's plastic handle



---

#### Note

The VHDCI SCSI connectors used on the MDR cannot support the weight of the attached SCSI cables. Always make sure the plastic handles are used to support the cable.

---

The MDR SCSI modules are marked with icons indicating whether they contain Low-Voltage Differential/Single-Ended (LVD/SE) or High-Voltage Differential (HVD) transceivers. Verify that the target SCSI devices have the same type of transceivers. The MDR will not be able to communicate with the SCSI devices if the transceiver types do not match.

Each SCSI module has one LED indicator for each SCSI bus. The LED is an indicator of activity on the SCSI bus. Depending on the level of SCSI activity, the LED could be off (indicating no SCSI activity) or it may appear to be solidly illuminated during high rates of SCSI activity.

If the LED appears on all the time but there is no apparent I/O activity happening:

- The SCSI bus may be hung.
- This could be an indication of a faulty SCSI module in the MDR.
- It could be an indication of a faulty component outside the MDR such as a SCSI cable or tape drive.

### **Power-On Self-Test (POST)**

Immediately after the MDR is turned on, it performs a power-on self-test to verify its internal hardware is working properly.

- The entire process should last approximately 90 seconds.
- When POST starts, all LEDs on the Management Module will illuminate and then go through a series of indicator patterns as various portions of the internal hardware are tested.
- After POST is complete:
  - The test is completed properly if the Management Module LEDs display a *single* LED moving back and forth.
  - If the LEDs illumination pattern *does not* change for more than 90 seconds, there could be a hardware failure in the MDR.

If this happens:

- ◆ Power off the MDR.
- ◆ Reseat all of the option modules, making certain to tighten the modules in with their thumbscrews.

If POST stops with a static LED pattern again, connect a terminal or terminal emulator to the MDR's DB-9 serial port to monitor POST progress and error messages.

The DB-9 serial port provides an advanced configuration and management interface. This serial port reports POST progress and system error messages, as well as providing command line interfaces for boot management and application management.

To monitor POST progress:

1. Connect one of your servers to the MDR by attaching the provided DB9 serial NULL MODEM cable to the serial port on the MDR.
4. Select *Start* → *Programs* → *Accessories* → *Hyperterminal* → *Hyperterminal* and enter a name for your hyperterminal in the New Connection dialog box.
5. Select an *icon* to represent the tape controller and click *OK*.
6. The Phone Number dialog box displays. From the *Connect using:* pull-down menu, select the communications port assigned to your serial port; click *OK*. (You might have to enter an area code.)
7. Enter the following information in the Properties dialog box; select *OK*.
 

a. Bits per second (baud rate)	<b>9600</b>
b. Data bits	<b>8</b>
c. Parity	<b>None</b>
d. Stop bits	<b>1</b>
e. Flow control	<b>None</b>
8. Select *File* → *Properties* → *Settings* and from the Emulation pull-down menu, select *VT100* → *OK*. The Power-On menu should display.
9. Press *Enter* if the Power-On menu does not display on its own.

## Troubleshooting the Network Storage Router

### LED Indicators

The LED indicators on the router are useful for diagnosing various problems:

- The Fibre Channel LEDs indicate Fibre Channel activity and link status. If one of these indicators does not light or stays continually lit without any corresponding SCSI bus activity, it may indicate a problem with the Fibre Channel link. Verify the Fibre Channel configuration.
- The SCSI Bus LEDs indicate SCSI activity. These indicators are lit during power up, configuration, and when the unit is transferring data. If the SCSI indicator stays continually lit without any corresponding Fibre Channel LED activity, it may indicate a problem with the SCSI bus configuration. Verify the SCSI bus configuration.
- The Ethernet LEDs indicate activity and link status. If one of these indicators does not light or stays continuously lit, it may indicate a problem with the network connection. Verify the network connection. The port must be connected to a 10/100BaseT Ethernet network to function properly.

### Verifying SCSI Bus Configuration

Items to check include:

- **Termination**—Problems with termination can cause intermittent or hard failures. A SCSI bus must be terminated on both ends. Termination problems are common when both narrow and wide devices are on the same bus.
- **Bus Type**—On an LVD SCSI module, SE and LVD devices can be connected to the same bus. However, if one SE device is detected during power up, communication to all devices will convert to SE mode.
- **Device ID**—Each device on a SCSI bus must have a unique ID. Verify configured IDs are not in use by other devices on the same SCSI bus.
- **Cabling**—Check SCSI cables to verify they are functional. SCSI rules for total length, distance between devices, and stub length must be followed. Connections should also be checked and reseated if necessary.
- **SCSI Devices**—Verify that the SCSI devices on a particular SCSI bus can be seen in the Configuration Menu of the router. If the router cannot see the devices, verify SCSI configuration, cabling, and termination.

### Verifying Fibre Channel Configuration

If SCSI devices are recognized on the SCSI buses, but do not appear to the Fibre Channel host, it may be that the Fibre Channel link is not properly established. Most hubs and switches have link indicators showing link status. When the router is connected and powered on, this link indicator should show a good link. If it does not, check the cabling or connections.

One method of verifying link integrity when connected to a functional host, involves disconnecting and then reconnecting the Fibre Channel cable. This

procedure should cause momentary activity of this indicator as the link reinitializes.

Additionally, verify that the media type of the router and the attached hub, HBA, or switch is of corresponding types.

---

**Note**

By default, the Fibre Channel port speed is set to 1 Gb/s. Changes to the Fibre Channel port speed must be manually set, such as for 2 Gb/s. If set incorrectly and plugged into a Loop or Fabric, the unit may receive Framing errors due to the incorrect Fibre Channel link speed.

---

## **Verifying the Router Configuration**

If you are in doubt about the configuration or about the location of the error, restore the router to the factory default configuration and proceed to configure the unit one step at a time, verifying the functionality of the configuration after each change is made.


## **Verifying Devices**

It may be useful to connect the SCSI target devices directly to a SCSI interface (for example, a host SCSI bus) to verify that the devices are functional.

## **Verifying the Host Configuration**

In some cases, it may be that the Fibre Channel HBA or host device driver may not be working properly. Check the configuration of these elements. It may be useful to check the release notes for the device driver to see if there are any specific issues or required configuration. It may also be useful to ensure that the current version of the HBA driver is being used.

## Checking the Library





## Typical Causes of Library Failure

When diagnosing hardware failures, observe the physical characteristics of the library with power off and power on. You can often locate failures during a visual inspection.

If, after all the troubleshooting, you find that the library is the source of the hardware failure, consider these common causes:

- Mechanical failure of tape drive
- Tape drive booted with tape inside
- One or more cables disconnected
- Cartridge lock assembly in wrong position
- Object blocking sensor path
- SCSI IDs not configured properly

## SCSI IDs

The order of the SCSI ID assignments for the robotics controller and drives is important when using the FCTC-II. The robotics controller must be set to SCSI ID 1. Drives must be set to SCSI IDs 2 and 3 at the hardware level. For example, the following table lists the IDs for a TL895 with seven drives.

Device	FCTC #1		FCTC #2	
	Bus 0	Bus 1	Bus 0	Bus 1
Robot	1			
Drive 0	2			
Drive 1	3			
Drive 2		2		
Drive 3		3		
Drive 4			2	
Drive 5			3	
Drive 6				2

SCSI IDs for TL895 with Seven Drives

## Library Troubleshooting Strategy

Conduct repeated tests and visual inspections of the library, and check cables and other connections to ensure that they are joined correctly. Do not simply rely on the diagnostics tools.

When you are not sure about how to proceed, remember to **run, watch, and record**.

1. **Run** one or more tests.
2. **Watch** the mechanisms.
3. **Record** the error messages.

## POST

POST is often the key to troubleshooting. During POST, watch the mechanisms of the tape library and note where they stop.

Observe what the unit does during the test. If an error is reported, what action was the unit attempting when the failure occurred? Visually inspect the mechanisms. Are all the cables connected? Does anything seem abnormal?



---

**Note**

Watch the unit perform normally as much as possible so that you can detect problems visually.

---

## Troubleshooting the SSL2020TL and TL891 Libraries

The following problems can cause the SSL2020TL and TL891 Libraries to malfunction:

- Platform problems
- General drive errors

### Platform Problems

Incorrect installation and configuration cause platform problems. In this case, the Library appears to be operating normally, but no data can be interchanged. You might or might not get an error code on the Control Panel. To identify an error caused by this type of problem, check your installation and configuration setup.

### General Drive Errors

General drive errors usually result from a miscommunication between a module's processor and a drive processor or a mechanical malfunction within the Library.

Both platform problems and general drive errors display an error message and a Fault System Code (FSC) on the Control Panel. Use the FSC to report errors to your Technical Support representative, or in some cases, to determine a recovery procedure.

A simple error recovery procedure displays at the bottom line of the Control Panel display. You can clear some errors by pressing the *Enter* button on the Control Panel or by cycling the power to the Library.

## Troubleshooting the MSL5000 Library Family

There are two main types of problems that can cause the MSL5026 Library to malfunction or fail to perform correctly: platform problems and general drive errors. Some errors cause Fault Symptom Codes (FSCs) to be displayed on the GUI touch screen, along with a description of the fault. Error Recovery Procedures (ERP) are available for each FSC. FSCs are described in the MSL5000 User Guide.

### Platform Problems

These errors arise out of incorrect installation and configuration. The most common characteristic of this type of error is that the drive seems to operate normally, except that no data can be interchanged. You might not get an error code on the GUI touchscreen. To identify an error as this type of problem, check your installation and configuration setup.

### General Drive Errors

These errors usually result from a miscommunication between the drive processors or involve a mechanical malfunction. In most cases, both platform and general drive errors report an error message and an FSC on the GUI touch screen. The FSC can also be used to determine a recovery procedure.

### Error Recovery Procedures

ERP Code	Corrective Action
C	Cycle power the library using the Power option on the GUI touch screen. Wait 30 seconds to power up again.
D	Turn off the power to the library and inspect the connectors and cables.
F	Invalid operation. Set parameters correctly and try again.
G	Call Technical Support.

### Fault Symptom Codes

Fault symptom codes, along with the suggested error recovery procedures, are listed in the MSL5000 User Guide.

## Troubleshooting the TL895 Library and ESL9000 Library Family

The following sections outline many of the messages and problems you might encounter during the initial configuration and operation of the ESL9000 Library:

- Messages
  - Status
  - Sense Data Value
- Problems
  - Start-up
  - Control Panel
  - Robotics
  - Operating
  - Tape drive

### Messages

Some events produce status messages on the control panel; others produce error messages called *sense data values*. Sense data value messages display on the Control Panel or host monitor, and include a number and description of the error that occurred. Information on messages can be found in the user guide.

### Problems

In addition to status messages, other problems might occur during library operation. The following tables list the most common problems and the steps you can take to resolve them. For problems that cannot be resolved using the information in this module, notify your system administrator or field service engineer.

## Start-Up Problems

Problem	Corrective Action
The library does not power on.	Ensure that the library power switch is set to the I (on) position and the power cord is connected to a grounded electrical outlet.
The library or tape drives do not respond on the SCSI bus.	Verify that each SCSI device on the same SCSI bus has a unique address and that the last device is properly terminated.
The library reports "not ready" during initialization.	Determine the failure type by checking any previous error codes returned to the host computer. Correct the cause of the error.
One or more tape drives fail to spin up during start-up.	Check all SCSI cabling and termination on the back panel of the library. If necessary, contact your field service representative about replacing the drives.
There is a cartridge in the gripper, preventing system calibration and inventory.	Manually remove the cartridge from the gripper. Then restart the library.
The library starts up in offline mode.	Press the Standby button. Verify that the library switches to online. You can use the Control Panel to select either online or offline mode at power-up.

## Control Panel Problems

Problem	Corrective Action
The Control Panel is blank or does not respond to touch.	Contact an authorized field service engineer.
An error or information message displays.	Write down the details of the message, including the SK, ASC, and ASCQ numeric values. Then press <i>OK</i> to clear the message. For instructions on resolving a possible error, refer to the user guide.

## Robotics Problems

Problem	Corrective Action
The robot does not move at power-up.	Verify that all internal packing materials (foam pads and tie wraps) were removed during the installation procedure. Check the Stop and Standby buttons; ensure the library is online and the Stop button is released. Ensure that all library doors are closed.
The gripper partially grips a tape cartridge.	Issue a <i>Move Cartridges</i> command to move the cartridge from the gripper to an empty storage element.
The bar code reader on the gripper fails.	Verify that nothing obstructs the reader. Then restart the library. If the problem continues, contact an authorized field service engineer.
The robot times-out or fails during an operation.	Confirm that the tape cartridge involved in the operation is properly positioned in the bin or drive and ready to be picked. Verify that the robot is not obstructed in any way. Retry the operation. If it still fails, contact an authorized field service engineer.
The robot drops a cartridge.	Open the storage array door safely. (To do so, press the <i>Standby</i> button first to place the library offline, and then unlock and open the door.) Retrieve the cartridge, orient it properly, and place the cartridge in an empty storage bin. (Do not try to place the cartridge in the gripper.) Afterwards, close and lock the storage array door and press the <i>Standby</i> button again to return the library to online mode.
A cartridge is in the gripper at start-up, when a move command is requested, or after a place command is executed.	Press the <i>Standby</i> button to place the library offline; then open the storage array door. Manually remove the cartridge from the gripper and place it in an empty bin. Then close the door and press the <i>Standby</i> button to return the library to online mode.
The gripper does not have a cartridge after completing a pick command.	Verify that a cartridge is in the source location. Retry the command. If the pick operation fails again, contact an authorized field service engineer.

## Operating Problems

Problem	Corrective Action
The host computer cannot communicate with the library.	This might be a SCSI bus time-out or a premature disconnect problem. Check cable connections, cable length, SCSI addresses, and termination. Restart the host and the library. If the host and library still are not communicating, contact an authorized field service engineer.
A cable or terminator is disconnected from the library back panel.	Reconnect the cable or terminator.
A tape cartridge (medium) is reported not present.	The gripper cannot sense a tape cartridge in a particular storage element even though the inventory reports that it is present. Verify that the designated cartridge is present. If it is, ensure that it is properly seated. (For a tape drive, ensure that the cartridge is completely unloaded.) Then retry the command. If the error persists, contact an authorized field service engineer.
A move command failed.	Check the source and destination elements. The source element should hold the cartridge to be moved; the destination element should be empty. Ensure that the gripper is empty and all actuators are free of obstruction. Also, verify that the library is online and the <i>Stop</i> button is released. Retry the command.
A flash memory error is reported.	Contact an authorized field service engineer.
A <i>maximum temperature exceeded</i> warning displays.	Power off the library and allow it to cool down. Lower the room temperature, if possible, and increase ventilation around the library. (If the operating temperature is too high, the library automatically shuts down until the temperature drops.)



## Tape Drive Problems

Problem	Corrective Action
The library is unable to communicate with a drive, indicated by a <i>Drive Communication Time-out</i> error.	Contact an authorized field service engineer.
Failure to mount or read/write with new or used cartridge.	The cartridge is bad. Retry the same operation with a different cartridge. OR the drive heads are dirty. Use a Cleaning Tape III head-cleaning cartridge.
Application software-generated initialize command fails with parity error.	The tape calibration failed. Use a Cleaning Tape III head-cleaning cartridge or try a different data cartridge.
Green <i>Operate Handle</i> indicator stays lit and tape does not move; yellow <i>Tape in Use</i> indicator is lit, but does not blink.	There was a cartridge loading error. Dismount the cartridge and inspect it for a mispositioned leader. If the drive leader is not in the correct location, call your HP authorized reseller or service provider. OR inspect the drive for a damaged, misplaced, or unhooked leader. If the drive leader is not in the correct location, call your HP authorized reseller or service provider.
All right-side or left-side indicators are blinking.	The drive failed the self-test or detected an error during operation. Ensure that the drive handle is in the down position. Try to clear the error by pressing the <i>Unload</i> button, or power off and then back on. If the error does not clear (the tape does not rewind and unload, and the lights blink), you have a hardware failure; call your HP authorized service provider.
The tape drive does not eject a cartridge.	Reset the library and retry the unload command. If the tape still does not unload, press the <i>Standby</i> button to place the library offline, open the front door, and manually unload and eject the cartridge.
A drive handle error occurs.	Contact an authorized field service engineer.
Unable to back up or copy data to tape.	The cartridge is write-protected. Set the tab on the cartridge to write-enabled or, if software is protected, modify that protection. OR no tape is in the drive. Insert a tape.
<i>Write Protected</i> indicator flashes orange.	Perform the head cleaning procedure. If the error repeats, try another tape.
<i>Tape in Use</i> and <i>Write Protected</i> LEDs flash rapidly in unison.	Eject the tape. Perform the head cleaning procedure. If the error repeats, replace the tape. Power off and power on the drive. If the error repeats, replace the drive.
After applying power, nothing happens. All indicators are off, except green <i>Operate Handle</i> indicator.	No tape is loaded. Load a tape.
Drive is not available to the system.	Check the AC power source to be sure power is getting to the drive. OR check the SCSI ID switch settings. OR ensure that the power cable is plugged in and that the SCSI cable is seated properly.

## Learning Check

1. The Media Robot Utility (MRU) can be used to perform which of the following?
  - a. Move a cartridge in a library from a drive to a slot in the library.
  - b. Load firmware into a tape drive.
  - c. Notify the administrator of faults with the library through email.
  - d. View errors on the Network Storage Router.
2. Error events caused by Fibre Channel are detected and communicated to a higher-level protocol without impact to performance in the secondary storage environment.  
☐ True  
☐ False
3. A LIP occurs when a:
  - a. Data frame is sent.
  - b. Job is restored.
  - c. Server is restarted.
  - d. Data frame is lost.
4. If you suspect that the tape drive is exhibiting faulty behavior, but it passes POST without errors, what should you try next to diagnose the problem?  
.....  
.....
5. A typical cause of a tape library failure is: (Select all that apply.)
  - a. One or more cables disconnected
  - b. An object blocking the sensor path
  - c. A GBIC removed from the library
  - d. A bad host adapter



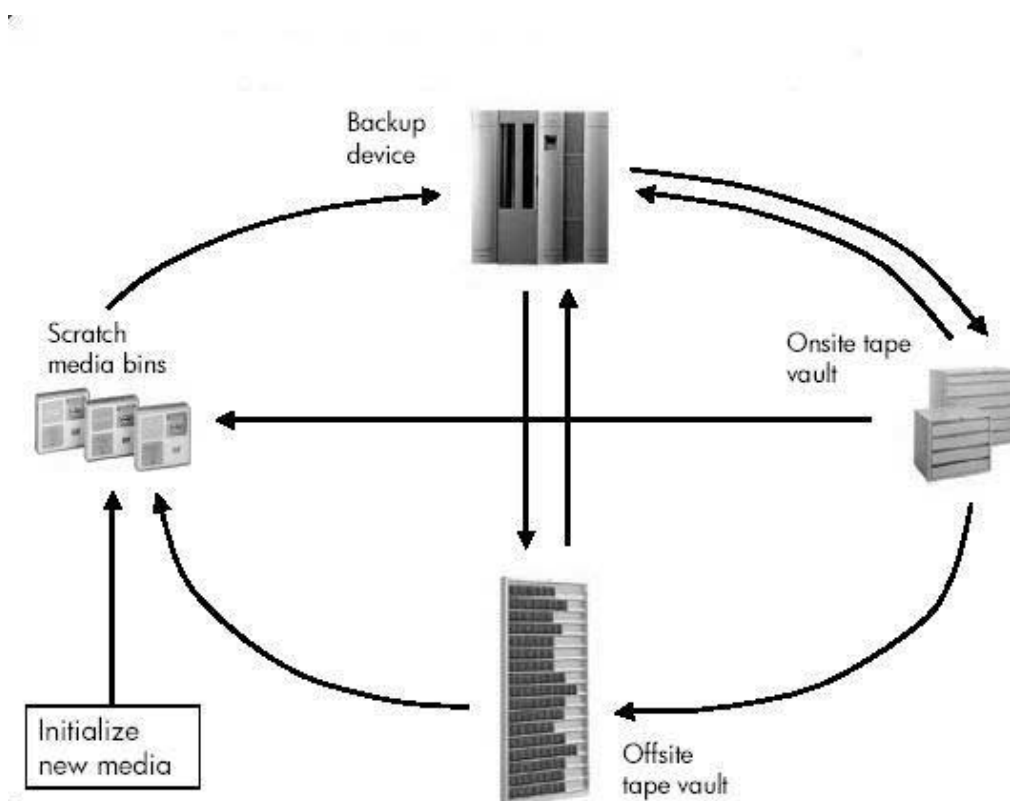
### Objectives

After completing this module, you should be able to:

- Explain the concepts and architecture of HP OpenView Storage Media Operations.
- Describe the required installation components for the Storage Media Operations server, database, clients, Data Protection Integration Agent, and Storage Media Operations licensing.
- Describe the objects and organizational units of Storage Media Operations
- Explain the daily operations for Storage Media Operations.

## Overview

### Typical Physical Media Lifecycle



Throughout their lives, all physical media pass through several levels of storage and archiving. The levels and the order and timing the tapes go through them depend on the data they contain as well as disaster protection policies. These are different from customer to customer. However, the stages and processes can be simplified to form a typical, physical media lifecycle.

At first, media usually get initialized in order to make them available to backup applications. Initialized media may be contained in scratch bins. For backup, they are loaded into backup devices, for example tape libraries. Media are likely to stay in the device after the backup in order to append additional backups or to perform restores. After a while however, they'll be unloaded from the backup device to make room for new media or because the version of the data they contain became unlikely to be needed for restores. In some case, media get unloaded immediately after the backup completes. Next, media will likely be stored in vaults. These vaults can either be at the same site as the backup devices (onsite) or some place else (offsite). An onsite vault is physically close to the backup device for example in the same data center or building. Onsite vaults allow for a quick return of media to the devices for but they don't offer media protection against major disasters, which endanger the life data on disks as well as their backups on media. This level of protection can be achieved by shipping media to separate sites, which are normally many miles away from the origin of the data. The disadvantage of offsite vaulting is the considerably large amount of time, which is required to return the tape onsite for restore or copy operations. If the backup data is subject to a

retention period, all media will expire at a certain point in time. Once expired, a medium may be re-entered into the cycle by putting them back into the scratch bin where they become available for new backup data.

## **Impacts and Challenges**

All media need to be managed in some way throughout their entire life. This may raise a quite a challenge if a company handles a high volume of for example several hundred tapes a day. How well these are managed has a significant effect on the availability of data for restore operations. The media management application needs to be able to quickly pinpoint the location of all media required and initiate their transport to the restore devices if necessary. Time to recovery is critical in all businesses. Downtime costs businesses up to several millions of dollars an hour. It's also necessary to monitor the quality of the media to ensure their readability when needed. This should be based on measures such as the number of overwrites as well as the lifetime of a medium.

The way media are managed also has a significant impact on cost. Media handling and tracking in general are very labor intensive. It may be optimized by a media management application if it arranges media mount and/or dismounts in data centers in order of the shortest way. Cost of ownership is also influenced by how effectively media are utilized. Ideally, all media should be utilized up to their maximum capacity.

## **Storage Media Operations Features and Benefits**

The impacts and challenges of the media lifecycle make media management a very important service. Establishing and executing the processes that move a medium thru the stages of its lifecycle is a complex task. It gets particularly complicated when a large number of media needs to be managed. HP OpenView Storage Media Operations automates the processes that make up the media life cycle. It tracks both storage and movement of media thus making it possible to quickly determine the physical location of media, which need to be returned to backup devices for restore or that are to be shipped to a different vault. Expired media will be returned to the scratch bins. Storage Media Operations is also capable of monitoring tape consumption of individual, reoccurring backups. Based on this information, it'll predict the amount of media needed in a backup device per day and order these media to be premounted. Also, it will provide information as to whether the tapes are utilized well or whether there are too many unnecessary blank tapes in the datacenters.

Storage Media Operations also supports barcode printers as well as scanners. The capability of printing barcode together with the tight monitoring of tape usage it can significantly lower TCO. Barcode scanners allow for accurate check-in or check-out of media into or out of backup devices and vaults.

## Components of Storage Media Operations

Storage Media Operations consists of three main components:

- Storage Media Operations Server
- Storage Media Operations Client
- Integration Agents


The Storage Media Operation Server is the central component of the Storage Media Operations architecture. It hosts the Storage Media Operations database, which houses all Storage Media Operations objects and procedures. Of course, a set of binaries including tools for database backup and maintenance is also installed to the server. It runs a Storage Media Operations Database service for the database and its procedures as well as Data Management Communication service, which allow it to communicate with Integration Agents of backup applications. A built-in web server provides connectivity for web client access. The Media Operations server provides a GUI that allows configuring core server parameters.

The Storage Media Operation Client is necessary to administrate Storage Media Operations. The GUI of the server doesn't provide this functionality. Administration is mainly about configuration of database objects and procedures. The Storage Media Operations Client is automatically installed to the server as well. It also supports to attach a barcode scanner to the client system for inputting barcodes of media while they are being moved from location to location.

The Integration Agents (IA) allow Storage Media Operations to automatically retrieve information on media, backups, systems, etc. from backup applications. The first IA available is the Data Protector Integration agent, which integrates with Data Protector 5.0 as well as Omniback II 3.5 and 4.1.



## Storage Media Operations Architecture



The Storage Media Operations Server is the core component of this client/server architecture. It houses the database with its objects and attributes. It also contains the business logic to process the administrator defined vaulting and scratch policies. The Storage Media Operations Server also performs Scheduling and SLA monitoring functions.

The Storage Media Operations Client is a Windows based client GUI, which is also referred to a Media Operations Manager. It runs locally on the server but can also connect remotely via TCP/IP.

The web client connects remotely from web browsers such as Microsoft Internet Explorer or Netscape to the web server, which is provided with every Storage Media Operations Server. The functionality of the web-based GUI is limited to media listings and daily operations such as vaulting. The web client doesn't support any barcode scanners for input of media barcodes.

The Data Protection Integration Agent provides integration into hp OpenView storage data protector and selected versions of its predecessor Omniback II. The IA uses the OpenView Foundation communication libraries, called BBC (Black Box Communications). The communication between the Storage Media Operations Server and the IA is entirely HTTPS based. The server will send all requests as HTTPS URL based requests, which act as XML command line calls.

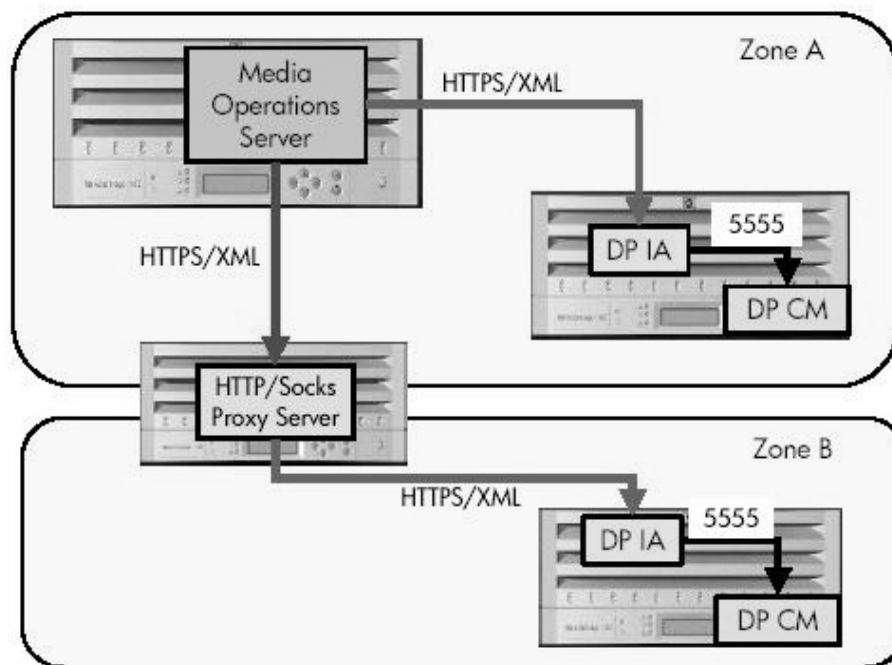
The Integration Agent will also return all resulting data as XML data via HTTPS. The agent does not require any special communication path to be set up as runs over standard HTTPS, which normally passes through firewalls. The HTTPS

communication of the Storage Media Operations Server with the Data Protector Integration Agent is secured using SSL. The IA holds a predefined private key and the MO server the matching public key.

The Data Protection Integration Agent will communicate with the Data Protector cell server using the internal Data Protector socket API protocol over port 5555 (default). If the port is different from the default, the Media Operations Server will add the custom port as a parameter to the XML call it submits to the IA. From a DP point-of-view, the IA acts just like another Data Protector command line utility except that it uses an HTTPS/XML interface. Besides command line calls, the IA also uses the DP port 5555 or the customized equivalent to retrieve configuration files such as backup specifications.

The Storage Media Operations Server can concurrently use multiple DP IAs to access one Data Protector or Omniback II cell manager. The load will be balanced among the DP IAs.

## Firewall Support



Storage Media Operations can work across firewalls. The Data Protector Integration Agent can be in a different zone than the Storage Media Operations Server. The HTTPS/XML requests of the server are sent via a standard proxy protocol (e.g. SOCKS). This way the Storage Media Operations Server can communicate with the DP IA across firewalls. Communication of the Data Protector Integration Agent with Data Protector or Omniback II cell manager is not possible across firewalls. The agent and the cell manager need to be in the same firewall zone.

## Installation Requirements

### Storage Media Operations Server Installation

The Media Operations Server can only be installed to systems running Microsoft Windows 2000 or Windows NT4 SP6. The installation requires an initial 100 MB of disk capacity including 25 MB for temporary file space. Of course, more disk space will be needed as the database grows. The biggest factor in database growth is the amount of media managed and media operations. Adding 10,000 media will cause the database to grow by approximately 100 MB.

The installation requires Microsoft® Installer 2.0 to be present on the system. If the system runs an older version of Microsoft Installer, the Media Operations Installation routine will automatically update it to version 2.0. In case of the automatic update, the system will require a reboot upon completion of the installation.

The Media Operations Server runs on a single Pentium III 500 MHz processor system or equivalent. We recommend upgrading the system to a dual processor system or to higher processor speed when reaching 25,000 media. This is especially true when the server shares a system with a Data Protector or Omniback II cell manager. The scheduled operations of the server require a lot of CPU time and might cause other applications, e.g. the Storage Media Operations Manager or Data Protector Manager, to hang for extended periods during the main processing loads. Therefore, we recommend keeping the Media Operations Server on a separate system. Due to a limitation of the database product from 4D, the MO Server is only able to utilize one CPU.

The server requires a printer to be set up on the system to perform the installation. The installation procedure will not only install the Storage Media Operations Server but also the client software as well as documentation and database maintenance and recovery tools.

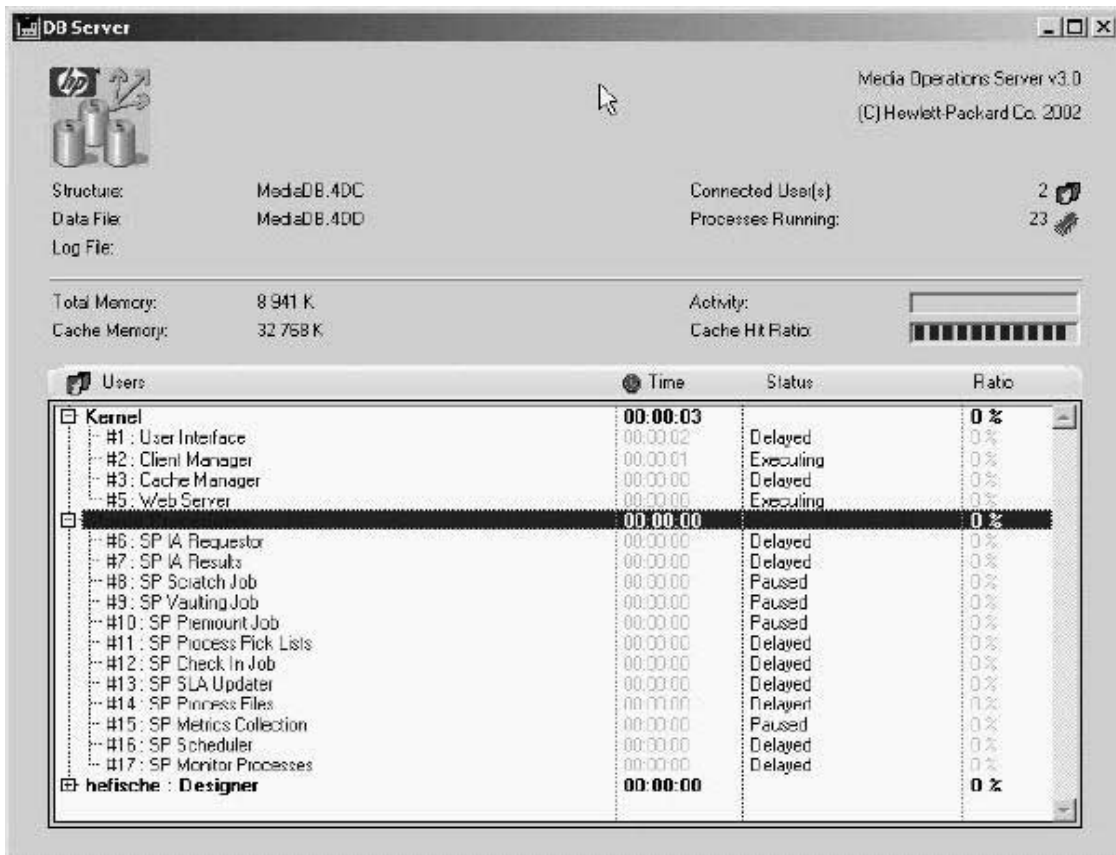
Run the setup.exe program from the installation media. During the installation procedure, you will be prompted to enter

- Path to the Media Operations binaries
- Path to the Data Management Communications module
- Path to the Media Operation Database
- Name and password of the service-level administrator user

Once the software has been loaded to the system, the installation routine registers and starts the HP Media Operations Service as well as the HP Data Management Comms Service.

The user name and password you entered will allow you to log on to Storage Media Operations.

## Storage Media Operations Server GUI



The GUI of the Storage Media Operations Server enables administrators to view and change parameters of the server. It does not provide any functionality to manage or track media.

The main window of the server GUI provides an overview of all kernel processes, stored procedures (SP) and users who are currently connected to the server. It also shows version and statistical information, for example the number of users currently connected, the cache hit ratio and the CPU load. Note that the server GUI itself counts for 1 connected user.

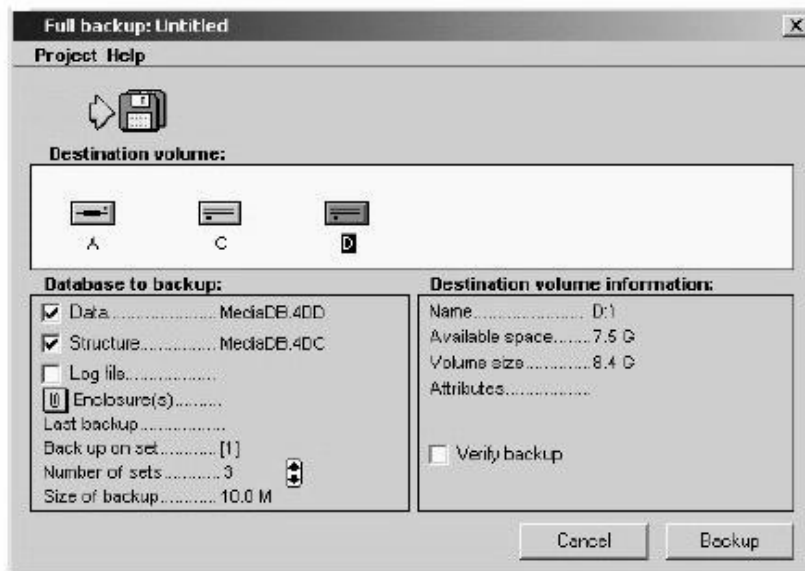

Besides the main window, there's a window for scheduling backups of the Storage Media Operations database. By default, no backups are configured.

## Storage Media Operations Database

The Media Operation database is a relational database. It is kept in the %MO\_DB% directory. It consists of the data file MediaOperations.4dd and the structure file MediaOperations.4dc. The data file holds the database objects and the structure file contains a structural layout of the entire database. The data files have a limitation of 2 GB each. They may be segmented to gain a maximum database size of 128 GB. The database can be expected to grow by about 100 MB with every 10,000 tapes added to Media Operations. MediaDB.4DR is the resource file of the data file and MediaDB.RSR is the resource file of the structure file.

After the initialization of the database, which is part of the Media Operations installation routine, the database will not log transactions in a transaction log file. This log file can be created from the Media Operations Server GUI. It will be necessary to recover the database from a potential corruption or disaster. Because of this, the creation of the transaction log requires a full backup first.

## Protecting the Storage Media Operations Database



The data kept in the Storage Media Operations Database is essential for the application to work. It contains information about the location of media. This information is critical in case of recoveries. Therefore, the Storage Media Operations Database needs to be protected from potential failures due to corruptions, hardware failures or human errors. The two methods of protection available are backup and mirroring.

## **Backing up the Storage Media Operations Database**

The Storage Media Operations Server provides a volume-based backup for its data file, structure file and log file. This backup may be run interactively or scheduled. It generates separate archives for the data, structure and log file. These files follow a default naming convention:

- MediaDBxxx.4BK – backup of the data file where xxx is a three-digit decimal number
- MediaDBxxx.4BR – backup of the structure file where xxx is a three-digit decimal number
- MediaDBxxx.4BL – backup of the log file where xxx is a three-digit decimal number

## **Mirroring the Storage Media Operations Database**

The Storage Media Operations Server allows scheduled or interactive updates to mirroring of the database. In order to mirror the database, the log file is required. A full backup is required for log file creation.

## Storage Media Operations Clients

There are two client interfaces available for Storage Media Operations: the Storage Media Operation Manager, which is a Windows based GUI that can be installed to Microsoft Windows 98 SE, NT4 SP6, 2000 or XP, and the web client, which just requires Microsoft Internet Explorer or Netscape.

The Storage Media Operations Manager provides access to all functionality of Media Operations given that the user has appropriate permissions. The functionality of the web GUI is limited to daily operations, reporting and media listings. All other Storage Media Operations objects cannot be viewed or configured.



### **Important**

The web GUI does not support media containers. It must not be used if there are media containers configured and used for media storage as well as media movement.

---

To install the Storage Media Operations Manager, run setup.exe and follow the installation instructions. The installation requires Microsoft® Installer 2.0 to be installed on the system. If there's an older version present, the installation routine will automatically update the system to version 2.0. The system will have to rebooted upon completion of the installation. It is not necessary to perform a separate installation on the Storage Media Operations Server because the Storage Media Operations Manager gets installed there automatically.

## Data Protection Integration Agent

The Data Protection Integration Agent (DP IA) is the interface that enables Media Operations to automatically retrieve data on media, systems, devices, etc. from backup applications. Currently, the agent support Data Protector and Omniback II.

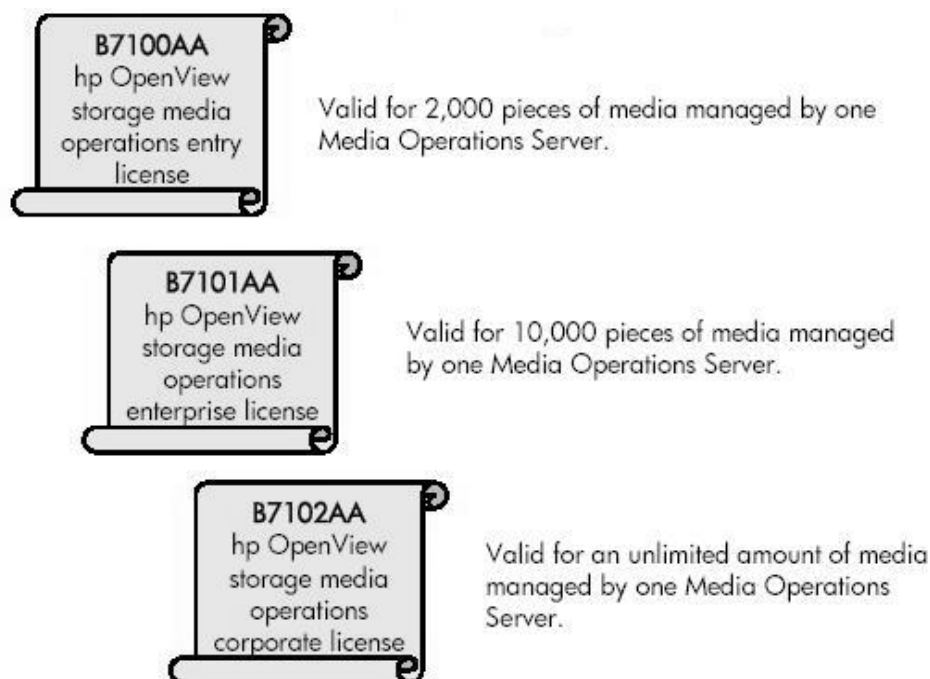
The Data Protection Integration Agent is supported to run on the following operating systems:

- HP-UX 11.x
- Sun Solaris 7&8
- Microsoft Windows 2000 and Windows NT4 SP6

It integrates with Data Protector and Omniback II 3.5 and 4.1. It supports the cell managers on all operating systems supported by the respective Omniback or Data Protector version. If the cell manager runs on an OS that the DP IA cannot run on, the integration maybe done remotely by installing the DP IA to a separate system, that runs an OS supported by the DP IA.



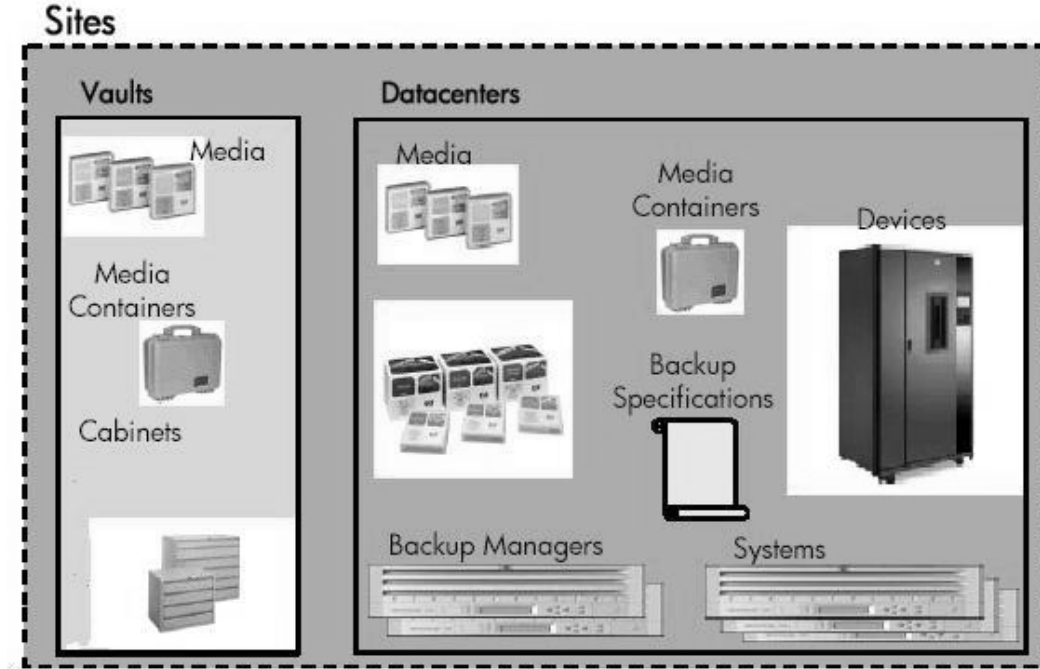
## Licensing



Upon installation, Storage Media Operations will run in a demo mode for up to 60 days. All functionality will be available for an unlimited amount of media. This is comparable to the 60-day Instant-On license of Data Protector and Omniback II. In order to unlock the functionality permanently, licenses need to be purchased and installed. The Storage Media Operations licenses are valid for a certain amount of media. The following licenses are available for Storage Media Operations:

- **B7100AA hp OpenView storage media operations entry license** — This license is valid for 2,000 pieces of media on one Media Operations Server. Multiple licenses can be installed to manage more than 2,000 pieces. It may be installed together with enterprise licenses (see below). One entry license must not be installed to more than one Media Operations Server at a time.
- **B7101AA hp OpenView storage media operations enterprise license** — This license is valid for 10,000 pieces of media on one Storage Media Operations Server. Multiple licenses can be installed to manage more than 10,000 pieces. It may be installed together with enterprise licenses. One enterprise license must not be installed to more than one Media Operations Server at a time.
- **B7102AA hp OpenView storage media operations corporate license** — This license is valid for an unlimited amount of media on one Media Operations Server. One corporate license must not be installed to more than one Media Operations Server at a time.

## Objects and Organizational Units



In order to utilize Storage Media Operations' capabilities, it is important to understand what objects exist in the product, what their purpose is and how they are organized.

- **Sites** — Sites are the biggest organizational units in Storage Media Operations and usually match a company's sites. They can contain several vaults and/or data centers. One Storage Media Operations server can manage many sites but requires at least one. Sites must be created and configured manually. There are no strict rules as to what a site is. From a media management and vaulting perspective, it makes absolute sense to compose a site of data centers and vaults, which are physically close together.
- **Vaults** — A vault is a place where media can be safely stored or archived apart from the systems and backup devices. In Storage Media Operations, a vault belongs to a site. Vaults must be created and configured manually. Vaults are onsite or offsite vaults. A vault is determined to be onsite or offsite from the perspectives of the media. Per definition, a vault will be an onsite vault if it's at the site that owns the medium. This site is also referred to as the home site of the medium. Vaults in all other sites are offsite vaults.
- **Data Centers** — Data centers are the places where the actual live data resides. They contain servers and storage devices including both disk and removable media storage. The data center is the place where data gets backed from disk to tape. In Storage Media Operations, a data center belongs to a site and can be organized in a grid-like layout in order to optimize vaulting jobs. Data centers must be created and configured manually.
- **Media** — Media are all supported removable storage media, such as tapes. They can either be added manually or automatically retrieved from a backup

application, if a Data Protection Integration Agent is available. The amount of media managed by Media Operations can be expected to range somewhere between thousands and hundreds of thousands. Because of this, media objects should be retrieved thru the Data Protection Integration Agent. Media can be in a data center, vault or in a media container. They belong to a media pool.

- **Media Pools** — Media pools are logical organization units that contain media. They hold media that are of the same type and have similar properties. Media pools do not really pertain to a data center or vault because their media can be in any data center or vault managed by any Storage Media Operations server. However, media pools have a Home Site to indicate where they were originally generated. Media pools may be created manually or automatically retrieved from a backup application if the DP IA is available. Media pools can be organized in media pool groups to simplify administration. For example, vaulting policies can be assigned to media pool groups instead of individual media pools.
- **Media Containers** — Media containers can be used to transport media between data centers, vaults and sites. They simplify the task of vaulting. All media entered into a media container are checked in and out all together rather than individually. Because media containers are intended for media transports only, they will only reference the site they are currently at but never be assigned to a vault or get a grid location in a data center. Unlike other objects in Storage Media Operations, media containers can only be created manually. Omniback II or Data Protector does not maintain media containers.
- **Cabinets** — Cabinets are media cabinets, racks, shelves or safes that are used to store media in vaults where they are separately protected from the live data in the data centers. In Media Operations, cabinets only exist in vaults. They need to be created manually.
- **Devices** — Devices are backup and restore devices that are used in data centers to backup or restore data. Storage Media Operations needs to know these devices in order to be able to create pick lists for vaulting jobs or for premount jobs. Also, Storage Media Operations can trigger a backup application such as Data Protector to automatically eject tapes from library devices. Devices can be created manually or retrieved automatically from backup applications if the DP IA is available.
- **Systems** — Systems are servers that are backed up to media in a data center. They may be created manually or retrieved automatically from backup applications if the DP IA is available.
- **Backup Managers** — Backup managers are special systems that run backup applications. For example, a Data Protector or Omniback II Cell Manager can be configured as a backup manager. Using the Data Protection Integration Agent, Storage Media Operations can automatically retrieve data and requests from backup manager. Backup managers pertain to a data center of a site and must be created manually.
- **Backup Specifications** — Backup specifications are backup jobs that are kept on backup managers and are used for backups that reoccur on a regular,

scheduled basis. Storage Media Operations can assign default vaulting policies to backup specifications. Media produced by backups will be treated according to the vaulting policies assigned to the backup specifications. Backup specifications can be created manually or retrieved automatically from backup applications if the DP IA is available.

## Daily Operations

There are four kinds of operations that Media Operations does to move media between scratch bins, devices and vaults: scratch-bin maintenance, premount jobs, vaulting and Check-Out Requests (COR).

Scratch-bins contain initialized or expired media that are available for backup operations. These media get mounted to backup devices during premount jobs. Scratch-bin maintenance includes moving expired media from the onsite and offsite vaults to the scratch-bins as well as prompting operators to initialize blank media for specific media pools.

Premount jobs are performed to load scratch media from the scratch-bins to backup devices. The list of media to be loaded is automatically calculated based on the media usage of individual backup specifications and the amount of scratch media currently mounted to the device. Premount jobs automatically include dismount listing, which determine all media that need to be dismounted from backup devices for vaulting operations.

Vaulting jobs move protected media, which were used during backup operations, from backup device to onsite or offsite vaults or from onsite vaults to offsite vaults.

Check-Out Requests are performed to retrieve media from onsite or offsite vaults for restore or audit purposes.

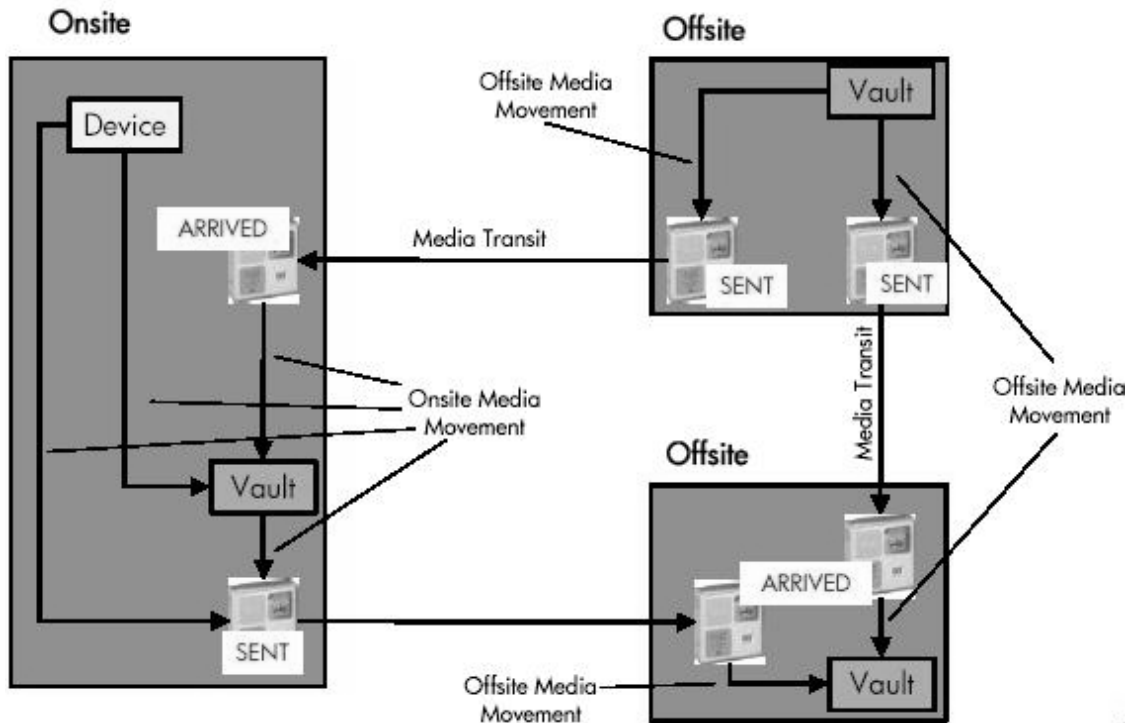
## Premount Jobs

Premount jobs ensure that there's sufficient media in all of the backup devices managed by Media Operations. Ideally, they run before the backups of the day respectively the next 24 hours. The amount of media needed is derived from the historic media usage of the backup specifications. Also derived from the backup specifications are backup devices that are affected. Only backup devices, which are referenced by backup specifications associated with a premount job, will be serviced.

At the same time, premount jobs also consider media that need to be unloaded respectively dismounted from devices because they're supposed to be vaulted. Premount jobs are associated with one site. It is not possible to run one premount job on many sites. It is very likely to have many premount jobs for one site though.

Media Operations can of course not take the media out of the devices or carry scratch media from the scratch bins to the devices and load them accordingly. These parts of the premount jobs are left to the tape operator. Media Operation only creates de-mount lists of media that are to be taken out of the devices and premount lists of scratch media that are supposed to be loaded to the devices. The list of media assignments can be organized in a walk-through order to optimize ways between the backup devices. If integrated into Omniback II or Data Protector via the DP IA, Media Operations can trigger the eject of media into mail slots of library devices or have them entered into the library from the mail slots.

## Vaulting Jobs



The concept of vaulting media, i.e. moving media from one location to another for different purposes, is to divide all media movements into separate vaulting jobs. There is a separate vaulting job assigned to the following objects:

- Onsite vault
- Data center
- Offsite vault destination (OUTGOING)
- Offsite vault source (INCOMING)

Offsite vaults in this context are offsite vendors and accounts, which can be configured under *Offsite Vendors & Accounts*. These jobs appear in the vaulting job list. They run according to their own individual schedule. There's a *Default Vaulting Job Start* time though, which is normally 8 AM but can be configured differently. Scheduled vaulting jobs will be inactive if they don't currently run.

Besides the scheduled vaulting jobs, there are interactive vaulting jobs. These are performed in order to move media between sites that are either on the same or on separate Storage Media Operations Servers.

### Vaulting Jobs of Onsite Vaults

Media, which are supposed to be moved from onsite devices to onsite vaults, get assigned to the vaulting job of an onsite vault. An onsite vault in this context is a vault that is on the same site as the device the medium is currently loaded to. The vaulting policy associated with the medium will not specifically tell what vault to move the medium to. The vault will be chosen dynamically and the medium will be processed by the vaulting job of this vault. Each data center that the media

come from gets its own sub-job in the vaulting job. The media in the sub-jobs are sorted by the device they come from so they get listed in the same order as on the premount list.

### **Vaulting Jobs of Data Centers**

Media, which are supposed to be moved back from onsite vaults to onsite devices, get assigned to the vaulting job of the data center the device resides in. Each vault where media come from will get its own sub-job in the vaulting job and the media are ordered according to the onsite vaults.

In case media reside in media containers in the vault, the media will also display the container ID in the vaulting job so the user who picks up the tape knows where to find it.

### **Vaulting Jobs of Offsite Vault Destinations (OUTGOING)**

Media, which are supposed to be moved from an onsite device or an onsite vault to an offsite location, get assigned to an OUTGOING vaulting job of an offsite vault. The offsite vault is represented by the respective offsite vendor and account. Where exactly in the offsite location the media will go to is not defined by this vaulting job. The exact location will be defined dynamically.

For media currently stored in media containers, the container ID will replace the media on the vaulting job. All media in a media container are treated as one single medium in vaulting jobs.

If the offsite location does not have an electronic media-tracking interface, i.e. it is not managed by another Storage Media Operations Server or any other media tracking software that Storage Media Operations can exchange data with, the vaulting job will also display the expected return date for all media that have are associated with a media vaulting policy that returns the media from the offsite location.

### **Vaulting Jobs of Offsite Vault Sources (INCOMING)**

Media, which are supposed to be moved from offsite locations to onsite devices or vaults or from offsite locations to other offsite locations, get assigned to the INCOMING vaulting job of an offsite vault. Its offsite vendor & account represent the offsite vault in this context. When the media is moved to an onsite vault, the vault is dynamically assigned. The destination vault is not part of the media vaulting policy. The media in the vaulting job are ordered first by the onsite devices and second by the onsite vaults.

When moving a medium inside a media container from an offsite location to another offsite location, the container ID will replace the medium on the vaulting job. All media in a container are moved as one single piece of media.

When moving a medium inside a media container from an offsite location to an onsite device or vault, the vaulting job will display the container ID alongside the medium so the user knows where to find the medium once the shipment arrives from the offsite location.

## Interactive Vaulting Jobs

Interactive vaulting jobs are used to exchange media between two separate sites, which are either on the same or on separate Storage Media Operations Servers. The interactive vaulting jobs are triggered by external requests and there are two basic types of those. The external check-in request is used when media is received from another site and supposed to be checked into one of the onsite vaults. The external check-out request is used when media is requested to be returned from an onsite vault to another location with the request coming from the offsite.

The external requests are validated by the username and password of the external vaulting user. A valid external check-in request will trigger the creation of an interactive check-in vaulting job on the site that receives the request. A valid external check-out request will trigger the creation of an interactive checkout vaulting job.

A Storage Media Operations Server can create an external vaulting check-in request internally, when media are to be moved between two sites controlled by this server. In this case, an already running vaulting job will trigger an external check-in request in the destination site with all media that need to be moved to the destination site from one site. If there are multiple sites on the server that are shipping media to the destination site, they will all trigger separate external check-in requests.

A Storage Media Operations Server can also receive an external check-in request from another Storage Media Operations Server. These kinds of requests are submitted via the HTTP-based external vaulting interface. The communication between the two servers is supported by the DMCOMMS. Such a check-in request will automatically create a temporary interactive vaulting job, which is used to validate the external vaulting username and password, and then immediately start the vaulting job. The parameters of the external check-in request include the username and password as well as the media and media containers that are supposed to be stored on the destination site.

Interactive check-in vaulting jobs can also be manually triggered. Storage Media Operations will prompt the user to enter the location and name of a so-called external check-in file, which has been created by a check-out request or vaulting job of another Storage Media Operations Server. The check-in file needs to contain the same XML data as a normal external check-in request received via the HTTP-based vaulting interface. If the check-in file contains a valid request, the interactive check-in vaulting job will be created. This method should only be used when there's no HTTP-based network communication possible.

## Assigning Media to Onsite Vaults

When Storage Media Operations generates the vaulting media lists of a site and assigns the media to vaulting jobs, some media movements will have "Onsite vault" as a destination but no specific vault, cabinet or even slot will be assigned. The allocation of vaults is done dynamically based on reservations of slots or slot ranges for certain media pools, the priorities of the vaults and of course the remaining slot capacity. Also, Storage Media Operations will always assign media




of moved from onsite devices to onsite vaults first and media coming from offsite locations second. Offsite locations in this context are offsite vendors and accounts. In all cases, reserved slots in a vault will have top priority when the assignments are made.

First, Storage Media Operations will attempt to assign media of media pools, which have reserved slots in the vaults of the site. It will start with reserved slots in the vault of the highest priority. Priorities go from 1 to 9 with 1 being the highest and 9 being the lowest priority. If there's insufficient capacity in the reserved slots of the first vault, it will allocate reserved slots in the vault with the next lower priority and so on. While media are being assigned, the number of reserved slots is reduced by the amount of assigned media. Note that Storage Media Operations reduces the number when it does the assignments and not when the media are actually put into the vault. This way, concurrently running vaulting jobs cannot assign the same slot to different pieces of media. The media in the vaulting lists for reserved slots are either sorted by the devices or alphabetically by the offsite location they came from.

Once all reserved slots have been assigned with the media of the specific media pools from onsite and offsite locations, Storage Media Operations will distribute the remaining media across all remaining free slots, which are not reserved by any media pool. It will again start with the vault of the highest priority and proceed with the vault of the next lower priority if needed. Media are sorted on the vaulting list by their devices or alphabetically by the offsite locations they came from.

If there's insufficient capacity in the vault to hold all media of a vaulting job, the vaulting job will not complete successfully. It will display an alert and all media, which could not be assigned to a vault will be put on the exception list. Once more vault capacity has been added, the user can use manual verification and movement function of the exception list to assign the media to vaults. The location of these media will be set to "Holding Bin" rather than "Unknown".

## Check-Out Requests

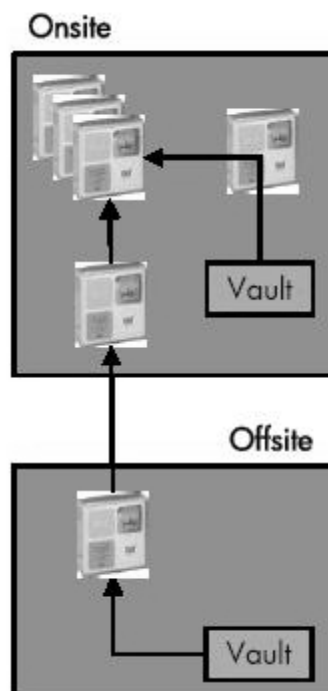


Check-out requests (COR) are performed to retrieve media from onsite or offsite vaults for data recovery, vaulting jobs or scratch bin maintenance jobs. Each COR contains a list of media to be checked-out as well as the source location and destination location of the media.

There are manual and interactive check-out requests. Manual CORs are created by users using the Storage Media Operations Manager or the web based GUI.

Interactive CORs are based on vaulting requirements. Vaulting jobs, e.g. offsite-to-offsite media movements, can request the check-out of media from another site thus creating an interactive check-out request at the current site of the media.

## Scratch-Bin Maintenance



Scratch bins contain scratch media. Scratch media are media, which are available for backup operations because their protection has expired or because they were initialized by a backup application but not yet used for backup. The purpose of scratch bin maintenance is to retrieve these media from onsite vaults or offsite vaults.

There's one scratch bin for each media pool managed by Storage Media Operations. As soon as the media pool is created either manually or automatically, the equivalent scratch bin is created as well. Unlike all other objects and object containers in Storage Media Operations, scratch bins are abstract. They can't be configured, browsed or viewed from the Storage Media Operations Manager. The location of the media will simply be changed to "Scratch bin" when they are removed from vaults. Storage Media Operations does not control where the scratch bin physically is or where in the scratch bin an individual piece of media resides.

## Learning Check

1. List four features of Storage Media Operations.

.....

.....

.....

.....

2. What is the function of the Storage Media Operations Server?

.....

.....

.....

3. What is the function of the Storage Media Operations Client?

.....

.....

.....

4. What is the function of the Integration Agent?

.....

.....

.....

5. What is the function of sites?

.....

.....

.....

6. What is the function of vaults?

.....

.....

.....

7. What is the function of cabinets?

.....

.....

.....

8. What is the function of backup managers?

.....

.....

.....  
9. What is the function of premount jobs?

.....  
.....  
.....

10. What is the function of vaulting jobs?

.....  
.....  
.....

11. What is the function of scratch-bin maintenance?

.....  
.....  
.....

12. What is the function of check-out requests?

.....  
.....  
.....